Closed Loop

Threads from CSGnet

Volume 1, 2, & 3

Closed Loop

Journal of Living Control Systems

Volume 4

Contents

Closed Loop was digitized by Dag Forssell in 2001 and 2009. Addresses and phone numbers have not been updated. Most are obsolete.

Posting on the Internet makes all the issues of *Closed Loop* available free to anyone interested.

Posted at www.pctresources.com

Table of Contents

Closed Loop

Threads from CSGNet

This reproduction of *Closed Loop* was created by Dag Forssell in 2001. Addresses and phone numbers have not been updated. Most are obsolete.

Posted at www.pctresources.com

Proofread as of March 6, 2001

As announced in the last CSG newsletter, CSGNet, an electronic mail network for individuals with control theory interests, was begun in August 1990. CSGNet now has about 40 participants in the U.S.A., Canada, Europe, and Australia, and it continues to grow steadily.

Since its beginning, CSGNet has been a remarkably active forum for the discussion of control theory. CSGNet has turned out to be an exciting and convenient medium for sharing ideas, asking questions, and learning more about control theory, its implications, and its problems. Among the more active CSGNet participants are Bill Powers, Gary Cziko, Rick Marken, Wayne Hershberger, Tom Bourbon, Chuck Tucker, David McCord, Dennis Delprato, and Hugh Petrie.

A serious shortcoming is that to date there are no clinical participants. This is most likely because most CSG clinicians are not affiliated with university or research institutions having access to either the Internet or the Bitnet electronic mail networks. Nonetheless, at least one commercial computer communications service, CompuServe, offers access to Internet (and therefore to CSGNet) for its subscribers. This means that independent researchers and clinicians who do not have institutional access to Internet or Bitnet can still participate in CSGNet. They just need a computer, modem, telecommunications software, telephone line, and money to pay for the connect time.

As this is written (January 1991), CompuServe's connect time charges are $6.00 per hour for 300-baud service and $12.50 for 1200- and 2400-baud service (call toll-free 1-800-848-8990 for up-to-date information). To obtain access via a local telephone number in most American and Canadian cities, Telenet is probably the least expensive telecommunications service to link to CompuServe. Telenet currently charges $12.00 per hour during prime time but only $2.00 during non-prime time (evenings and weekends). To make the most use of CSGNet at minimum cost via a commercial service, participants should connect to the service only for uploading and downloading mail. In this way, messages can be composed and read off-line. Two lower-cost services that do not now have access to Internet but might have access by the time you read this are GEnie and Prodigy. GEnie currently charges only $4.50 per month for unlimited access to its basic services, including electronic mail. There are no sign-up or connect time charges for participation on CSGNet itself.

CSGNet's Bitnet address is "`CSG-L@UIUCVMD`" (use no quotes in this and the following addresses); "`CSG-L@VMD.CSO.UIUC.EDU`" is the address for Internet. The messages sent to CSGNet via these addresses will be forwarded automatically to all participants. Use the address "`>INTERNET:CSG-L@VMD.CSO.UIUC.VMD`" to reach CSGNet via CompuServe. To become a CSGNet participant, initially send a note to the network manager, Gary Cziko, at "`CZIKO@UIUCVMD`" (Bitnet) or at "`G-CZIKO@UIUC.EDU`" (Internet).

Gary Cziko
217-333-4382

Last October, when CSGNet was just taking off, Bill Powers and Tom Bourbon and I talked about publishing a sort of patchwork newsletter stitched together from the Net's conversational "threads." I agreed to consider the feasibility of such an undertaking and, if interest seemed high, to start it off with the understanding that Tom and others would provide assistance or take over as they were able. I've been impressed by the highly creative and substantive dialogue on the Net, and I suppose that its quality will continue to flourish as more participate in the discussions. At least to date, there has been plenty of material worthy of preserving and disseminating in a "digest," which would also allow Net participants—at the behest of the editor—to clarify and expand their comments in light of reactions to them by other Netters.

The question, of course, is whether anybody else is excited by the possibility of a CSGNet digest. The following "threads" from the Net will give an idea of what can be expected in a digest, except that for this "sample issue" of the digest, I didn't ask participants to elaborate on their original statements. Please let me know if you think the project is worth pursuing to the extent of at least one full-size issue. Would you pay $10.00 per year for two issues of *Closed Loop*, each about 100 pages (like these) long? Would any of your non-CSG colleagues and / or local libraries be willing to pay $20.00 per year? I'd appreciate any and all comments and suggestions.

Greg Williams
606-332-7606

The Uses of Control Theory

Rick Marken: Many people have the idea that the true test of the value of a theory is whether it is "useful." This seems to be particularly true in the field of psychology. One unquestionable reason for the popularity of behaviorism is its apparent usefulness: it tells you how to cure "behavior problems;" raise children, manage people, etc. I think a case can be made for the proposition that cognitive psychology (and its variants) really came into its own when it learned how to sell itself as "useful." Thus, the popularity of Al (with helpful expert systems), human-computer interface engineering (my own field), neural nets, fuzzy logic (the Japanese use it in washing machines!?!), etc., all of which are related to cognitive psychological theorizing. Even Freudian, Jungian, and other "clinical" theories are popular because they promise to show you the source of your own problems: they claim to be useful.

I am often asked, when I present control-theory ideas, "so, how can I use this; what will the theory buy me?" I don't think that I happen to be running into an unusually utilitarian group of people. I think all people look at ideas in terms of what those ideas can do for them: after all, people want to be able to control things better; they are control systems. The success of science in general (and of scientific theories in particular) is typically presented in terms of "look what science has made it possible for us to do (control)." Science is seen as the handmaid of control; not as a window on understanding.

I submit that people's interest in "usefulness" puts control theory at a huge disadvantage in the public eye. Things that are useful help us control. But control works best when practiced on objects that are not themselves trying to control. Control theory tells us that people are trying to control. Unfortunately, people have the nasty habit of mistaking "other people" for the kind of objects that can be controlled. One of the main goals of control theory (as I see it) is to teach people that other people are not that kind of object (the kind that is more familiar to physicists). In fact, control theory suggests that efforts to treat people as though they were controllable objects are likely to lead to conflict rather than success. When there is conflict, there is no control on either side.

I don't think that the message of control theory is "just leave people alone and everything will be all right." But the message is definitely not "if you understand control theory you can get people to behave just the way you want." Many of the people who have asked me about the uses of control theory have definite goals regarding how they want people to behave. These people tend to ignore a theory if it doesn't say "in order to get behavior Y you do behavior X." It is difficult to convince them that, in the long run, they will be able to achieve their goals more successfully if they are more selective about what they try to control (non-living systems) and what they try to cooperate with (living systems).

So, what do you think? What is the use of control theory? How would you communicate its usefulness to, say, an experimental psychologist, the manager of a business, a plain old ordinary person?

Gary Cziko: I think control theory can be very useful for education, management, and clinicians (as demonstrated at our meeting), but there is also a scary side as well.

If control theory tells us that attempts to control other people using "peaceful coercion" ultimately lead to conflict and violence, then why not start with conflict and violence from the beginning? Saddam is now controlling the oil production of Kuwait quite successfully by using force.

Rick Marken: I don't think the control theory message is that "peaceful coercion" will necessarily lead to conflict. In fact, peaceful coercion could be quite successful. Actually, it seems to work all the time. I want to eat and this company is willing to give me money so that I can. I understand that my role is to "work" for them—where "work" can mean spending my time doing something that I prefer to do less than other things. I'm willing to make this exchange—the company "controls" what I do, and I control the amount of money I get. It works because, so far, we are both willing to accept a little error—I don't get nearly as much money as I want, and they probably don't get all the work they would like to get out of me. But we're both happy.

Control theory just says that when you deal with a person, you are dealing with a control system. The result of that "dealing" depends on how you deal with the control system and what the control system's current configuration is. But it is true that if you try to control the control system "arbitrarily" (that is, without taking into account its purposes), there is a good chance of conflict. For example, if the company decides that it will only pay me if I work in a certain way, and if it's the only company in town and I have no alternative means of getting money, then there are likely to be problems if, for some reason, I don't want to work in that particular way. If the company tries to control me —meaning it will only accept seeing a particular kind of behavior on my part—and if that behavior is something I just don't want to do, then there is conflict.

Most people deal with other people as people—they act as though they understand that the other person is a control system and they show respect. You get into problems with very "purposeful" people who have to have people behaving in just a certain way—no attempts at cooperation. These people treat people as objects. When I control a hammer, I want it to move exactly as I want it to move. I don't want to compromise and say, "well, if you want to land a few millimeters closer to my thumb then it's OK with me—I understand that you have needs too." I don't say that because the hammer has no needs or wants, and I can control it perfectly—we never have conflicts. But if I act the same way with my daughter, son, or wife, I am probably looking at significant conflict.

You brought up our current crisis with Hussein. How would you analyze the situation from a control-theory perspective? Obviously, Saddam is an example of the kind of person I described above as "purposeful." He clearly wants something, and he is willing *to* engage in conflict in order to get it. I argue that conflict can never be a good solution, even for the victor, since strong control systems will prevail over weak ones in a conflict. Conflicts are most interesting and obviously debilitating when both parties are about of equal strength (or skill or whatever). But even the winner of the conflict is a loser (in the long run). It is very seductive—winning a conflict looks like successful control by the person who does win. But I argue that it is a fool's paradise. The winner then imagines that control can always be achieved by force (not true), and the loser never really goes away.

I admit that there are many instances where the havoc being wreaked by a control system is so bad for other control systems that there seems no option other than forcible conflict (Hitler comes to mind, slave owners, and possibly Saddam). But can't we think of ways to avoid getting into these situations? I just can't believe that there are that many "evil" control systems running around.

David McCord: Rick, your interesting remarks suggested to me a potentially very useful aspect of control theory—conflict resolution. Conflict situations are often those in which two parties are controlling the same input quantity around different, incompatible reference levels. From a control-theory perspective, though, we know that those reference signals are merely the means to ends, outputs of control loops one level higher. Conflict resolution typically involves "going up a level" in order to identify higher-level goals of each party that are not fundamentally incompatible. While this technique is included in many different approaches *to* conflict management, control theory provides a unique understanding of why the technique works.

Chuck Tucker: I believe that the major argument for the usefulness of cybernetic control theory (or what I call Sociocybemetics) is that it is a model of how a system and process work. This is the point that we have made over and over again in our meetings—the model tells you and everyone how living systems both individually and collectively work—how they do what they do-how to fix something when it goes wrong—how to make it possible for a system to destroy itself (positive feedback)—how to suggest a system solve problems-how problems can be located—and much more. This is basically the argument for the type of model we use and it differs drastically from the types of models (theories) that are used by almost everyone in the life, social, and behavioral so-called sciences. Now perhaps we need to catalog or collect illustrations, examples, and stories about how the model has worked,

so we can have them handy to present to persons with whom we interact. I suspect that this network would be a good place to begin our list of *working examples of CCT. How* about it, mates???

"Revolutionary" Control Theory?

Chung-Chih Chen: I have read "A Manifesto for Control Theorists" by Powers. It is really very interesting. I like the idea of being a revolutionary. That is always what I want to be. But it seems to me that it's very apparent that a living system can be regarded as a (feedback) control system used in engineering. So I am very surprised that the manifesto claimed that it is a new idea for life science. I wonder why life scientists didn't discover it before.

Rick Marken: What is new, I think, is that the control of perception (which is what feedback control means in organisms) is the fundamental organizing principle of living systems. It is the fundamental organizing principle because what living systems do, at all levels of organization, from the cell to the organismic level, is carry out purposes—i.e., they control. It is the fact that organisms control, rather than what they control, that is of central importance to control theorists. Control theorists are more impressed by the fact *that* organisms control than by *what* they control. It is just as amazing that a cat controls the texture of the food it eats as it is that a person controls the network of contingencies that produce checkmate in chess. It is the organizing principle that is revolutionary: behavior is the control of perception.

AI types seem to be more impressed by the kinds of complex variables that people can control than they are by the phenomenon of control itself. This is certainly understandable. I'd rather watch my kid play chess than watch my cat chew. It is the content of control, rather than the organizing principle, that interests AI and cognitive science types, in my opinion. But Al types certainly know about control theory and some have a pretty good feel for what it is about. I was just looking over Minsky's *Society of Mind* book. He has a couple of chapters on "difference engines" which reflect a definite understanding of the purposeful nature of their behavior. (A difference engine is just a feedback control system.) He definitely understands that these systems produce goal results in the face of disturbance. But he doesn't really grasp the idea that this means that they are controlling perception, not "output." So near, yet so far.

Ultimately, AI and cognitive science seem to have concluded that control theory is just a subcomponent of a more overwhelming model of human nature. I think if you look carefully you will find that this overwhelming model is some form of external causation—where "external" could mean in the environment or in the brain / nervous system. Just like the behaviorists, AI people often get very close to the underlying principle of control (purpose) and then go off and do something else instead. Still, much of the AI / cognitive work is relevant to control theory. I see it as explorations of some program-level perceptions that people control and how they might control them. They also are more explicitly concerned with control of self-produced perceptions (those not produced via the external loop through the environment), such as memories and imaginings.

So, finally, the control revolution is really based on taking purpose seriously and understanding that purpose must be organized around the control of perception. For research purposes, this means that a large part of understanding the human mind must involve learning the nature of the perceptual variables that it controls.

Wayne Hershberger: Chung-Chih Chen, welcome! I understand your incredulity. I still do not understand how psychologists can fail to recognize the fact that animals control their environments, to the degree that they are able. Indeed, we are all puzzled; read the introduction of William T. Powers (1978) Quantitative analysis of purposive systems: Some spadework at the foundations of scientific psychology. *Psychological Review 85,* 417-435.

Tom Bourbon: When Chung-Chih Chen expressed surprise that the life sciences don't embrace control theory, Rick replied that they are close to the model, but are not quite there. I'm not sure I agree, at least not entirely. It seems to depend on which sources you read. If you look at accounts in physiology and in "neuroscience" of the control of movement via skeletal muscle, then there is little doubt that few life scientists appeal to control theory as an explanation, and that many of them reject the control model.

But the picture is quite different when the discussion shifts to internal variables. There, for several years, many physiologists have used a fairly good control-system model. Not the old, rather static models of "homeostasis," but models in which the "set point" (our "reference signal") is compared to a negative feedback signal from sensors that detect the present state of a controlled variable. And the present state of the controlled variable is a function of the output of the system (they now recognize that the external variable, not the output function, is important) plus the effects of disturbances of all sort. If you want a good representative text, try *Human Physiology,* R.F. Schmid and G. Thews (Eds), Springer Verlag (1983). There are many more. This version of a

control process is so widespread that most authors do not even cite a source—it seems to be taken for granted.

The biggest differences I see between their models and ours are these: they still refer to a comparator as a controller; the error signal is still called a command signal; and the perceptual signal is their negative feedback signal. And they do not yet realize that the perceptual signal is the variable the system really controls. Of course, we don't help the situation very much with our terminology—calling the external variable the "controlled variable," then chastising people when they do not realize that the system controls its perceptual signal, is not terribly fair on our part.

As for cognitive models...! If there were any remaining doubts that they reduce to S-R models in I-O model clothes, those doubts are over. Read "What connectionist models learn: Learning and representation in connectionist networks," S.J. Hanson & D.J. Burr, *Behavioral and Brain Sciences* 13(3),1990, 471-518. On page 473 is a re-creation of Egon Brunswick's old "lens model" in which many environmental "inputs" converge on, and are "focused by," a lens (now called "unit processing"), then there emerge many expanding outputs. The inputs are now called "fan in;" and the outputs, "fan out." I'm not sure the model explains anything more than Brunswick's did.

More important, the authors dearly identify the goals of con- nectionist modeling, as they see them: to show how the "hidden layers" in the model allow it to match outputs to inputs. There it is, clear as day, the thing we have known all along, but were criticized for saying: most "cognitive" models reduce to stimulus-response models by another name. The implications of this fact are great, given that cognitive-neuroscientific theorists declare behaviorism "dead," and their models both superior and ascendant. And a majority of them view control models as just another version of cybernetic feedback models, able to account for only a portion of "mere" sensory-motor coordination, if even that. (See especially their remarks on p. 472, right-hand column, and p. 481, right-hand column.)

Rick Marken: Tom, I think we agree more than you think. I do think that the life sciences are often close to control theory (in my perception of closeness) but, in science, a miss, even a near miss, is a mile. The reason they are close (in my perception) is because a stimulus-response model can look an awful lot like a control model. It can even behave like one! And, as you correctly point out, the model that the life sciences are ultimately trying to defend is some version of a stimulus-response model.

A stimulus-response (or response-selection) model works when you *define the stimulus in a way that implicitly includes the reference condition.*

The stimulus-response model works because behavior is occurring in a closed loop. So the model can be called a stimulus-response model, but it is really a control model with the reference signal implicitly set to zero. An excellent example of this same thing can be found in some work on computer animation that I have stumbled across. Here are some references for those who are interested: J. Williams and R. Skinner (1990) Motion Control: A notion for interactive behavioral animation control. *IEEE Computer Graphics and Animation*, May, 14-22; V. Braitenberg (1984) *Vehicles*, MIT Press; C.W. Reynolds (1987) Flocks, herds and schools. *Computer Graphics (Proceedings of SIGGRAPH)*, 21, 25-34. These folks have built little control systems that follow things or move to targets on the screen. But they don't think of them as control systems; they have sensors and effectors, so they "must be" stimulus-response devices. The devices exhibit some pretty impressive, goal-seeking activity. These researchers are sure that they are S-R devices with no inner purposes. But they are actually control systems. The sensor input does affect the effector output, but the effector output also affects the sensor input; there is a closed loop. The loop is stable because there is 1) negative feedback, because they have set up the S-R rule so that the output nulls the input, and 2) proper dynamics; there is slowing of the output effects of the sort that we use when we write our models of control. That is, the output at time t is proportional to the integral of the stimulus over time.

These "stimulus-response" devices are really control systems. They will reach their targets even when there is disturbance. But they illustrate what I mean by "close, but no cigar." These people are building control systems and watching them behave purposefully. But the researchers don't see this because they are guided by the unseen principle that behavior must be guided by external events.

One thing that might be fun is to build some of these simple organisms, but put in an explicit reference signal. This should be a variable reference signal, and, for now, it could just vary slowly and randomly. Now we have an organism that is still "S-R" in the sense that these researchers imagine, but one which is always clanging the definition of the stimulus on its own. The random changes in the reference produce "spontaneous" behavior that cannot be controlled by an external observer. But it is possible to demonstrate that the behavior is still purposeful and organized (nonrandom) by applying disturbances and seeing that they are resisted.

Tom Bourbon: Rick, as for how close the life sciences might be to an understanding of control, look again at the reference I cited as an example. There are *many* similar examples. These people are *not* talking about motor control. Instead, they are describing the control processes

for internal variables. In that field, the understanding has progressed dramatically since only a few years ago–so much so that we risk alienating a very large community when we say, flatly, that the concept of control is not understood in the life sciences. The reference signals (a.k.a. set points) are explicit, not implicit; the output is not the object of control, rather, there are *clearly identified* controlled variables (external to the control system); disturbances affect the controlled variables; and so on. Obviously, these are not the people who reviewed our manuscripts!

Look at the Schmid and Thews reference or at one of Mountcastle's more recent editions of *Medical Physiology*. What you will see in no way resembles the literature on motor control, or most of the literature on "cognitive neuroscience." I think you will be pleased: it is science, not seance.

Chung-Chih Chen: Thanks for all comments on my surprise. I am looking for the suggested papers and studying the feedback control system from the beginning. I am still not sure if control theory is a revolution. I will tell you when I understand better.

Closed Loop # 2

Threads from CSGNet

CSGNet, the electronic mail network for individuals with control-theory interests, is a lively forum for sharing ideas, asking questions, and learning more about control theory, its implications, and its problems. The following "threads" stitched together from just two of the Net's many ongoing conversations exemplify the rich interchanges among Netters.

There are no sign-up or connect time charges for participation on CSGNet. The Bitnet address is "CSG-L@UIUCVMD" (use no quotes in this and the following addresses); "CSG-L@VMD.CSO.UIUC.EDU" is the Internet address. Messages sent to CSGNet via these addresses are forwarded automatically to all participants. Via CompuServe, use the address ">INTERNET: CSG-L@VMD.CSO.UIUC.VMD" to reach the Net. Initially, you should send a note to the network manager, Gary Cziko, at "G-CZIKO@UIUC.EDU" (Internet) or at "CZIKO@UIUCVMD" (Bitnet); Gary's voice phone number is 217-333-4382.

Greg Williams
606-332-7606
May 1991

The Method of Levels and Internal Conflict

Bill Powers [in reply to comments by clinical psychologist David Goldstein]: When a client expresses confusion and frustration, I would ask him / her to tell me how that feels. "Tell me what it feels like to be confused or frustrated" (or whatever words he / she uses—you can ask him / her if those are the right terms). "Is there some feeling that goes with this in your body? Does it feel like a mental confusion? Is it like being afraid? Some other feeling? What kinds of thoughts go through your mind while this is happening? Is there something you're thinking about it right now?" And so on. Of course, when he or she has spent enough time describing these things, you try to pick up on the *next* level as it comes into view.

I probably haven't explained this very well in previous writings. What you're looking for is really being *acted out* as much as described although usually there is verbal content that goes with displaying the

attitude. What you're looking for isn't the subject-matter under discussion (the method itself, in this case). It's something that is *about* that discussion. You mustn't get suckered into joining the conversation. If the person responds to your request by saying "I don't know what you're asking me to do," you don't respond by explaining in more detail what you're asking the person to do. You ask the person to describe how it feels not to know what you're asking for. You ask for thoughts that go through the mind when that not-knowing is occurring. You ask what feelings go with it. And the person will tell you. You don't need to explain much, because what you're doing *illustrates* what you mean.

If there is anything general to learn about therapy, it has to be at the level of principles where all people are alike. All people control. All control systems resist disturbances of their controlled variables. All reference levels, nearly, are specified as part of some higher-level process. Reorganization follows attention. These are the things that control theory has to say, provisionally, about therapy. The CT therapist uses these principles to guide the exploration of a person's organization, to lead the person to see how that organization works or fails to work, and to help the person find a point of view from which effective reorganizations can be generated.

All roads, therefore, lead to the Method of Levels. The aim is to trace the hierarchy of control upward to the point where there is a control process that ought to be working but isn't. Then you have to help the person see why it isn't working right. By "working right;" we must mean "working so as to achieve still higher-level purposes." This is the only way to define a control problem that doesn't assume some one objectively right way for all people to be organized. The problem must always be that some high-order goal is not being met. The place where reorganization is needed, as far as therapy is concerned, is somewhere in the middle, between the person's highest levels and the lowest. The lower-level systems, most generally, will be working correctly if there is no organic problem. The highest-level systems are seeking the therapist's help and are on the therapist's side (or the therapist should be on their side). In the cooperative exploration known as therapy, two people learn just which processes aren't working so that higher systems can use them. And one person reorganizes them until they do work.

Not every human problem, given this understanding, is a therapy problem. Therapy will not provide missing higher-level systems. It will not cure goals that are set at the highest levels in ways that guarantee conflict with everyone else. It will not provide the things that education provides: understanding of the world, of other people; acquisition of skills. It will not provide what spiritual searches provide: the sense of harmony and beauty that makes a person feel whole, that makes life worth bothering with. Those things concern us all, and no one of us is more than a learner in those regards. The end of the line in therapy is not becoming a superbeing, but becoming an ordinary person capable of entering the struggle along with the rest of us—getting up to speed, as it were, for continuing a journey in a direction that is not clear to anyone.

My initial interest in control theory came at a time when I finally realized that it's necessary to understand how people work before you can help them (on purpose). I don't doubt that people are sometimes helped somewhat by existing psychotherapies. But the therapists don't understand why (they simply assume that it was their method that worked). Therapy takes far too long and, as far as I can see, doesn't get to the real issues giving people trouble. You can certainly cite individual cases that go against my generalization, because some individual therapists do have a knack for helping, but you can't show any case in which the result could be predicted or explained. Not by any theory that I could believe.

I think that the Method of Levels contains the essence of what is effective in psychotherapy: putting a person in a mental position from which internal conflicts can be resolved. If this is all there is to it, vast numbers of patients now undergoing psychological treatment should be released from treatment. I think that is exactly what should happen. If a person's problem is ignorance and lack of skill, that person needs education and training, not psychotherapy. If the person has organic damage, that person needs medical help (which, unfortunately, will probably not be up to the task, either). If a person's problem is a lack of respect for the opinions, feelings, and rules of others, that is a political problem and has to be worked out through negotiation, with both sides taking equal responsibility for the problem. The concept of people as autonomous control systems requires a completely new approach to human interactions, including "helping."

I don't think that a control-theoretic approach to psychotherapy can be developed unless we simply give up on all the older approaches, throw them in the trash-can along with the theories they are based on, and start over. Maybe what we come up with will turn out to resemble different aspects of different older methods. Who cares? If that happens, it will just show why other methods didn't fail *all* of the time, instead of *most* of the time. We need to get rid of the bad guesses, the fairy tales, the plausible ghost-stories, the irrelevancies that just confuse the issues of therapy, and try to pare the process down to something that works for reasons we can understand, and with some degree of reliability.

What I'm *proposing* is the following: If a person is having some sort of psychological difficulty, the normal thing to do is to reorganize and

resolve it. When a person has the same difficulty for a long time, clearly, reorganization isn't working. A "difficulty" shows up in CT as an error signal that isn't being corrected, or at least as an unreducible error signal that shows up every time the person tries to use a certain control process. If an error signal exists and no action takes place to correct the error, then something is preventing the action from taking place or having its normal effect (manipulation of lower-order reference signals). The only strictly *psychological* way for this to happen is for a second system to come into action every time the first system attempts to correct its error, the second system canceling the output of the first system. In short, conflict. Nothing can prevent an otherwise competent control system from correcting its error but a second control system that is opposed to it. If, that is, the problem is of the sort we would call psychological, and that is amenable to treatment through cognitive interactions.

Now, what could keep reorganization from working? Only the failure to bring it to bear on the systems responsible for setting up the conflict. As these systems are necessarily of a higher level than the systems in direct conflict, the locus of reorganization must be moved, somehow, to those higher systems. The Method of Levels is one way to do that. There might be others, but I don't know what they are. Some successful methods might be nonverbal. Not all reorganizations that are needed would be at cognitive levels.

Control theory suggests that the core of any psychological problem is conflict. I do not believe any other explanations that I have ever heard.

There is nothing about control theory that requires you do to anything to a patient that either you or the patient finds unacceptable. You are always there, observing and aware of effects of what you do. Effects that you are unaware of will happen no matter what you do. You can see whether the observable effects are what you hoped for, just as you can when you test any other theory. And perhaps uniquely to control theory, you can see whether the process you have attempted to put into practice has actually taken place—for example, whether your attempts to get a person to move up a level have actually resulted in the person's speaking as if from a new point of view. So you can distinguish between failure of the process and failure to get it working properly.

The point of the Method of Levels is not for the observer / listener to make clever guesses that are correct. The point is not to discover what *might* be going on in the person's head, but what is going on. The point is to draw the attention of the subject to the background processes, whatever they are. The observer gets no points for guessing correctly. The speaker is the only one who knows what the background material is, and the only one who benefits from noticing it. All the observer can do is guess. A wrong guess is just as good as a right one, if the speaker corrects it. The speaker is the ultimate authority.

Furthermore, the background thought does not have any necessary connection to the foreground subject matter. A person describing a picture might say “I see a barn,“ while the background thought is “What is this joker trying to get me to do?” The listener can only guess usefully when the speaker makes specific allusions to some background process: “I’m looking at a barn—is that what you want me to say?” (Are you wondering what I want you to say?). When I give examples, I use examples like that because I have to use something. In practice, the background thought, when revealed by the speaker, is often a total surprise.

Fortunately, when engaged in conversation, people often do make allusions or side-remarks that drop hints about the operative background thoughts. They can also drop hints by the way they speak—correcting themselves, hesitating, looking disturbed. Sometimes there will be a silence; you can ask, “What was going through your mind just then?” Or “What were you feeling just then?” That would be hard for an after-the-fact analyst to do, especially when looking at a typescript.

The real problem that demands therapy, it seems to me, is the inability of a person to take good advice, to change goals, to be more realistic, to abandon fruitless actions, even when the person knows that doing these things would help. People seek help when all the obvious things have been suggested, when they’ve tried to change their bad habits and their bad feelings, when they’ve struggled and lost. They come in when the normal processes of healing and learning have bogged down.

If a person is pursuing too many goals at the same time, it will do the person exactly no good to be told “You need to cut down on the number of goals you’re trying to achieve.” That piece of advice might be a perfectly true statement, in that if the person *could* cut down the number of goals, life would be less complex. But the person is most likely to be seeking help because the person *can’t* cut down on the number of goals. All of them seem important. And some of them can’t be abandoned because they’re holding other goals in check—to relax one side of the conflict would be to allow the other side free play, which the person has reasons for not doing. If the person *could* just take the advice and drop some of the goals, that person might need a wise friend but wouldn’t need a therapist. There’s no harm in offering good advice, but if the client has anything like a serious problem, don’t expect it to work.

Each person is unique and finds a unique way of achieving multiple goals at multiple levels. Within one person, finding an appropriate goal and defining it in terms of specific sub-goals requires achieving a balance among multiple processes of control which interact with each other, and all too easily conflict with each other. There is no way for another person to help in this multiple balancing act. It can be

done only within and by the person in question.

This is why I have always been interested in finding approaches to therapy that do not depend on giving advice or trying in some way to rebalance another person’s control systems through direct intervention. The method of levels is the only approach I have seen that acts primarily to facilitate natural processes of reorganization without attempting to direct their effects. This method is noncommittal about what is actually wrong in the person and what the person needs to change, and it does not attempt to make the change for the person.

Tom Bourbon: I am not a therapist, but I suspect that much of what happens in education is similar to some of what happens in therapy. What I usually find is that the problems a student reports to me are not the major problem, rather, they are what catches the student’s attention when viewed from another level. And what a teacher does certainly is not to use magic words that go directly to the symptoms, but to encourage the student to locate the real problem, then find a way to deal with it. (A not-infrequent solution is to leave school, which is what the student wanted to do all along.)

Bill Powers: I’m about to overrun my limited area of expertise concerning therapy; my last post verged on pontificating (I hope it only verged).

I believe I said that the conflict explanation (and the Method of Levels) applies primarily in therapy based on “cognitive interaction”—talking therapy. Of course with enough of that sort of hedging, this amounts to saying that the conflict explanation always works best except where some other explanation works better.

On the other hand, “mental retardation” and “schizophrenia” aren’t explanations. They are names for fuzzily-defined constellations of symptoms, and I don’t see any reason, a priori, to reject the idea that such symptoms could arise from severe conflict. They could *also* arise from physiological causes, but unless you’re a medical person who will believe only physiological explanations, there’s no way to decide on the basis of symptoms whether the problem has a physical or a psychological origin.

Even finding that a drug treatment affects the symptoms does not prove a physiological origin, because generally the “psychoactive” drugs used affect functions of the brain that are also affected by normal brain activity. If, for example, there is a dearth of dopamine in some part of the brain, this is because the normal sources of dopamine, a neurotransmitter, are not active. You can either supply dopamine artificially, treating the symptom, or find out why the normal sources have dried up, treating the cause. The cause could be a physical malfunction in the neurons themselves, or it could be lack of activation from other systems that normally send signals to those neurons. In the latter case, I see no reason why the explanation could not turn out to be conflict that is canceling the normal output of a control system somewhere else.

As to the other kind of symptom, people are often labeled mentally retarded when something else entirely is wrong with them—even diabetes. Of course, while they suffer from the diabetes or other condition, they *are* mentally retarded. That says nothing about what is causing the retardation. Severe conflict about learning or reasoning could easily result in retardation. Symptoms are just that: symptoms. They do not by themselves give you any clues as to causes.

It is, I think, vitally important to consider the psychological explanation in all cases, because misdiagnosis can lead to giving palliative treatments only, and can doom a person to a lifetime of unnecessary dysfunction and even misery. If conflict therapy could lift the internal suppression that would allow a "retarded" person to begin functioning normally, the advantages over the normal treatment of retardates would be obvious and enormous. Overlooking that bet would be just as serious a mistake as trying to use psychological methods to cure AIDS.

There might be other ways than the Method of Levels that will direct reorganization to work where it is needed. I just don't know of any other way. As to the direct assault on symptoms, I'll repeat myself. If it works, then there was really no serious psychological problem to begin with. It ceases to work when you run into a conflict. Then the person says, "Oh, yes, that's a good idea," but is unable to do anything with it. Something else is saying it's a bad idea. If you now push to get your advice taken, you will just arouse the other side of the conflict more. I have no objection to solving a person's easily solved problems by giving advice, getting them to try plans, and so on. But when those are all taken care of, either the person goes away satisfied, or as Portnoy's analyst said, 'Now we begin."

Despite my occasional diatribes about the general ineffectiveness of psychotherapy, I recognize that some psychotherapists do help people (*some* psychotherapists help them a lot), and that control theory, as imagined by an engineer to apply, will be only marginally helpful until the theory is translated into practice by the people actually doing the work. And I don't ever forget that the practical application might well result in information that says the theory needs revision. I think that simply understanding behavior as control gives the psychotherapist a new place to stand from which to view the therapeutic interaction. I don't really *need* to offer any suggestions about how to use control theory in this context, other than those having to do with the basic understanding of control theory itself.

This will not discourage me from trying to get therapists to do some reorganization of their own. It's impossible for a theoretician to suggest a new approach, such as the Method of Levels, without suggesting that the *present* methods a clinician uses could be improved upon. That, of course, is a veiled criticism, implying that the clinician isn't doing as well as he / she thought (and it can also be taken as veiled bragging by the one offering the suggestion). It's a disturbance, isn't it, to a large complex of control systems developed over many years, aimed at giving effective help to people? The effect of a disturbance is to call forth a countereffort, isn't it? If we can dispose of the criticism aspect of all this (it will help, of course, if I stop criticizing), we might be able to get somewhere with evaluating the concept of level-raising as a therapeutic tool, either to find out why it doesn't work even though it ought to, or how to apply it effectively.

I *totally* agree that the only feasible policy regarding reorganization is *hands off*. Reorganization can't be directed from outside. But here comes the theoretician with his "but." My way of applying the method of levels is pretty direct and blunt. In real therapy, it would probably get me killed, eventually. In the demos I have usually been dealing with friendly strangers, but not clients, knowing nothing in detail about them. So onlookers haven't witnessed any great empathy, haven't seen me giving advice or doing anything to help resolve any problems that might be described. I haven't done any of the things that would show the basic moves of this method embedded in a wider context of therapy with a familiar client in an atmosphere of trust. My way of using this method probably ought to be restricted to people who are in good enough shape not to blow up in my face. Fine, I don't want to be a psychotherapist. I'm content to demonstrate a principle and let others who are better qualified rework it into something of practical use.

So here's the "but." My way of applying this method essentially ignores the content of the ongoing conversation, treating every statement, every description, as nothing more than possible evidence about a higher-level point of view that's in operation, in the background. One thing that feels very strange to a victim of this process is that the questioner never really seems to make contact; it's like talking to someone who seems overly fascinated with the way your mouth moves and isn't showing normal reactions to what you're saying. In learning to do this, I have learned how not to let remarks pass as if they were just a sort of innocent accompaniment to the main theme. For the speaker, the background attitudes and thoughts are silent and hardly noticed at all, the way you don't realize that you keep looking at your watch because you're in a conversation that has to end before your plane leaves. These background processes are there in consciousness, but only a very little bit, not enough to warrant full attention. The whole point of the Method of Levels is that the *listener* can't do the same thing—realize vaguely that the background thingie is there, but let it pass without giving it full attention. The listener has to notice those glances at the watch.

I have no doubt that all good therapists use the Method of Levels in some way. But using it knowingly might work better yet.

And finally, if you're going to teach anything about control theory to a client, the first thing should be how reorganization feels and why it's OK. Just about the only thing a person can do consciously to shut off reorganization is to shift attention to something that doesn't cause the errors that result in reorganization. I think that this is basically why people get stuck: when they pay attention to an area where reorganization might do some good, the conflicts come into play and the person feels worse. I should think that the course of therapy would be smoother if a person could learn to recognize a certain kind of "feeling worse" as a sign that something is happening, at last.

The failures tend to be people who focus on problems outside themselves instead of seeing that something they are doing needs changing. When I read that, I get a sudden picture of some people I have known, including myself, who were in just that position—wanting to solve some external problem and not seeing any way to change it—and not realizing that the only effective change would be internal. If you think about this the right way, you can understand how this comes about.

Suppose someone comes in for help and says, "There's a kid next door who I am certain is being abused. I can hear yelling at night and the kid screams and sounds terrified. I lie awake at night hearing it, and I can't sleep. The cops won't do anything. I'm turning into a nervous wreck because of it. I can't stop thinking about it. Something has to be done."

Now just imagine how this person would react if you said that he needed to deal with his feelings about what is going on, and that after enough therapy he wouldn't be bothered any more. He'd say, "Why are you trying to change *me*? I need help in changing this awful situation that's getting me down. Somebody has to help that kid! Don't you think that what is going on is wrong? I don't *want* to feel better about it—I want them to stop abusing that kid. Are you telling me that I'm imagining it? Well, I'm not!"

This is what it's like to be so focused on an external problem that you're totally unaware of where you're coming from. Everything tells you that you're completely justified in needing to solve the problem, that something very bad is happening out there, and that you need to do something about it. That might be the exact truth, in that any normal person would feel the need to do something. In the background, however, there are all sorts of conflicts that keep you from thinking of an effective action, so all you know is that you feel helpless and overwhelmed by the problem and need help with *it*, not with yourself. This sort of problem is a real attention-grabber. The only thing getting reorganized is what you imagine to be going on behind the scenes, what actions you imagine taking (and immediately give up on because every one of them arouses some sort of conflict).

Ed Ford approaches problems like this head on. He says, "What are you doing about it?" And Chen, "Is it working?" This is really a version of the Method of Levels, because it brings into the picture what the person is doing in addition to what's going on out there. Once the person begins to examine what the person is doing, and evaluating it, the level from which awareness is working *has* to have changed.

I would think that the toughest cases would be those in which the person starts to go along with this shift of viewpoint, realizes that it's taking attention away from the external problem, and flatly refuses to do it. And I think that this is where the subtlety of the Method of Levels comes into play, because in refusing to do it, the person will be telling you the higher-level reason for the refusal. This is the real barrier: the reason for the refusal. I think that if you can be just insistent enough at this point you might be able to get the person up one *more* level to talk about the reasons for refusal. You simply ignore the act of refusal itself —while accepting it—and go for the real conflict.

Teaching control theory is probably a good idea for any client prepared to learn it. But I think we agree that before this teaching can even start, you have to get the person moving in some direction and out of the clutches of the "presenting problem." If you can jog the person up a few levels, maybe that will prepare the ground enough.

Ed Ford: I have read with some interest the discussions about conflict. First, I believe internal conflict is at the heart of all human problems. I have come to control theory as a reality therapist, and, happily, I have found control theory opening doors to a much more efficient way of helping clients. Not just in plan-making, but in the entire process of *helping others to deal with their own individual worlds.* All we therapists do is teach them an efficient way of dealing with their world by teaching them how to organize their systems within the framework of the control theory model.

I see symptoms as just that, symptoms. They're not problems, only evidence of problems. Unfortunately, the two are easily confused; that's because we see the symptoms. The conflict is all internal. All symptoms do is give evidence of conflict. The real problem is that somewhere within a client's system there is conflict, or a lack of harmony. When clients come to me, they are obviously reorganizing (who would pay a private counselor when their world is in harmony, when their goals—read reference signals—are being satisfied?). When people begin to sense relief from the pressure of reorganization, then they know they are getting somewhere. (Isn't the job of a therapist to teach clients how to reorganize more efficiently?) And I certainly don't believe in disturbing anyone's system (you do violence when you push on a control system, right Bill?). To find happiness or internal harmony, clients have to be taught how to deal with their world by learning the process of controlling for what they want (and not controlling for those things over which they have no control). And this can be done. In short, *the goal of the therapist is to help the client develop a belief that his/her system can be used to reduce conflict. The second part of that goal is to teach the client the skills of dealing more efficiently with his/her internal world and re-establishing and maintaining harmony within it:* Proof of the validity of the model is the use to which clients put it, and especially its effectiveness in reducing error. I see this happening not only in my clients, but also in my graduate students as they work with their clients, and more interestingly, with those who come up to me and express how much better they understand themselves after a lecture on control theory and stress.

First, I teach them how they control for input. I teach them that they deal with people and what they say according to how they're perceived, including all the various categories that go to make up that perception. When it comes to learning about the variety of reference conditions, I learned one heck of a lot from control theory. I see systems concept as where we set our values, beliefs, the way we think things ought to be. At principles level, I see this as where we establish our standards, which should reflect and be in harmony with systems concept, the highest level. At program level, we make decisions hopefully based on our standards, which are based on our values or beliefs. If I decide to have an affair with a woman (program level) and I have a value that says that's a no-no (systems concept level), then I create conflict within my system.

There must also be harmony within each level. If my job has a higher priority than my wife, and I don't find satisfaction in the application of this prioritization of goals, I will again experience a lack of harmony and begin to reorganize until a better idea presents itself (establishing my wife at a higher priority than my job).

There are two more serious sources of internal conflict. Incompatibility of goals is the most common, for example, a single parent's conflict between the responsibility towards raising his / her children and the social demands for adult companionship. More difficult are the conflicting demands of the abused woman, between her abusing spouse / boyfriend, who is perceived as the only source of love and security, and the shelter which offers safety for her (and her children, if there are any) along with a sense of worth (from being treated humanly and through

finding and maintaining a job).

The other area of serious conflict is when we want something over which we have no control. Persons come to see me, all filled with frustration, sometimes crying, but always upset (reorganizing inefficiently). After a short chat, I ask them to tell me their various goals (systems concept level) which are presently important to them. Invariably, four out of five of these goals are things over which they have no control. Examples such as "my children to get off drugs; ' "my spouse to show me more affection," "my boss to show me some appreciation," "loss of a loved one in death," and "I'm getting old and not appreciated by my children any more." Need I continue? The attempt to satisfy impossible goals is classic. The greater the intensity of desire (I guess some of you would say the stronger the electrochemical signal), the greater the misery and the more intense the reorganization.

As for problems of the mentally retarded and the schizophrenic, they certainly evidence conflict in my experience. The mentally retarded certainly have goals, rather simply defined perhaps (although obviously I can't see into their created worlds), and certainly they have a view of the world (they do recognize it and deal with it, although on a limited basis). They certainly experience frustrations, and they often work things out and evidence harmony quite a bit. No matter what the presenting problem, and no matter what the condition of the presenter (read client or patient), they all have the same kind of world. The job of the therapist is to figure out (a little reorganizing on our part) how to teach the client to use his / her system according to his / her capacity and willingness to learn.

Now the schizophrenic. I worked for two years in a hospital for the criminally insane as a consultant. I worked on the wards dealing with patients, training the staff. My perception of so-called mental illness is that it is chosen. I found that in my contact with patients, they reacted quite well to this approach. I believe that patients arrive at various choices of acting through reorganization. People, when they reorganize, don't always choose the most efficient way to deal with conflict, but they will make a choice that reduces error. It might not reduce the errors of others (a child's tantrum comes to mind), but if it reduces their error (the child gets what he / she wants), then a new method of reducing error has been learned. It might not be the best, nor bring the most satisfaction, but it works well enough to reduce error, so they use it again and again. And many people tantrum right to their grave, if need be.

In summary, Bill, don't revise the theory, it's working quite well, thank you (and I've spent 10 years learning it). Also, I have found that level raising does work. Finally, people shouldn't be listening to the therapist. That's because the therapist shouldn't be doing the talking. The job of the therapist is to question the client and listen, listen and watch the way the client is dealing with his/her world. That's the way you teach people to think (a rare experience in school these days). I question people about their world (reference levels and perceptions to you scientists) and ask them if their worlds compare favorably (perceptual error or no perceptual error). If they don't, I ask them if they want to set a reference condition for working at another way of getting what they want, and then I teach them (because now I'm perceived as a teacher) to get what they want, making sure in the process they establish measurable goals that can be easily compared with internal reference signals.

Bill Powers: Ed, good to hear from one of the people on the jury. Control theory has to make sense to non-theoreticians and practitioners outside academia if it really has something to say about human nature (although in your teaching capacities you aren't really outside academia except in spirit). I think you've demonstrated that it is teachable in a useful way, and that teaching it to clients can at least offer them a helpful framework for restructuring their lives. Even if the applied version of the theory is still subject to revision and criticism, as it stands, it probably makes more sense than the theories most people bring with them into a counseling session.

We theoreticians and academics in the CSG are grateful to Ed for his common sense and his willingness to put our abstract notions to the ultimate test: trying them out (sometimes with a degree of faith we don't deserve) in real life. Ed can't be accused of using control theory with easy cases.

"Conditioning"

Gary Cziko: I wonder if someone can help me to understand one of the building blocks of "scientific" psychology from a control theory (CT) perspective, so-called classical or respondent conditioning.

I have yet to come across a CT account of this which I can understand as well as I can understand what behaviorists call operant conditioning. I have read Wayne Hershberger's account in the *American Behavioral Scientist*, but I find the notion of anticipatory phenomena a bit troubling. I've gotten the feeling that Bill Powers doesn't like anticipation or feedforward either, but I can't quite see how classical conditioning phenomena can be handled by present time higher-order control systems.

By the way, has anyone done an experiment something like the following? Take a "conditioned" Pavlovian dog and fill its mouth with a working load of saliva before presentation of the conditioned stimulus. Does it then salivate at the bell? CT should say it doesn't.

Rick Marken: Gary, I don't think I have anything too original to say about classical conditioning and CT. I'm sure others will handle it just fine, but, I agree, the idea that prediction is going on seems unlikely to me. The organism just controls a higher-order sequence perception. No feedforward, only feedback. I think you are also right about water in the mouth reducing conditioned salivation.

Bill Powers: As I understand it, there are some responses that are unconditioned (meaning that they occur every time the unconditioned stimulus is present), and some that are conditioned (the response does not initially occur, but must be induced through an experimental manipulation).

The unconditioned stimulus can be viewed as a disturbance that tends to alter a controlled variable that is very reliably controlled by a given species. One would tend to think of such reliable control as resulting from built-in rather than learned control systems—the so-called reflexes. Dick Robertson, on the other hand, has data showing that unconditioned responses are not as reliable as advertised. But let that go.

An example of a conditioned stimulus would be a bell that rings just before the unconditioned stimulus (a puff of air on the eye) occurs. The bell alone initially is not followed by a blink. After some number of trials, the blink occurs at the bell instead of waiting for the puff. Since the response has already occurred, it's irrelevant whether the puff now also occurs. The puff can be discontinued and for a while at least the blink will occur on ringing of the bell.

The CT explanation entails making a model, which properly ought to be done in the context of a systematic experiment. First, we guess at the controlled variable. Perhaps the effect is based on a variable that would be disturbed if the blink did not occur. To understand what that variable might be, we can try converting to continuous variables. A blink in response to a puff of air is the instantaneous version of squinting in a stiff wind that blows directly into the eyes. Preventing wind from blowing directly into the eyes might ire learned as a consequence of drying of the eyeball, or of getting dust blown into the eyes. Or, since this is such a common experience, such a control system might be built in or come into operation just through maturation. Hard to guess. Now, the blink in response to a puff becomes the action of a continuous control system presented with a very brief disturbance. It responds, but a little too late to counter the puff; an instant later the puff is gone and the eyes open again.

Now we need to bring in the conditioned stimulus—the bell. At first the bell elicits no response, but the immediately following puff of air does. The system experiences the bell followed by the puff's sensory effect that occurs before the eyes can shut. This is an event (a short fixed pattern of lower-order perceptions). Presumably, the effect of the puff is still unpleasant. Reorganization takes place and the perception of this event is assigned a reference level of zero. When the output part of the system becomes organized, the error resulting from occurrence of this event (with a zero reference setting, any occurrence is an error) is routed to a lower-level system that can counteract the effect of the disturbance. In a natural setting, the person might raise a hand, turn the head away, close the eyes, or do all three. The error appears as soon as the first element of the event occurs, the bell. The resulting action of the lower-level system now prevents the puff from having any effect, so the second element of the event is prevented from happening, if the delay is long enough. Perception of the event, and thus the event-error, is reduced, but not to zero because the higher-order system can't correct for instantaneous disturbances and can't anticipate the initial component of the event, the bell.

In general, interpreting the logic of classical conditioning phenomena tells us what kind of variable and what level of control might be involved in particular cases. It's probably best to try the lowest-level variable possible first. In the case of "anticipatory" responses, I don't see any way to do this below the event level.

"Conditioning" is a circular term when used as an explanation. In fact, this term refers to the procedures carried out in a conditioning experiment. The result of the procedures is that a neutral stimulus becomes effective in eliciting behavior. This result can't be explained by attributing it to conditioning, because it is the effect of conditioning (a procedure) that is to be explained. Only by proposing a model of the behaving system can you come up with a real explanation. And doing that converts conditioning from something that the environment appears to do to the organism into a skill or capacity that the organism has. Given two organisms, one with this skill and the other without it, both subject to exactly the same conditioning procedures, only the organism with the required internal abilities will demonstrate the phenomenon, protecting itself against the disturbance.

Gary Cziko: Bill, I appreciate your control theory interpretation of classical (respondent) conditioning and can follow the argument when you talk about air puffs on the eyeball. A reference level of zero puff on eyeball makes sense.

But could you try this out for something like the startle reaction to a sudden loud sound? What good does jumping out of your chair do when someone pops a balloon right behind you? In fact, the startle reaction also includes an eyeblink. Is this just a useless side effect of some behavior which is in some way more functional? Perhaps just "priming the pump" to get the systems going for flight or fight?

Bill Powers: Gary, remember to try converting to a continuous-variable basis. If you hear a loud roaring right behind you, wouldn't you like to increase your distance from whatever it is before you bother to look? It might have teeth. Of course a *bang* is just the beginning of a roar (or whatever) and is gone as soon as it appears. So whatever action you were about to take disappears just as fast. You can't judge what a control system is for by watching it operate under unusual circumstances. Watching a system designed for continuous control but subject to an impulse-disturbance isn't going to tell you much (unless you're set up to record transfer functions). Most "reactions" of this sort occur in circumstances set up by experimenters who are thinking strictly in terms of discrete events. Bang. Jab. Flash. Puff. Jump. Twitch. There is very little of the world or its organisms that behaves that way, except in experimental psychology laboratories.

It occurs to me that I may have given the impression that stimulus-response reactions are *impossible*. That is certainly not so—just look how the nervous system is hooked up. An electric shock that you can't fend off will excite lots of sensory neurones, and that will disturb lots of circuits, which can easily result in activation of many muscles. That's an open-loop reaction to a stimulus if I ever heard of one.

But we have to ask how important in the overall picture such reactions are. Maybe we should make a list of all the interesting, important, or complex stimulus-response reactions that we can think of, so as not to slight that mode of operation. I'll start it off. Let's see—there's the patellar reflex, the pinprick reflex, the eyeblink reflex, the salivation reflex, the startle response,—uh—the vestibular reflex (although that one is really a slow control system), the sneeze, the—uh—equation-solving reflex... well, over to you.

More seriously, we should not reject the SR explanation on principle. If we do reject it, we should do so, case by case, because we can show it is a wrong or inadequate explanation of what is observed, or because we can show that it is only a special case of a more general control process. The corollary is that we shouldn't claim that any behavior is a control process unless we have some reason to think that the Test would be passed. This isn't a religion.

Wayne Hershberger. I am disappointed, Gary, that you found my control-theoretic account of respondent conditioning difficult to understand ("Control theory and learning theory," in the special issue of *ABS* edited by Rick Marken: 1990, 34, pp. 55-66). The audience I had in mind while writing that paper was the psychologist who is familiar with learning theory and conditioning phenomena, but I had supposed that what I was saying would also be clear and convincing to readers familiar with control theory. I have also been cheered by the reprint requests I continue to receive for that paper, believing that my readers understood my message. Perhaps neither assumption is warranted—what a discouraging thought.

The question about the salivating dog is appropriate, Gary, because the dog *would* salivate to the sound of the bell, even though the increased salivation would generate, rather than reduce, error.

It was virtually *always* the case that Pavlov's dog had "a working load of saliva before presentation of the conditioned stimulus." There are many salivary glands, and Pavlov postulated only one or two at a time, so that the control of the saliva level in the dog's mouth was not compromised. In classical conditioning, whatever the unconditional reflex, it is generally the case that the subject is at equilibrium or steady state when the CS is presented. However, an experiment reported by Kimble and Ost (1961) looked at the effects of a CS when presented along with a UCS (an error-generating disturbance). I cited that study in my ms, but it was cut in the editing necessary to shorten papers. I am including the unedited passage below:

Classical Conditioning

Although endogenous disturbances in the form of "noise" are generally detrimental, not all self-generated disturbances are bad for control. Disturbances may actually facilitate control by offsetting each other. For example, the slope of a roadway may offset the effects of a crosswind, leaving the driver with less of a net disturbance to offset. Since it is the net disturbance which the negative-feedback loop offsets, a reduction of the net disturbance is generally beneficial. By generating such compensatory disturbances of its own, a control system can, in principle, facilitate its control Indeed, some control systems, natural and man-made, actually employ such a mechanism. In engineering, the mechanism is generally called feedforward. In psychology it has been called classical Pavlovian conditioning.

[Endnote: Many things categorized as examples of Pavlovian conditioning today (e.g., autoshaping) have remarkably little to do with Pavlov's original work (Rescorla, 1988). However, the feedforward mechanism being discussed here appears to be part and parcel of the phenomena originally observed by Pavlov in the context of his classical conditioning paradigm, particularly his observation of the temporal contiguity of a conditioned reflex (CR) with its "reinforcing" stimulus (UCS). Although this CR-UCS contiguity is related to the CS-UCS contiguity, thought by some to be essential to Pavlovian conditioning (cf. Wasserman, 1989), the two are not the same. Feedforward involves the former type of temporal contiguity, but not necessarily the latter.]

Whenever an environmental disturbance to a controlled variable is predictable in its onset and extent, the control system may offset the environmental disturbance with a compensatory disturbance of its own, providing that it can synchronize the self-generated disturbance with the environmental one. The self-generated disturbance is a com-

ponent of output which will actually generate error unless the anticipated environmental disturbance offsets it. That is, it is a genuine, albeit self-generated, disturbance, and not merely error-actuated output. The compensatory endogenous disturbance does not reduce an extant error; rather, it co-opts, or preempts, an anticipated error. Therefore, the mechanism is called feedforward rather than feedback.

In Pavlovian psychological terms, an environmental disturbance is an "unconditional stimulus" (UCS), which automatically, or unconditionally, elicits an error-actuated compensatory output or "unconditional reflex" (UCR). Pavlov (1927) discovered that if a neutral stimulus (i.e., one that does not disturb the controlled variable in question), is predictably paired with a UCS, this neutral stimulus becomes a "conditional stimulus" (CS), which is capable of eliciting a "conditional reflex" (CR) resembling the UCR. Pavlov found that if a delay is interpolated between the CS and the UCS, the CR will be delayed, so that it occurs just before the UCS. That is, the CR is an anticipatory output which is not only synchronized with the anticipated UCS, but similar to the UCR. The CFA, therefore, acts as a self-generated compensatory disturbance.

[Endnote: In his authoritative review of classical Pavlovian conditioning 28 years ago, when behavioristic learning theory was still very much in vogue, Kimble (Hilgard & Marquis, 1961) noted that 'The views held most commonly have been that the CR is either a fractional component of the UCR, or that it is a preparation for the occurrence of the UCS" (p. 53). From the perspective of contemporary psychological control theory, it appears to be both.]

Consider again the example of steering an automobile: Let us suppose that the driver is already an expert; that is, his steering control system automatically offsets environmental disturbances (UCS) with error actuated output (UCR). Also, for simplicity of argument, let us suppose that there is no wind, and that the roadway is straight, smooth, level, and two lanes wide. Finally, suppose that our driver is going South and a convoy of large trucks is going North. As each truck passes, a pressure wave pushes the automobile toward the shoulder of the road. The skilled driver's steering control system nips each of these disturbances In the bud with error-actuated output. That is, the driver steers down the middle of the Southbound lane with the car swerving ever so slightly as each truck passes.

The scenario is set for classical conditioning to take place. The sight of each approaching truck is a CS, which is predictably paired with a UCS (pressure wave). After a few trucks have passed, we should find, according to Pavlov, that the driver begins to anticipate each exogenous disturbance (UCS) with an offsetting endogenous disturbance of his or her own (CR). To the degree that the CR cancels the effects of the pressure wave (UCS), the car will now swerve less than it had before. This, of course, makes the CR and its effects virtually invisible. In order to see the endogenous disturbance (CR) dearly, we need to occasionally remove the exogenous disturbance (UCS). That is, suppose that an occasional phantom truck appears ((S) which generates no pressure wave. Since there is no environmental disturbance to offset the endogenous disturbance (CR), the CR would manifest itself by generating error: the car would swerve toward the phantom truck. But, of course, the skilled driver would nip this endogenous disturbance in the bud with error-actuated compensatory output, just as he or she would offset any exogenous disturbance. So, the CR would appear as a brief swerve toward the center of the highway whenever a CS is presented alone (i.e., whenever a phantom truck appears). If the driver perceives the endogenous disturbance on these occasions, it will likely be mistaken for an exogenous one: the phantom truck will seem to pull or suck the car toward the center line (e.g., see Hershberger & Misceo,1983).

The key feature of classical Pavlovian conditioning is anticipation. It is as if the conditioned individual imagines the impending exogenous disturbance before it has actually occurred (as Pavlov suggested). And since an exogenous disturbance is perceived in terms of the compensatory output which it elicits (see the section below: Perceiving

Disturbances), the imagined exogenous disturbance comprises a form of covert output, which, if disinhibited (Pavlov's term), will yield overt output. To the degree that such a disinhibited imagined-disturbance (i.e., elicited output), matches the impending exogenous disturbance, the generation of real error (and the attendant UCR) is preempted (Kimble & Ost, 1961, actually noted the absence of the UCR); however, to the degree that it does not match the exogenous disturbance, the endogenous disturbance merely generates error of its own. That is, a CR is either adaptive or maladaptive depending upon whether it is followed by an appropriate UCS. Accordingly, Pavlov observed that the UCS reinforces the CR; that is, if the CS is repeatedly presented alone, the CR fades away or extinguishes, but if the UCS makes a timely appearance, the CR persists and is strengthened.

Control theory predicts that the (R which a UCS reinforces will resemble the UCR to that UCS, only insofar as that UCR is a compensatory output offsetting a disturbance to a controlled variable. For example, Pavlov often used dry food powder injected into a dog's mouth as a UCS. Although dogs routinely masticate food presented in this manner, this chewing does not constitute an offsetting reaction to a disturbance; rather, the presentation of the food powder merely enables the instrumental act of eating, which the dog proceeds to do. However, the dry food should disturb the controlled salivary equilibrium in the dog's mouth, in two ways: (a) the powder absorbs saliva, leaving the mouth drier than normal (i.e., a sensation of "wetness" which is below the normal set point or reference level), and (b) the taste of food probably elevates the set point regulating the "wetness" that is to be maintained during the act of eating. Since both of these factors would tend to generate error-driven output, the increased salivation which the UCS precipitates should be reflected in the corresponding CR. That is, in response to an effective (S, the dog should salivate, but not necessarily chew. This is in fad the case (Zener,1937).

As for your being troubled by anticipatory phenomena, I am afraid you will have to take that up with God almighty, I'm not responsible. The fact that a conditional reflex anticipates the unconditional stimulus which reinforces it is not my doing. I am just trying to understand the phenomenon.

One of the keys to understanding classical conditioning is a recognition of the fact that a control system might sense absolutely none of its disturbances. None. In other words, an unconditional stimulus (a disturbance) need not be sensed to be effective. Therefore, it is presumptuous to suppose that the occurrence of an unconditional reflex implies a prior registration of an unconditional stimulus. Further, even when a disturbance is perceived, it is presumptuous to suppose that it was perceived before the reaction. I tried to make this point in another passage that was edited from the above ms; the passage follows:

Perceiving Environmental Disturbances

Although the individual disturbances need riot be sensed to be offset, they may be monitored collectively after the fad, because they are mirrored collectively in the organism's, or mechanism's, compensatory output. For instance, the weather is mirrored in the fuel bill, and the crosswind is mirrored in the degree to which the driver crabs the front wheels to stay on the road. Hence, by monitoring output after the fad, a mechanism or organism can appreciate the magnitude of the disturbances it has been offsetting. For

example, by looking at last December's fuel bill one is reminded of the severity of the weather at that time. Or, a driver can discover the force of a steady crosswind by noting how much the car veers when it enters a tunnel (where there is no crosswind) and the car's direction of motion suddenly reveals how much the front wheels had been crabbed to offset the wind. Of course, the monitoring of output need not be delayed; the output may be monitored as it occurs. For example, before the advent of power steering, drivers could constantly "feel the force of a crosswind through the steering wheel"; that is, they could feel the muscular force required .to rotate the steering wheel so as to offset the effects of a crosswind on the car's direction of motion. Similarly, we might judge an object's weight by monitoring the force (Misceo,1983).

The notion that neural efference (output) can be monitored or sensed is not new; it is as old as experimental psychology itself. Wundt (1863) referred to sensed efference as "innervation sensations;" and von Helmholtz (1867 / 1962) spoke of the "effort of will." (For historical reviews, see Scheerer, 1987, 1989). Helmholtz argued, for example, that the perceived visual direction of a fixated object (an object imaged on the fovea, or line of sight) depends upon the intended rather than actual direction of regard, because the fixated object appears to lie in whatever egocentric direction the individual intends to look, even when the extraocular muscles are paralyzed.

This is not to say that any or every efference can be monitored by an organism. Indeed, there is some reason to believe that efference in "the final common path" (i.e., in the fibers directly innervating the muscles) might never be registered perceptually (cf. Hershberger & Misceo, 1983); for this reason, Wundt's expression, "innervation sensations," which connotes final common path, is less appropriate than Helmholtz's "effort of will." Helmholtz's volitional language, on the other hand, is very well taken, because of the two types of efference that seem actually to be monitored, one comprises neural reference signals, such as Helmholtz's "intended eye orientation." The other type comprises neural feedback signals of the type Sperry (1950) called "corollary discharges" and von Hoist and Mittelstaedt (1950) called "efference copies." Although both types of monitored efference (neural reference signals and neural feedback signals) appear to play important roles in the primate oculomotor control system (Robinson, 1975), the perceived visual direction of a fixated object appears to correspond to the individual's intended eye orientation (a neural reference signal), just as Helmholtz hypothesized over a century ago (Hershberger, 1987b). Thus, just as we tend to judge an object's weight by monitoring the force required to heft it, so we tend to "seer' fixated objects as being localized in whatever direction we intend to gaze. (In a well articulated field of view, the retina might also provide information regarding direction of gaze; Matin, et al, 1982).

Gary, just as the sensed efference comprising an unconditional reflex might, in principle, mediate perceptual impressions of the unconditional stimulus, so might the sensed efference comprising a conditioned reflex (reinforced by an impending disturbance) mediate an anticipatory perceptual impression of that impending disturbance. But this would not mean that the anticipatory perception precedes or anticipates the action, the conditional reflex in question. Rather, the reflex would precede / mediate the perception. This idea is not new with me. I believe it can be traced back to the ancient Greeks. It is also the theme of an entire book by Taylor recently mentioned on the network—although Taylor did not recognize that his viewpoint (a motor / output theory of perception) presupposed control of input.

However, I am inclined to think that some, if not most., of the efference comprising *conditional reflexes* goes unregistered; that is, the nervous system does not take conditioned reflexes into account in registering disturbances. For instance, a student (Giovani Misceo) and I had subjects judge the weight of a 4 pound cylinder dropped abruptly into their hand (they were cupping the cylinder in their hand before it was dropped). An indicator light flashed each trial for 500 ms, starting either 500 ms before or 500 ms after the cylinder dropped. The cylinder appeared to be lighter on the trials preceded by the flash. The subjects arms were not dropping as far on these 'lighter" trials because of a conditional reflexive contraction of the biceps, of which the subjects were unaware; hence, the illusion.

Generally, reference signals comprise the only type of "output" which could mediate veridical perceptions; unregistered conditional reflexes could serve to keep such reference signals "calibrated." For instance, persons wearing wedge prisms (bases out) before their eyes must converge their eyes more than normal, and, consequently, they see things as being closer (and smaller) than they are; but only initially. Very quickly, the subject begins to experience what is known as perceptual adaptation. With time, less and less of the prism-induced innervation of the medial rectus muscles is registered in the subject's perception of space. Things eventually look normal—until the prisms are removed, whereupon, things appear for a time to be more distant (and larger) than they are. Note that the polarity of the oculomotor feedback loops is *not* altered by the prisms. This adaptation is not the restoration of control per se. And, it appears to involve a type of efference which goes unregistered—whereas convergence normally registers as distance of regard.

It seems likely to me that (a) the convergence which registers as distance of regard is represented by a reference signal (in the Paramedian Pontine Reticular Formation) that controls the neural signals (or efference copy) sent to the extraocular muscles, and (b) the unsensed innervation of the medial rectus muscles is added to these signals. When one then considers the feedback loop through the retina, the unsensed innervation is a sort of endogenous disturbance offsetting the exogenous disturbance (prisms). Since the prism is a constant, the constant innervation amounts to biasing the output. However, when one wears bifocal prisms (different prism diopters), one above the other, vertical eye movements jog at the border, even after the glasses are removed. This conditioned reflex (or abrupt change in output bias) is *not* error-driven.

Gary, I am arguing that an anticipatory conditional reflex is triggered by the CS which precedes it and not by an anticipatory perception of the impending UCS. The reflex can, in principle, *cause* or mediate an anticipatory perception of the impending UCS, but there is no reason to think

that the reflex is triggered or caused by an anticipatory perception of the impending UCS. I hope this helps.

Rick Marken: Wayne, I still don't believe in feedforward or re-efference. I won't believe it until I see a working model. I think it might be worthwhile for you (and / or one of you students) to build a working model of conditioning based on your principles. Bill already has a nice working model of operant conditioning. You definitely know the most about classical conditioning; you know the phenomenon, so you should develop the model. I really think it would be worthwhile. After all, classical conditioning is one of the staples of introductory psychology courses. Why argue about how it can be explained—just make a model that can do it. And take the approach to modeling of a control theorist —that is, identify the variables involved and make sure that the model behaves in an appropriate representation of the relevant variations in the external environment.

It might be a nice way of getting us into models that control variables that are defined over longer periods of time (longer than the brief integration periods for position perception, for example).

Bill Powers: Wayne, congratulations on a perfectly beautiful piece of work. I think you have classical conditioning nailed down. In my previous post on this, I mentioned some of the factors you brought up, but you have it organized much better and more completely in addition to having the experimental evidence to back it up. Have you considered publishing a paper on just this subject in the psychological literature?

I'd be willing to accept "feedforward" if everyone could mean by that term exactly what you said. It is, of course, still evidence of feedback at a higher level. As you say, an anticipatory perception doesn't precede the response—we can still only perceive what has happened or is happening. But the effect of perceiving the right thing can be a response that anticipates the disturbance. If the response occurs either too soon or too late, it will *cause* error instead of correcting it. A higher system (or reorganization) has to adjust the timing until it's just right.

Rick Marken is working on modeling behavior at the transition or the event level. This is going to take us outside our familiar little diagrams, particularly in controlling events, because we get into timing and delays, and the output function has to do more than just send a steady signal to lower levels. Maybe Rick can work up a demonstration of classical conditioning, using your (Wayne's) analysis.

Nice work.

Gary Cziko: Wayne, thanks so much for your detailed response to my question about classical conditioning.

It's going to take me a while to understand your perspective thoroughly, but already I am beginning see more clearly where before there was just confusion.

I'll get back to you after I've had more time to read, digest, and ponder.

Rick Marken: As Bill mentioned, I am starting to work on a model that controls a higher-order variable—probably an event. I think this is what is going on in classical conditioning; the animal learns to control an event (CS-US) rather than just control a variable to which the US is a disturbance. The means of control involves salivation. The event is multisensory—sound, chewyness, swallowing—all of these things must happen in a particular "shape" for the reference level of the event to be achieved. The reference level of this event is influenced by many outputs besides salivation. The animal can be affect the "shape" of the event by varying its position relative to food and sound, varying its salivary output, varying what combination of stuff it puts in its mouth, etc. The more restrained the animal, the fewer means it has to control this event.

I think it is very important to remember that a static perceptual signal can represent the state of a time-varying event. Many of the most interesting perceptions we control are .defined by lower-level perceptions that occur over time. The notion of feedforward, I think, only becomes necessary when we think of a present-time perceptual signal as the representation of a present time event. But the perceptual signal could be the output of a "time computation window" that is 'looking for" some pattern of events that occurs over time (like physiological "motion detectors"). Past, present and future are all represented in this window simultaneously. A temporal pattern that "fits" the window's template consists of past, present and future events that were "expected" by the window. There is no need to control based on future prediction or real time computations of what "might" occur (feedforward). Just look at what "is" occurring; the current value of the perceptual signal represents the degree to which a particular temporal event is occurring.

Wayne Hershberger. Bill, thanks very much for the kind words. Control theory is the *only* theory that I know of wherein the distinction between elicited and emitted output (the reflexes and responses of classical and instrumental conditioning, respectively) is *not* gratuitous or ad hoc. In this sense, control theory is the only theory which promises a parsimonious accounting of *both* phenomena.

When I've thought about modeling conditioning, I have done so in terms of your little stick man who reaches out as if to touch visible targets. Suppose the little man could not see his finger; say he is reaching for a luminous target in the dark. The stick man, as is his wont, locates the target by orienting his head (a la an owl or preying mantis). The orientation of the head could be used to calculate a reference signal for the desired orientation of the arm, which the little man could realize while in the dark. Then, suppose the light comes on and the man uses the retinal error signal to null his pointing error (which is how he now works). Further, suppose that that visual error signal also calibrates the function relating head pointing and arm pointing. That would be a form of classical conditioning. I would be delighted if you, Rick, Tom or Greg would help me model the process.

Rick, I do not dispute the value of modeling the classical conditioning phenomenon, only who should do it. It seems to me that you could accomplish in a few days what might take me many months to do.

Bill Powers: Rick, one point Wayne was making is that in order for a UCS to exist, there must already be a control system. The unconditional stimulus disturbs the variable that is under control; hence you always get a response to it.

I think you and I agree that a likely candidate for the CS effect is to be found at the event level, where either a "CS-UCS" event or a "CS -[response]" event comes to be controlled. The CS starts out as some neutral perception initially unconnected with the CS. We have to account for how it becomes connected, and then for the actual control process that produces what looks like a conditional response to the CS.

Closed Loop # 3

Threads from CSGNet

CSGNet, the electronic mail network for individuals with control-theory interests, is a lively forum for sharing ideas, asking questions, and learning more about control theory, its implications, and its problems. The following "thread," stitched together from just one of the Net's many ongoing conversations exemplifies the rich interchanges among Netters.

There are no sign-up or connect time charges for participation on CSG-Net. The Bitnet address is "CSG-L@UIUCVMD" (use no quotes in this and the following addresses); "CSG-L@VMD.CSO.UIUC.EDU" is the Internet address. Messages sent to CSGNet via these addresses are forwarded automatically to all participants. Via CompuServe, use the address ">INTERNET: CSG-L@VMD.CSO.UIUC.VMD" to reach the Net. Initially, you should send a note to the network manager, Gary Cziko, at "G-CZIKO@UIUC.EDU" (Internet) or at "CZIKO@UIUCVMD" (Bitnet); Gary's voice phone number is 217-333-4382.

Greg Williams
606-332-7606
July 1991

Competition, Morals, Religion, and Science

Bill Powers: I woke up this morning wanting to write a nut letter or an essay. I hope the result is the latter. The trigger was hearing last night that the Gross National Product had dropped last quarter by "2.8% annualized," which I take to mean 0.7%. It occurred to me that something is drastically wrong, not with our "economy," but with our conception of it. It is simply not possible that the American people are incapable of sustaining an acceptable standard of living for themselves, through their own efforts. But the impossible seems to be occurring.

I think the villain is competition. This might seem like heresy in a free society, and perhaps it would be if competition were working the way it did in the 19th Century, when there was still a place to go when you got squeezed out. But I think that between population growth and running out of uncommitted territory and resources (because we are finally up against the fact that we live on a sphere), we are now faced with a degrees-of-freedom problem. What one person or group does to control for the things which matter disturbs what other persons or groups are controlling for, and adjustments which ease the conflict are becoming harder to find. I think that this process accelerated some time in the 1940s. I've been watching it get worse, therefore, for 50 years. It's been getting worse, of course, for much longer than that, but not as fast.

It used to be that when competition for jobs was fierce, the losers could somehow manage to find different but equivalent jobs or move to places where jobs were more available. When a company went under, another company would spring up to take its place, in an area where workers and managers could still apply their skills, but where the competition wasn't overwhelming. This worked for a long time (with ups and downs); in fact, it led to a mystique in which competition itself was lauded because it seemed to energize people to try harder. What wasn't so obvious was that this "trying harder" is a form of conflict: we are "trying harder" against each other. A lot of the energy created by competition accomplishes nothing more than cancelling out someone else's energy, leaving no net benefit for anyone. While there was still room to expand, while solutions to conflict still could be found, the energizing aspect of competition had a net positive effect. But there have always been hints that this is not the best way to organize a society: people always try to find a way to get out of the impasses caused by competition. Left to themselves, they seek the least-conflict state.

The basic idea behind social organizations like businesses or governments is that when people work together they can accomplish more for themselves than they can when working separately. This remains true as long as competition doesn't occur. Competition occurs naturally, through failures of coordination or through a desire for freedom. Failures of coordination can be corrected, because coordination is usually someone's job and people can learn to do a job better. But the desire for freedom, which is a necessity for autonomous systems like human beings, leads to competition through conflicts of goals, and no person can alter another person's goals in the same way a coordination plan can be altered. Conflict of goals can arise when individuals who are supposedly working together no longer subscribe to the same coordination plan. When that happens, either people leave the group or they begin to apply some of their efforts to resisting the efforts of others in the group. The group becomes less effective in either case.

When conflicts arise, some of the people in a group can leave to pin another group with goals they find more to their liking. As groups become larger, having wider effects on the shared environment, the potential for forming new groups diminishes, and conflict arises between groups. As that happens, the advantages of group effort over individual effort diminish. More and more of the group effort goes into cancelling the effects of other groups' efforts.

One solution is the coalescence of groups. But because these groups have disparate goals, mechanisms have to be invented to deal with conflicts without resolving them. The systematic application of group-sanctioned coercion arises: law. Law exists because of individuals who pursue goals conflicting with those of the majority, but who do not or cannot leave the group. The degree of coercion used in a society is a direct reflection of the disparities of goals in that society, and a direct indication of the degree to which that society is failing in its primary purpose of enhancing the capacity of each individual to control his or her life better. It also reflects a loss of degrees of freedom; there is no longer a way to get out of a society with which one disagrees and find a situation more to one's liking. One must therefore either change one's goals or risk coming up against massive coercion.

As conflict increases, the efforts of individuals to satisfy their own goals also increase; they must, if the goals are still to be met. But a large part of the increased effort is simply defensive; it is necessary only because someone else wants something incompatible, and it accomplishes nothing but maintenance of the status quo. Life becomes harder to sustain, but it does not get any better. Eventually, the efforts increase even further, and life gets worse. The escalation of mutually cancelling effort has a natural upper bound: we call it war. On a smaller scale, we call it violence. Violence is the all-out application of one's maximum possible force to achieve a goal, winner take all. As competition increases, so does violence increase. Violence becomes less and less a fringe phenomena seen among people whose goals are the most extremely different from the average, and it creeps in toward the center.

I think the lessons of control theory are clear: competition is not the basis for a healthy society. What a better basis would be I do not know, but I know that this one can no longer work. The next phase in human societies will be invented when the current phase loses its support. I think the understanding of human nature provided by control theory already tells us that we are not on the right track, and will help in the formulation of new approaches which do not automatically generate self-destructive violence. Nobody is going to hand us the new ideas engraved on stone tablets. We will invent them, and survive, or wait for someone else to do it, and perish.

Izhak Bar-Kana: About competition, etc., I can quote Churchill: "Democracy is the worst, except for all other alternatives." To blame the conflicts and violence on free competition is a little bit too much. Maybe a less understanding attitude toward violence could help more, especially in this country.

Chuck Tucker: What Bill has done is to present a theory of society based (of course) on control theory. Basically, I agree with his characterization, with one minor alteration: I still believe that there is far less competition for those who "make it" than most of us suppose. I still hold to the idea that there is Capitalism for the Poor and Socialism for the Rich even in the so-called Socialistic countries. There must be some way to incorporate this phenomena into the model (unless I am wrong).

Ed Ford: Bill, I question whether competition is really the problem. An article in today's local newspaper on the new U.S. moral code states that "Americans are making up their own rules and laws. We choose which laws of God we believe. There is absolutely no moral consensus in this country, as there was in the 1950s and 1960s."

When I was a child, my family used to vacation in northern Michigan. In the small town near us, there were two gas stations. They closed alternately on Sundays, allowing each a day off every other week. Closer to home, my wife is in competition with numerous poster shops and yet, when she desperately needs a poster, she calls one of her competitors, and they sell it to her at their cost.

I don't believe it is our conception of the economy, but rather our values and beliefs upon which we establish the standards for our decisions and how we deal with each other (including how we compete). The real villain is the lack of consensus on the moral principles which came from our ancestors. As I reflect on the hundreds of people I have seen in my counseling practice, few have included faith in what recovering alcoholics call a higher power when they reveal those things which are important to them. The solid Judeo-Christian values which permeated my childhood environment seem to have disappeared.

What has made the CSG such a great organization is the very thing missing where people associate and/or deal with one another. We respect each other and what each one of us has to offer. In short, our values are very much the same.

Rick Marken: Powers' theory of control not only helps me understand the (usually simple) phenomena of control which I can easily demonstrate. It also provides a framework for understanding more complex control phenomena, such as what happens when two or more control systems interact. The theory makes predictions about what we would see if people were organized as hierarchical control systems. I believe that in this spirit Bill Powers brought up the topic of social systems and the problem of competition. Bill's model makes some interesting predictions about what happens when people interact in a world where there are fewer degrees of freedom available than those needed to be varied by all systems in order to achieve their goals. One of the most interesting predictions, to me, is that it is not *physical* degrees of freedom which limit control, it is *perceptual* degrees of freedom. This means that even though the environment might provide enough degrees of freedom for n people to satisfy their goals simultaneously, it is possible for the people to perceive the environment as though it had only n-1 or fewer independent degrees of freedom. That will create conflict and competition—even though the competition is not intentional.

I think this "degrees-of-freedom" problem should be fleshed out better; but I think it is one aspect of many of the problems we appear to have in our society—the ones Bill alluded to, among others.

I agree tentatively with Izhak and Ed that the apparent value our society places on competition is not necessarily a big contributor to our problems. I think people verbally extol "competition" more than they actually practice it. I think competition—real competition, the kind where people act to deprive others in order to have for themselves—is a side effect of the degrees-of-freedom problem and the way certain people end up perceiving the world. One piece of evidence for this, I think, is that the most fierce advocates of competition will happily collude (cooperate) with the competition (and even break the law to do it) if it is to their mutual benefit.

I don't agree with Izhak's and Ed's proposed solutions to whatever problems we perceive in society. Izhak says we should tolerate violence less—but I haven't met many people who tolerate it. Violence is competition (which I believe is a side effect of the degrees-of-freedom problem) in a runaway condition. Killing all perpetrators of violence might cut down violence a bit—but, I think, because doing so would free up some degrees of freedom for the survivors. I'd rather find ways to increase the degrees of freedom available to all systems. As to Ed's solution, I don't see how it is informed by the control model. How does faith in a "higher power" improve the ability of control systems to cooperate for their mutual benefit? My experience has been that, since faiths are based on verbalisms rather than phenomena, people tend to perceive the meaning of the words slightly differently. Since many of the faithful have goals about what they want to perceive others believing, we see efforts at corrective action to bring people to the "true faith"—i.e., theirs. It took years for Western societies to free themselves from this source of conflict. Of course, we are not completely free of it. Faith might be great, individually—I can't participate because my thought processes keep getting in the way—but I think it ranks with economic ideologies as a singular cause of social problems.

In summary, I want to suggest that the value of theory is that it provides a framework for understanding complex phenomena based on a model of simpler phenomena. I think the control model is relevant to understanding complex phenomena like competition in social systems.

I think we should base a discussion of competition on the model, rather than suggesting solutions we could have picked up as easily from conservative newspaper columnists or Sunday evangelists.

Joel Judd: Through verbalisms we interpret, convince, and confabulate what we perceive (is this too far off the mark?). We do so, at least initially, according to patterns and interpretations which come to us from family, friends, and society. Narratives reveal the way we justify, explain, and account for disturbances to canonical concepts we have learned through verbal and non-verbal perceptions, and they are used to convince others that our perceptions are valid, or to go further and convince them that our interpretations of the world are the correct ones. This is one area where conflict arises among members of society.

This leads into recent comments from Ed and Rick about what we base our values on. I've withheld commenting about religion so far, as I've enjoyed comparing control theory with my own beliefs privately. I think it's OK to propose that something like control theory can provide information about societal problems and solutions to them. But I don't rule out the idea that higher-level reference levels could be adopted from a "higher authority" instead of "evolving" by trial and error, or arising by some other method. I don't see faith in a higher power as inherently problematic, nor does faith automatically translate into cooperative, loving control systems. If the faith inspires system concepts of the sort which foster peaceful coexistence and mutual cooperation, where's the harm in that? If there is only lip service being paid to the values, then we have what's commonly called "hypocrisy." Unfortunately, we do deal with higher levels in "verbalisms," so what I perceive by "love thy neighbor" might not be exactly what you perceive. However, there are ways of judging the way others perceive values, one of them being "by their fruits ye shall know them."

Another problem Rick presents is the tendency which humans have, once they feel they have the "truth," to try to convince / coerce others to perceive things the same way. This type of behavior is not all that different from fanatics of political ideology or any other ideology. It has two effects: 1) to attempt to take away another's free agency (i.e., control), and 2) to discourage one from looking to religion *at all* for answers about our existence. A related comment is that if we were to consider the possibility that there might be a worthwhile religious organization somewhere on earth, we would still have to face the fact that running it and belonging to it would be the same old imperfect control systems we find everywhere else. So one should be careful not to throw out the system because of the people who are involved in it.

I think anyone familiar with such matters would agree that faith has to be an individual matter; I can't "give" it to you any more than I can give you good manners. But I think Ed's comment gets not to the proposal that a particular *religion* would solve society's problems, but that certain *values* might. And control theory explains how and why these values might—they provide a high level of control. I don't understand a separation of the two.

Bill Powers: Rick asks: "How does faith in a 'higher power' improve the ability of control systems to cooperate for their mutual benefit?" The method of levels might have something to say on this subject. One of the unjustified postulates behind this method is that awareness usually operates as if from some particular level, which gives form to the current point of view. What you see from this point of view is the set of all perceptual signals of lower levels, with the current point of view projected into them as an attribute of this apparent external world. So if you're working from the category level, it seems that all of the relationships, events, transitions, configurations, sensations, and intensities you experience are exemplars of categories. You aren't conscious of categorizing; you just see that the categories are there, as if they existed objectively. So you're unaware of the operation of the level currently occupied by awareness. You're aware of the lower levels *through* it. This is all very metaphorical, and I don't know what it's a metaphor *for,* but pragmatically it seems to reflect experience.

Working this metaphor in the other direction, the implication is that you are also unaware of the operation of control systems of *higher* level than the "occupied" level (the level in the state we call conscious, to be slightly more operational about this). In particular, you're not aware of what is setting the reference signals at the occupied level: they are experienced simply through realizing that some perceptions are in the wrong state (you feel an effort to change them) and others are OK. You see a square with one side bowed out, and that looks *wrong*. You want to push it straight and make it into a better square—a better exemplar of squareness.

As far as consciousness is concerned, then, the definition of OK and not OK is given, not chosen. If you happen to be conscious at the logical level, the next thing which happens is a lot of reasoning about where this OK-ness is defined. Ah... it is clearly coming from a Higher Power. And that is perfectly correct: it is coming from higher levels, principles and / or system concepts, systems running automatically in the forms they had after the last reorganization—but not consciously.

Which brings us to the next sentence in Rick's comment: "My experience has been that, since faiths are based on verbalisms rather than phenomena..." Not so fast. What I've just been proposing is a phenomenon which a lot of people might have experienced throughout history. They don't have to be theoreticians to experience it, but if

they are theoreticians and don't have any constraints on their theories like science, they are free to propose any explanation they like. One of tire theories is that this advice from above about what is OK and what is not OK comes from a supernatural power outside of you (perhaps acting on your insides, but basically existing in a universe larger and more powerful than yours). Moses came down from the mountain with 10 principles engraved on tablets. Could this be a story reflecting the first conscious human experiences at the principle level? Moses' theory, of course, was that the principles were handed down from a Higher Power—which we, of course, recognize as the system-concept level. Moses heard a Voice commanding him. If the highest organized level in which your awareness can reside is the principle level, the *reference* principles will seem to come to you out of nowhere, but that doesn't stop you from trying to devise a Where.

One of the constants across religions is a belief in the power of prayer or submission to divine guidance. Instead of thinking about the content of prayers, think about the attitude behind them. One has to deliberately seek a state in which guidance is sought and accepted. In other words, the rational system (if that is the highest conscious level) has to find a logical way to accept that it is not the highest level, and so not resist any changes *in itself* which it can't explain rationally (or, more generally, can't characterize in terms of its typical mode of perception, evaluation, and action). I think this is an attitude which fosters going up a level, because it encourages you to observe the conscious level, rather than just interpreting the world through it. You begin to experience it as a level, and you can't do that *from* that level.

Of course, the next level has to exist if any of this is to happen, and it has to be functioning at least a little bit. I think that theories are proposed most flexibly when the next level up is still forming and isn't working very well. It's possible that the principle level formed in historical, or at least legendary, times. And it's possible that we are still in process of forming the highest level I have any inklings of, the system-concept level. Control theory is a system concept, surely. Where did it come from? Don't ask me: there it was. There must have been a time in the history of *Homo sapiens* when *no* system concept would have made any sense, *no* principle, *no* program. It's hard to imagine how the world would have looked when the highest level was sequence.

Human beings have been thinking about system concepts in an organized way for fewer than a few centuries, I would guess. Maybe that's an exaggeration, especially as it implies that everyone develops the next level simultaneously. But just look at the way people have been trying to model human beings since the 1940s. There has been an explosion of conjecture, with all sorts of new ideas showing up out of nowhere. There has been a quantum change in the very way we ask questions about organized systems. So it might be that our system-concept levels have just started becoming functional on a wider scale. No wonder we aren't very good at this kind of control.

And another implication is that a new level above system concepts is starting to bestir itself, poking random reference signals into the existing system-concept level, saying, "let's try this one on, or that one, or maybe that other one." What's it going to be about? There will probably come a time when people begin to get a strong sense that something is telling them to choose *particular* system concepts and avoid others: something which speaks to them from a direction they cannot comprehend any more than the first flint-knapper comprehended where the idea of sharpening stones came from. They are bound to wonder where that advice is coming from. There might be human beings alive now who wonder why I am having such a problem imagining why we pick one system concept rather than another.

So, Rick, I think there is a phenomenon, and that religious and philosophical drinkers have been trying to comprehend it. I don't agree with their theories, but I don't claim that they have been theorizing about nothing, or just verbalizing.

Rick Marken: In response to Joel Judd, and at the risk of offending everyone, let me share my own thoughts about the relationship between control theory and religion. Religion, from a control-theory perspective, is just something people do. In the model, religions are system concepts. The particular religion you follow is (according to the model) determined by the highest-level references in the model. So, in theory, there is no way to change references for religious system concepts other than by reorganization—and given the rather remarkable shifts I have seen people go through in their searches for spiritual fulfillment, random reorganization seems to be how it works. A religion is a perception derived from lower-level perceptions of principles (values, morals), programs (rituals), relationships (worship, prayer), etc. Different religions represent different combinations of these lower-order variables controlled at different reference levels.

So "being religious" is something that a 10-level hierarchical control system can do, like "being a Dodger fan" or "being a control theorist" (though don't ask me to build a working version of a religious control system this weekend—give me about 300 million years). I don't believe there is some "right set of values" for getting along in life or getting along with others any more than I believe there is a correct way to hold your right hand. There are certain values (rules) and rituals (programs) which are right if you want to perceive yourself as a "Catholic" or a "Buddhist" or a "Dodger fan," just as there is a correct way to hold your right hand if you want to say the Pledge of Allegiance correctly.

"Right" for a control system means "matching a reference signal"; the reference signals defining a particular religion are set by the system-level religion control systems.

Since nothing really sets the reference for the highest-level systems (other than reorganization due to intrinsic error), there is no experience of anything saying "be Catholic" or "be a secular humanist," so, I think, we have the experience that we take our system concepts "on faith"; they just are true; they are what we like. You might attempt to rationalize why you want to maintain a particular system concept, but ultimately, if it is really a system-level reference (and not just, for example, a program-level perception you are controlling in order to, say, "please your parents"—a principle-level perception), then there is really no more of "you" left to adjust system level references to satisfy any higher-level goal. Some system concepts (the religion ones) are sometimes thought of as more important than others (the sportsfan ones), but I'm not impressed that this is anything other than a historical accident; if things go on as they are in soccer fandom, there will soon be as many people who died (and killed) for the home team as died (and killed) for Yahweh (or Christ or Mohammed or whomever).

I don't want this to be taken as anti-religious in any way. Control theorists just want people to behave "up to specs" (in Bill's wonderful phrase)—that means to be able to control the variables they need to control without interfering with other people's ability to control what they need to control. Many people seem to get great satisfaction, inspiration, and spiritual fulfillment from faith (i.e., controlling religious system concepts), and they do it without messing up other people. That's just great. All I want to argue is that the control model should be able to explain all of human behavior, and that certainly includes behavior called religious. The control model implies nothing about the best set of principles for people to adopt in order to live best and get along best with others. There is reason to suspect that many different sets of principles will do. However, there are certain principles which will lead to problems—not because god said so (though s/he might have—s/he just never says much to me), but because they are inconsistent with the nature of human nature. So a principle allowing a person to enslave other people (a principle, incidentally, which god never saw fit to condemn—the Hebrews started enslaving people, apparently with god's blessing, shortly after they themselves were freed from slavery) might work for some time (it did), but it's not a good long-term basis for running a society, because the slaves are control systems, and they will always try to get as much control as they can. And people waste much of their productivity doing what is needed to keep the slaves slaves. It also violates the "up to specs" rule, since a slave probably has a hard time finding the set of references which eliminates intrinsic error.

I hope that control theory might be able to give a theoretical basis for understanding the best way for people to get along with each other and do the best for themselves as well. If the result of this theoretical exercise says "thou shalt have no other gods before me; then I shalt not.

This partly answers Bill's complaint about my claim that religions are based on verbalisms rather than phenomena. I agree that that claim of mine was wrong. As a matter of fact, I have had religious experiences (perceptions of religious phenomena) myself (almost always while listening to Bach, Mozart, or Beethoven). What I meant to describe (and what I will stick to) is my impression that many institutionalized religions, which take "scripture" very seriously when it comes to articulating their principles, tend to mistake the words for whatever wisdom (phenomena) those words might articulate. If you need to read a book in order to find out that it is wrong to kill and steal, then let me be the first to encourage you to keep reading that book.

Joel Judd: Rick, you didn't offend me. But I can tell that when the discussion gets rather "far afield," most people would rather "stay in the house." Talking about higher levels seems kind of ethereal I guess; not terribly scientific.

Mark Olson: Rick, you didn't offend me either. Joel, ethereal, maybe. Scientific, maybe not. Surely interesting, though! It's hard to conceptualize a systems-level analogy of a tracking task. It sure would be nice to make the ethereal scientific.

Anyway, the idea that the systems level is a recent (a few thousand years old) development is interesting. Could we develop a classification system of the animal kingdom based on the number of hierarchy levels each species possesses? My guess is that we would find a relationship between the amount of "rights" we give to a species and the number of hierarchy levels that species possesses. This idea just occurred to me, and, no, I am not particularly interested in animal rights as a topic in itself.

We shouldn't avoid this topic because it sounds unscientific—talking "unscientifically" often leads to an idea which, when tested, "revolutionizes" science. In other words, another variable means to an agreed-upon end.

Rick Marken: Mark, I agree—Bill's idea of a recent origin of the systems level is extremely interesting. I don't believe it, because I have this notion that the levels of perception are structurally imposed by the nervous system, and, thus, result from evolution rather than learning.

I read the physiological evidence as pointing in this direction; that is, there are cells, for example, in the lateral geniculate (I think) which look for patterns (configurations) rather than for other classes of perception (transitions, etc). I think the type of configuration a cell sees can be learned—a curve, rather than a line, maybe. I don't know of any evidence for this learning capability in cell receptive fields. But I think that such learning would be within a class. If the control model represents, to some extent, both the functional and structural organization of the nervous system, and if there is a systems-concept level up there in the cortex, then that's what it perceives—systems. Any kind of system, maybe, but just systems. If there were a level higher than systems, then I think it would have shown up by now. On the other hand, maybe it has always been there—it just didn't have much material to work with until now. Maybe that's why the systems level appears to show a historical development. It was always there (in *Homo sapiens*, but maybe it just didn't have much to work with early in the going.

Bill Powers: When the issue of religion, higher power, faith, and so on appeared on this net, only a couple of voices were heard against a vast silence. This is interesting. I happen to know that there are some strong opinions out there, a few favorable and many unfavorable, on this subject. I jumped right into it with a control-theory-based conjecture about the way religious perceptions and phenomena fit into the control model, and Rick, after expressing his views along the same lines, noted that we seem to have hit a touchy subject and offered to change it (not that we're limited to one subject at a time). And Joel Judd might have expressed more than one person's view when he said: "Talking about higher levels seems kind of ethereal I guess; not terribly scientific"

The interesting aspect of Joel's comment is that it is a higher-level point of view. To say that something isn't terribly scientific is to imply that we try to say things which *are* scientific. From this I deduce that one can perceive the degree of scientificness of a discussion. If the degree is less than some desired degree (very scientific), something must be able to detect the difference between the actual degree of scientificness and the desired degree. This difference, I take it, is the basis for whatever action is taken concerning the discussion, such as writing a sentence saying that it's pretty ethereal. Clearly, there must be a system concept about what "scientific" means, and there seems to be a control system related to it.

It seems to me that for those who consider stick wiggling boring and want to get into the more interesting higher-level aspects of the control-system model, we have here a wonderful laboratory in which to explore the real system, the one we carry around in our heads all the time. If I say something bearing on religion, your first reaction to it is evidence about the system concepts you have and are willing to defend. If it is possible for you to observe those reactions and bring out a fuller description of them, then you will have one foot in the point of view from which you can evaluate system concepts as a phenomenon, using a real live example. As you observe this example of a system-concept control system in action, you will see how control actually works at this level, and gain a deeper understanding of the way system concepts guide and use lower levels of organization, such as those having to do with principles and programmatic thinking—logic.

Of course, in order to do this, it is necessary, at least for the moment, to cease identifying with any particular system concept—that is, treating it as your own point of view. I would wager that very few of those who saw the "religious" topic go by did anything but identify with whatever system concept was operable at the moment. The disturbance was successfully counteracted; the incipient error was kept small. If the topic had switched immediately back to one of the other popular low-level topics, there would have been a little sense of relief, of relaxing the guard. The disturbance would have gone away.

And now here it is back again. So what's happening now? Same sense of error again? Same generalizations about why it's not a good topic? Same strategy for making it go away? Have you been here before? If so, why not observe what's going on this time? You don't have to identify with a system concept to do that. It's just a system concept, a phenomenon. It relates to principle thoughts and logical thoughts and familiar words and phrases hooked up into familiar sequences. When you're just observing it, it isn't a good concept or a bad concept; it's just what it is and it works the way it does.

Phenomena first. Theory second. Hearken to Marken.

Mary Powers: Wonderful! Along comes this new thread—religion—which I can't keep my hands off. We're talking about a bunch of systems concepts here—organized religions of various flavors, God, and what Ed referred to, as the 12-step groups do: a Higher Power.

I don't hold with organized religions any more than Rick, and for similar reasons—they don't do anything for me, and, in their names, people have done and do horrible things to each other. The latter is not so much a flaw of religion, though, as it is a result of the human bias to consider only as truly human the members of one's own group—those, you treat with the Golden Rule, etc., but for those others (unbelievers, heretics, etc.), anything goes (but that's another thread).

I don't believe in God either, simply because giving a concept like that a name concretizes it, and soon you have paintings of a man with a white beard zapping Adam into life. I love myths and fairy tales, but I don't believe them as explanations of how things came to be. I prefer

stories which work—models—to explain things: cosmology, evolution, continental drift.

But when that kind of story is eliminated, there is a major part of religion still left, and that is concerned with the principles one lives by. I'm not in favor of buying any particular religion's list, but I am in favor of spending some time thinking about such things and whether what one is doing with one's life is relevant and consistent with them. (Ed is concerned with what he perceives as a decline in morality—I am impressed by the huge jump in the last couple of years in books on ethics which have come into the library where I worked until recently.)

Of the three concepts I listed in the first paragraph, the one which makes the most sense to me in terms of control theory is the idea of a higher power. God, as they say, is everywhere, which means inside as well as Out There. Acknowledging a higher power is to recognize that there's a lot more to oneself than one's conscious self. Think of that forgotten name which appears an hour after you stopped trying to remember it, or, more seriously, the new idea or a solution to a problem (which can be intellectual, artistic, emotional, spiritual, moral, or whatever) which just appears, again not through conscious effort. One must consciously prepare the ground, but the answers come from a higher level than where one is consciously at, and it's no particular surprise that in a religious context they are called gifts from God.

It seems to me that this kind of thing happens best with practice, and the practice is letting go (the 12-steppers say, "Let go and let God"). The letting go is often done by sleeping. I take long hot baths. Many people do it by prayer and meditation. The interesting thing to me is that effortfully trying to get an idea or solve a problem looks very much like pushing on a conflict. As was discussed in the psychotherapy thread, control theory says that you cannot force a solution to a conflict, but resolve it by—whaddaya know!—going up a level. To one's higher power, or certainly to a higher level in oneself.

Whether or not doing this eventually leads one to being a more decent, moral person I do not know, but it seems likely to me. Over the last few millennia, the religious life has produced (in addition to bureaucrats, power freaks, and sadists) some very mellow souls, and it's worth looking at what they have to say—because they are talking (obscurely and metaphorically, usually) about levels of the mind which control theory, coming from the bottom up, is as yet only pointing at.

Izhak Bar-Kana: I respect the religions, at least those which I know, for trying to teach people that if *you* are not God, neither is anyone else around here.

Rick, I object to your arguing with arguments which are not mine. I am not sure we use the same names for the same things. When you mix free competition with stealing, something is wrong. Bill Powers might have a good theory, and I might see that people respect him for that. I might see that he is very successful, at least within this small universe called CSG. I might try to do better, and this is all competition is about. If I try to steal his ideas, then I am a thief. I might try to call him names, I might become violent, but this has nothing to do with free competition. Maybe this is related to the modern trend in sociology: "Why ain't I entitled to the same things?"

I am not interested in the public opinion about violence, as I am not interested in the public opinion about education, drugs, etc., *especially* in this country. As a simple engineer, I am interested in deeds. My friend, you might be killed in front of a lot of people, and no one would interfere. Even worse, they would run away... from the police, so they would not get involved, become witnesses, etc. The amount of violence in this country, which people seem to get used to, is unbelievable.

Ed Ford: My reference to a higher power or religion was only to establish *an example* of a system of values (systems-concept level), a system which varies with each individual, from mere lip service, to control or to harm others, to genuine concern for others. Within our Group, we have established an unusually high degree of rapport because we have all accepted similar values and standards. It isn't the values themselves, but our (to quote Bill) attitude or perception of our individual goals and wants which determines how each of us deals with each other.

And yes, faith (maybe a misused word) can be based on fact. My belief that George Washington lived is based on fact. So is my belief in the basic message and messenger of the particular religion I adhere to. That also is based on fact (just look at today's date).

Rick, your comment that it took "years for Western societies to free themselves from this source of conflict" is most interesting. Our faith in a higher power doesn't improve our ability to deal more equitably with others unless we translate those values to standards and decisions in a way which respects the internal control systems of others. Unfortunately, people have used these ideas as an excuse to control, abuse, and manipulate others ("even the devil can cite scripture to his means"). For a control theorist, what makes any living-systems concept valid is that it has as its basis a respect for the choice-making abilities of others —for the control systems residing in all of us. I really intended to use my words as an example of a systems concept in my discussion about competition, not to create an issue about the validity of religion.

Joel Judd: Since this topic is still alive, I'll repeat what I said last fall about the initial attraction of control theory, and that is its inherent respect for one's autonomy. Apart from the practical and conceptual shortcomings of behaviorism and cognitivism, what I dislike the most about them is the way they ultimately tend to take away one's choice, or at least responsibility, since we are just reacting to stimuli. My own religious beliefs are centered around the concept of "free agency," and control theory just confirms my belief that we are all free to choose. Freedom, of course, doesn't mean "anything goes," but it's in deciding what goes and what doesn't that groups of people get into trouble.

Chuck Tucker: I have found the discussion begun by Ford's answer to Bill's discussion of competition to be very useful, and I think that an ethical standard can be constructed from the exchange of posts and some reference to previous writings. I shall briefly support my suggestion with comments about the posts.

It was Ford who suggested that CSG members get along so well because "we respect each other and what each one of us has to offer. In short, our values are very much the same." But notice that rather than focusing solely on this aspect of Ford's post, Rick mentioned a "higher power" and the phrase "verbalisms rather than phenomena," and he disagreed with Ed's suggestion that we need more faith. Then Joel brought the conversation back to Ford's original point by saying that „faith" and „values" rather that a particular religion can be used as higher level concepts to bring about cooperation. Bill made comments on "higher power;" demonstrating that a "higher power" can be part of a control system and used cooperatively as a phenomenon. Rick followed with a discussion of his view of religion, and while noting that he was not taking an "anti-religious" view, he did end his post with a recognition that he was wrong in noting that religion was just verbalisms. Then Rick, after noting he might have offended someone, suggested that the subject be changed. But Bill returned to the discussion by making the concepts of religion, science, and logic almost on the same level. Mary noted how control theory can use higher-level concepts like "higher power" without a particular religious organization's "spin" on the concepts. She also mentioned that it is "to one's higher power, or certainly to a higher level in oneself" which one goes to resolve a conflict.

Now, what I make of these exchanges is an ethical standard at the highest level used by those who use and believe in control theory. This standard is: all human beings are self-regulating control systems and should be respected as such. Ride is correct when he says that most religious leaders (and their religious doctrines) do not respect humans as self-regulating control systems and try (rather unsuccessfully in most instances) to coerce / force / bribe others to follow their rules (which many do not follow themselves). I claim (see Bill's Chapter 17 in *Behav-*

ior: The Control of Perception) that control theory contains this ethical standard, and that religions would do much better if they would also use this standard. Thus, we in the Group get along so well because we place a value on and find as important the fact that each of us is a self-regulating system.

Rick Marken: For reasons I cannot understand, I count myself as one with a belief which makes it far more satisfying to know the truth than to be right. I tend to distrust and fear control systems which prefer being right to being truthful (or, since we rarely, if ever, get the latter, which admit that their "rightness" is tentative). It seems to me there have been, are, and will certainly continue to be control systems which want only to be recognized as having the right idea—an idea we would probably call a system concept. The methods of showing that these system concepts are right have too often included violence.

I argue that there is only one system concept I know of which has, explicitly, included as one of its working principles the principle that it is more important to know the truth than to be right. I think this principle implies a willingness to subject one's beliefs to the test to observation, logic and reasoning—i.e., *falsifiability*. Scientists who act as though this principle is not part of their system concept are no longer—from my point of view—scientists (even if they say they are and they do a lot of math and a lot of experiments). They are just ideologues—religious fanatics like the rest. I don't think any ideology (religion) other than science contains this principle of "truth over right" as part of its system concepts. The very essence of religion is revelation—"I know what's true no matter what logic or my experience says." What could be more dangerous? When I meet a religious person (or the exponent of any other ideology—i.e., a system concept not including falsifiability as a central tenet) who says, "gee, I might be right but I'm willing to change based on the evidence;" then I'll be greatly impressed. I might even join the religion.

Joel Judd: Most serious religionists, or at least ones I admire, would argue that the search for meaning, God, etc., is the search to be both true *and* right. I don't see the mismanagement and abuse of religion as negating any possibility that there are Truth and Rightness together somewhere. The problem, or paradox, is that I don't believe inquiring minds want to know; rather, there has always been the desire to prone God, etc., "scientifically," and I don't see that happening in the near future. That is why scientists argue against "religion" as Rick does: "The very essence of religion is revelation—'I know what's true no matter what logic or my experience says.' What could be more dangerous?"

Or what could be more sublime? I find it interesting that Rick uses the word "revelation;" because in my beliefs that happens to be a key concept. It refers to the idea that God communicates with man (which of curse assumes there exists God, etc.). No, it's not amenable to logic, but yes; I do believe experience can bear out one's perceptions of "revelation," if you mean the same thing by experience as I do. Revelation to me might just be "'luck," "good fortune," or a "timely decision" to you. There is no way I can "prove" to you it is right, or true.

One last thread which has run unexpressed through most of the "religion" polemic concerns the idea of "selflessness," for lack of a better word. Most major religions include some form of the doctrine that a human being reaches greater heights by thinking less of self and more of others. In Christianity, the paradox was expressed by Christ when He spoke of "finding" your life by "losing" it, explaining that serving others is somehow more divine than serving yourself. Included in this self-subjugation is obedience to God, with the understanding that He has had more "experience" and is in a position to suggest how we might make the most of being human. I would bet that a lot of tie people we admire fall into this characterization, whether or not they believe in a higher power. It's great to recognize your potential as a fully functioning control system, but I think it's even greater to reign in all that power and place it in the service of others to help them reach their potentials. While I'll never be able to "prove" that, that's the interface between science and religion for me.

Chuck Tucker: My point was that control theory, as I understand it, has an ethical principle which is on the same level as religions, theories, ideologies, or meta-meta-instructions. The principle is: respect each human being as a self-regulating control system. I also tried to make the point that most of those I know who hold to some religious doctrines do not use this principle, and that occasions much conflict, anger, despair, and other disturbances even more profound.

Rick Marken: Joel, you make me feel a bit like Scrooge McScientist. I think my hostility toward some aspects of religion masks my real love of many things which would also be called religious. (In fact, I realize that I keep posting on this topic because I am so drawn to, well, spiritual topics). It's hard for me to have a consistent attitude about a system concept (or set of them) which has brought us everything from witch hunts to what Bach wrote. There are some beautiful sentiments in the *Bible.* I love Ecclesiastes (by and large), and the stories of the *New Testament* are great. I love the character of Jesus. I love a great deal of Western mythology—Greek, Norse, etc. I'm not a big fan of the Eastern mythologies—but that is a matter of taste.

The problem with religion—what spoils it for me—is what you might call "literalism" or "fundamentalism." I think it's what is also called "faith;" I'm afraid. It's the part where you have to "worship" something or "believe" that something "really" happened or that something "really" exists although there is no evidence for it. There is no faster way to corrupt the sublime, from my point of view, than by making the "rightness" of it mandatory. The problem, I think, comes from the fact that religion (Western religion, anyway) filled at least three roles, two of which are now handled much better by modern disciplines.

One role of religion was explanation of what was observed—this is what Genesis and many mythologies try to do. Now we've got science —we understand that the wonderful imagination which created the "explanatory" myths is only half of the process of explanation—there must also be the discipline of observation and testing. But some people still want the "explanation-of-phenomena" role for religion—to give it legitimacy, I suppose. Hence we get creationists, flat-earthers and other, basically harmless, crazies.

The second role of religion is to express the unexpressable—the nature of the human spirit. This is now handled by art—poetry, music, etc. The *Bible* has some of the best prose and poetry going. It is art—some of the most inspired art of all time. Biblical art is a subset of a vast expanse of songs of the human spirit. But it is not special (other than in terms of how well if achieves its artistic goals of expressing the human spirit). It has no more privileged place in the art world than Shakespeare or Chaucer (or name your favorite poet). But there are still some who want to maintain that biblical writings are special—inspired by God. This leads to book burners and banners. These crazies are dangerous and quite unacceptable.

The third role of religion (and there might be more) seems to me to be rather unique to Western Judeo-Christian religion. This is the ethical role. Apparently, at some time long ago, some Hebrew tribal person realized that there was no obvious reason why s/he was being a nice person. And if s/he had no reason, then nobody else had a good reason, and they might go haywire at any time. S/he realized that s/he needed to tell people there was a reason why they should continue to be nice to each other—it's because they have 11th-order system-concept control systems watching to make sure that they have selected the right references for their principles. S/he just called these references "God." Not leaving anything to chance, s/he made sure that everyone knew that if they didn't set their principles appropriately, then they would suffer an error signal—eternal damnation in the fires of hell (catchy new name for an 11th-level error signal).

I suppose civil laws could be considered replacements for the written ethical standards (backed by threat of coercion) which had been provided by religion; but I don't think they are, quite. I think what Hugh Gibbon is doing in trying to analyze the system concepts underlying the law and our sense of justice is the start of a rational approach to understanding the ethical basis of our behavior. Chuck Tucker suggests that there might be an ethical principle which is part of control theory itself—but I don't think so. I think control theory can explain why we do (and don't) behave ethically, but it boasts no ethics of its own.

Because there is no really convincing modern discipline to replace the ethical role of religion (although I do believe that control theory might start to help—but don't expect anything interesting for a few decades), the crazies in this area of religion have been particularly prevalent and destructive. Nowhere else has religion caused more misery to innocent people than in the ethical bullshit it has imposed, based on the "wisdom" in ancient texts. I think the creationists are amusing and the book burners are annoying, but the ones bringing "God's rules" are just flat-out evil. I have had many homosexual friends whose lives I've seen made miserable and difficult because of the religious prejudice against this practice—because God says it's wrong. We have a massive overpopulation problem in the world, partly due to the fact that some nut cakes have divined that God doesn't like anything to come between semen and ovum (this one, alone, will probably be sufficient to end any hopes of a civilized society). From what I read, it seems to me that Jesus was the kind of guy who wanted people to find their highest degree of personal human fulfillment. He didn't get mad at prostitutes (who hurt no one, save possibly themselves) or homosexuals (again, who hurt no one except, possibly, themselves) or masturbators or birth controllers. Not even an adulteress. I think Jesus knew the difference between helping people achieve their own personal goals and helping people achieve his goals. I love selfless giving—but remember, that's *self less*. If Christians were really Christian, they would be out there trying to help homosexuals find the mates they want—not the mates the Christians want. Of course, these values of mine must be all wrong because they are not written down on an ancient parchment. Ah well.

Anyway, when it comes to religion, I think the aspects of it which really are wonderful can only be kept wonderful if they are brought back into the bosom of art, where they belong—where they will not be corrupted by the ugly drive for "rightness" tainting discussions of ethics.

Bill noted that discussions about religion, and our reactions to them, constitute hints about the nature of our own system-level reference signals. If you can get past the fact that the substance of these beliefs is considered "true;" you will notice that they are perceptions which you are trying to defend at particular references. Thus, our arguments, if analyzed properly (I bet Bill could help), are themselves a laboratory for study of control of the highest-level perceptions in the control hierarchy—definitely more interesting than watching control of the position of a cursor on a screen.

Another reason that religion is relevant to control theory, I suggest, is for the same reason that it is hard to keep religion out of discussions of the origin of life. Control theory, like evolutionary theory, is trying to deal with aspects of human existence which were once the sole purview of religion; with evolution, it is the origin of people; with control theory, it is the nature of the soul. Of course, regular old psychology treads on religious issues, too. But control theory gets to the "soulful" aspects in a particularly deep way. Control theory explains (rather than explains away) one aspect of people which most deeply defines our human nature—our purposefulness. Suddenly, teleology is no longer a spiritual mystery, but an understandable characteristic of closed-loop, negative-feedback organizations of matter. Most importantly, religion itself is an understandable part of the control model—it is a system-level purpose, an intention to perceive certain principles, relationships, categories, etc. This doesn't make gods or religions go away (just as evolution did not make gods and religions go away) but; like evolution, control theory certainly requires a thoughtful reevaluation of this system concept. There is just no getting around it. I can't help but feel that, to the extent that control theory is an improved model of human nature, reevaluating one of the most important aspects of human nature in the context of this model cannot help but be for the best.

Bill Powers: Joel, this is the point where in ordinary conversations I would say "Oh, sorry, I didn't mean to tread on your beliefs." This isn't an ordinary conversation. It's a scientific conversation, meaning that the participants are assumed to be more interested in improving their explanations of natural phenomena than in defending them. So when you say, "I find it interesting that Rick uses the word 'revelation,' because in my beliefs that happens to be a key concept. It refers to the idea that God communicates with man (which of course assumes there exists God, etc.)," I can only take this to be a scientific report. You are reporting a phenomenon (and in conversations of this sort, one main ground rule is that all reports are honest and taken to be honest). The phenomenon is "experience can bear out one's perceptions of 'revelation,' if you mean the same thing by experience as I do. Revelation to me might just be 'luck; 'good fortune; or a 'timely decision' to you."

The theory I propose to account for the phenomena of revelation, taking it as given that revelations do occur, is that (1) higher-order systems in the brain, operating at levels higher than the normal level which is conscious (whatever that mans), can inject reference signals which appear arbitrary and sourceless to the conscious systems; and / or that (2)

the process of reorganization can alter (at random) the way the conscious systems operate, including the way they perceive, so that sudden new understandings and new methods of acting appear, as if from nowhere. I would argue that there is no reason to think that such changes in the conscious world are due to any factor outside the brain —i.e., a supernatural being. On the other hand, there is no evidence that such supernatural intervention does not occur; we do not have the ability, now, to tell the difference between supra-conscious processes originating inside the brain and supernatural processes originating outside the brain—our only evidence is the experienced result.

Now, you go on to say: "There is no way I can 'prove' to you it is right, or true." You are referring, I take it, to the proposition that such revelations originate outside the brain I agree; I see no way to construct a compelling argument which would persuade any reasonable person of the truth or falsity of your proposition, or of mine. So, in terms of scientific knowing, we would have to agree that we do not know which is the coned proposition, if either. In such cases, we must choose something as a provisional belief, to take the place of knowledge. The question then is which belief to choose, not on grounds that it is "right" (because we do not know which is right), but on whatever practical grounds we can find.

One possibility we must entertain is that sudden changes in the conscious world sometimes might be due to normal reorganization or to the action of higher-order systems in the brain, and sometimes might be revelations from a higher power outside (or larger than) the brain. If that possibility exists, then we must ask about the consequences of making a mistake: of mistaking a brain process for a revelation from God. Suppose you suddenly get the thought, crystal-clear and compelling and as if from a higher source, "All of your troubles are being caused by the Jews. You must therefore kill all of the Jews, and purify the land." If you are convinced that this thought is a product of your own organizing processes, you will evaluate it in terms of all of your other concepts and understandings and goals, and quite probably will dismiss it as just another of those bright ideas which would not work out very well. But if you decide that this sudden idea is a revelation from God, you have no choice but to obey. The theory of God does not allow for ignoring God's word, or reevaluating it.

I think we must accept that thousands upon thousands of people have received sudden thoughts which they attributed to God, and as a result have committed what I at least consider to be unspeakable evils, thinking that they were acting under Divine Orders. In many theologies, the answer to this problem is not to say that such sudden thoughts arose from internal reorganizations and were simply not evaluated appropriately, but that they originated in *another* supernatural power:

Satan, the god of evil. The theory of God, in combination with observations which seem to attribute unacceptable characteristics to God, requires introducing the theory of Satan, who is responsible for the unacceptable "Divine" orders.

The *Koran* states quite plainly that God commands loyal Muslims to convert the infidels, and if they will not convert, to destroy them as the forces of Satan. I imagine that there have been many faithful Muslims who have undergone a crisis of the spirit over this teaching: God says you must kill these innocent people, while reason and compassion say that to do so would be evil. The power of faith, however, can overcome mere human reason and feeling. The good Muslim would subjugate his personal thoughts and feelings to the commands of God, and do what the Divine Word says he or she must do. I'm no expert on the Muslim faith, but I think this is not a grossly unrealistic scenario.

In this country, of course, our God (of Christianity or Judaism, to speak only of the majority beliefs) does not command us to kill the infidels (although not everyone would agree with that). So we have the case where in one part of the world, divine revelation contradicts what divine revelation says in another part. A crisis of the spirit in a soldier from the U.S.A. in the Persian Gulf War might lead him to decide not to kill an Iraqi soldier in his sights, while another crisis of the spirit in an Iraqi soldier might lead him to decide to kill the American who is in an equally helpless position Both reject what personal inclination demands, and submit eventually to the Word of God—with opposite results.

The theory of God keeps getting more complicated as problems like this arise. This theory, to say the least, lacks universality. It must be clear to the adherents of different faiths that their beliefs differ radically from those of others who also lay claim to belief in God. The only solution which does not lead to God contradicting Himself is to decide that one's own faith is the *right* one, while the others are in error on the points of dispute—they have mistaken their own thoughts for revelations from God. In countries where freedom of religious belief and expression are considered extremely important, this leads to the odd situation in which a constitutional edict requires distortions of the True Word of God to be tolerated. In other words, one must figure out how it is all right for other people to go against the word of God, while it is *not* all right for oneself to do the same thing.

All in all, I think that my theory makes more sense. It allows us to understand the experience of revelation in a way which does not require all people to experience the same, or even consistent, revelations. It does not in any way deny the reality of the experience of revelation: it merely explains it in a different way. In a context allowing equal consideration to all varieties and details of religious belief, I think my proposition remains free of contradictions and entails the postulation of the fewest entities, whereas the theory of God requires the multiplication of entities and the maintenance of principles which differ from believer to believer—all of them True.

Rick Marken: Bill, what can I say? Pretty strong stuff—a theory of the 11th order. I think your point about constitutionally mandated religious tolerance was great I've always wondered how it could really work, since it does require (if you believe in the "god theory") that you allow other people to go against the word of god while you don't. I think it is becoming dear that it can't work. It's not going to work in India any more. It's barely holding on in the U.S. It seems to me there are only two possible solutions—one (which I think Ed suggested) is to accept the god theory and hope (or require) that everyone agrees on just which god is really out there; the other is to give up the god theory and try an alternative—possibly brain theory: the theory of 11th-order control systems. I think that the latter is quite unlikely, ever. Pretty depressing. My rule of thumb, however, is to always try to live in the society having the least institutionalized commitment to a particular version of the god theory. I hope America can hold out for a while longer—but it looks like, after a brief period of enlightenment, the world is prepared to dip into another millennium of besting for the correct god theory. Oy vay.

Joel Judd: Rick and Bill: "I wanted out, but they keep pulling me back in." (Al Pacino in "Godfather III")

At the risk of turning this into a forum for personal beliefs, I want to mention some fundamental notions in order to respond to your comments. Assume (and I know this is a big assumption) the following scenario: there exist a couple of Gods (it takes two to have kids, you know) who have some offspring and want to offer a physical / mortal existence to them (for reasons I won't go into fully). This existence requires a place to live and the niceties of mortality—birth and death Part of the reason for sending the children away is to let them learn to make choices concerning—that's right—Good and Bad, Right and Wrong. Following the mortal part of this plan, the children will continue on immortally in different states of "maturity" and "knowledge" according to their actions on earth. Now, as soon as this plan is presented, two people offer to help carry it out—right again, Lucifer and Christ (both sons of God, by the way). (In case you think I'm making this up, check out Isaiah and Revelations, among other sources.) However, they quibble over an important issue: Free Agency. You see, Lucifer, being a good guy and a little bit selfish, offers to make sure that *all* of God's children make it back safe and sound—by forcing them to make good choices. Christ, on the other hand, says he will let everyone have a say in the matter, allowing them choices (and, knowing that children inevitably goof sometimes, he will do his best to allow everyone to make up for their mistakes, and show them how to do so). Well, we can find out how this (mythical) story turns out by looking at Christian theology. Lucifer becomes the bad guy by resenting God's rejection of his offer, and he and his followers leave without tasting mortality.

Returning to science, I try not to get worked up about science / religion (dare I say S-R?) arguments, because of conclusions like Bill's: "In such cases, we have to choose something as a provisional belief, to take the place of knowledge. The question then is which belief to choose, not on grounds that it is 'right' (because we do not know which is right), but on whatever practical grounds we can find." I believe that the crowning principle of mortality is freedom (as do you all, but perhaps for different reasons), and from my point of view, part of the reason for being here is to see what we'll do without that convincing certainty that "Dad" is always looking over our shoulder. However. "If that possibility [revelation] exists, then we must ask about the consequences of making a mistake: of mistaking a brain process for a revelation from God. Suppose you suddenly get the thought, crystal-clear and compelling and as if from a higher source, 'All of your troubles are being caused by the Jews. You must therefore kill all of the Jews, and purify the land.'... But if you decide that this sudden idea is a revelation from God, you have no choice but to obey. The theory of God does not allow for ignoring God's word, or reevaluating it." This and Rick's comments along the same lines point out many people's worst fears about religions. But religion can suffer from the same confusion as science. For example, the characterization of the "theory of God" given above assumes that anyone is justified in professing revelation and recruiting others to help. This is not the pattern in Christianity, where one person is called at a time to speak for God (as "Prophet"). Nor can a prophet say whatever he or she wants to say and get away with it. There are any number of checks and balances on people's behavior by which we can judge—"by their fruits ye shall know them," "do unto others,," etc. We all can think of worst-case scenarios where God, Christ, and others have been invoked in the name of genocide, purification, education, and other causes. But I don't think any of those crusades spread peace, goodwill, and cooperation, the hallmarks of God-like behavior. We can judge religion and religionists with a few almost common-sensical standards, like the couple just mentioned.

Bill says: "The theory of God keeps getting more complicated as problems like this arise. This theory, to say the least, lacks universality. It must be dear to the adherents of different faiths that their beliefs differ radically from those of others who also lay claim to belief in God .... whereas the theory of God requires the multiplication of entities and the maintenance of principles which differ from believer to believer—all of them True." Unfortunately, this is one of the best ways *to* turn people off about something: provide too many contradictory choices. Returning *to* the scenario laid out above, and assuming it were correct, wouldn't this be a great way to turn people off about religion / God?

There are two other issues I'll dangle. One concerns the idea of Spirit / Body (the soul). That revelatory communication (if it occurs) would take place at levels we generally talk about as *lower*, I find intriguing. I wonder about the Spirit / Body interface and how higher levels might relate to / communicate with things "spiritual l" as opposed to the more physiological functions of lower levels of the hierarchy. Of course, if you don't entertain notions of immortality, then such issues are not interesting.

The second issue concerns the perspective on life obtained from belief in God and belief in Man. I almost never bring this issue up, because it directly addresses the worst fears examples which always come up in discussions of religion If one is focused entirely on mortality and birth and death as the bookends of one's existence, then life often becomes overwhelmingly precious and something to be maintained at all costs. If, on the other hand, one believes that "life" began long before birth, and extends long after death, then the mortal part of this picture becomes almost a "drop in the bucket; " as it were. That *does not mean* that life is valueless or worthless, only that it is not *everything*. When someone whips out an *Old Testament* "myth" and shows how this beneficent God drowned thousands of Egyptians in the Red Sea, or murdered thousands of Sumerians in the Middle East, I tend to look at the context of the story (what we *don't* know about the situation as well as what we do), and consider the Big Picture. And when a child dies of malnutrition and disease in Bangladesh, or a family is wiped out in a Kansas tornado, I don't curse God, or complain that if God existed He certainly wouldn't let such things happen. Instead, I try to do my part to see that the corner of the world I can influence is made better.

God is not around to babysit us every second any more than most parents are around their 50-year-old children—but they certainly are available to give advice and offer solutions, *if the children ask* (and sometimes when they don't).

None of this is very scientific or convincing experimentally. But it's how I make sense of the world, and my life in it. *That* can be explained by control theory, as Bill and Rick and others have pointed out. But it probably can't be proved. Back to more mundane matters.

Bill Powers: Rick, before we get any further into showing the defects of various god theories, let's pause and figure out what we're doing. Control theory is not going to settle the question of the existence, nature, or purposes of God. That question isn't even interesting from the control-theory point of view. What is interesting is the fact that people support such beliefs and that the beliefs play some role in determining their principles, strategies, procedures, categories, and so on. If we wanted to play games at the system-concept level, we could make up our own stories about why we're here and what it's all about. We could seek converts, start a church or a political party, and go around claiming that our system concept is better than anyone else's. We could even have our own war once we got the hang of it. It's been done lots of times before.

Speaking strictly as a control theorist, a position from which I've been straying lately, what I'm interested in are the system concepts underlying the various god theories. I want to know if there are sets of principles from which they are drawn; if the principles guide logic and reasoning; if logic and reasoning select sequences of actions; if the sequences are indeed composed of symbols (category-perceptions)—and so on. In other words, I want to know if the hierarchical-control-theory model actually works as an explanation of human experience and behavior. As a control theorist, it isn't my business to offer free advice concerning which system concepts are the best.

As I said, I've been straying from this course. Straying from it involves saying things like "How can your system concept be the only True one when I know of many people who believe in a different and even contradictory one?" That amounts to trying to tell someone his or her system concept is no good. If people are control systems, and if they all have 11th-level (system-concept) organizations, and if they each develop in a fundamentally autonomous way, then of course they are going to end up with different system concepts, even when they think they have the same system concepts as others do.

In fact, it is very hard for people to agree on system concepts even when they try. It isn't so much that they resist having their system concepts modified to fit the group, but that they really have only a foggy idea of what the "group system concept" is supposed to be. Perceptions of this level are extremely hard to communicate. Religious and political groups keep forming and fragmenting for this very reason: the people develop divergent perceptions and goals, get into conflicts, and split up into smaller groups to eliminate the conflict. This happens in every case where people try to share important system concepts, not just in religion. If anyone gets fanatical or fundamentalist about control theory, it will happen here, too. The more important the goal (meaning the smaller the error that is tolerable), the less difference in interpretation is required to create a significant conflict.

There are many things we can say as control theorists about system concepts without getting into judging their substance. The point of a hierarchical-control model is to account for all of the levels of human functioning we can identify. We certainly have to consider an important subject like religious belief, because it is a phenomenon of human experience. We are even interested in the content of those beliefs. But the interest does not have to do with the correctness of the content; only with its relationship to lower levels of control.

So if I say, as I'm inclined to do, 'Joel, I don't believe the story you tell,' I am not speaking as a control theorist but only as a human being who prefers his own stories. I'm willing to argue on this subject as long as anyone feels like participating, especially if there are things I really should be doing but don't want to do, but if I do so I won't be talking about control theory. I'll just be telling you how William T. Powers is organized—one-five-billionth of the human race. Maybe I'm doing that when I talk about control theory, too, but I'm a heck of a lot better organized in that field than I am in the field of spiritual subjects.

So, Joel, it's quite unimportant whether I believe your story or not—as long as we agree that we're here to talk about hierarchical control theory. If you could analyze the story into system concepts, principles, programs, sequences, and so on, we could talk about how well the hierarchical model fits the way these perceptions work together and the way a person might behave to maintain them at their respective reference levels. Then we might come to understand something about belief itself, instead of trying to decide which beliefs are correct. I understand that, from your standpoint, your beliefs are true and right. From my standpoint, so are mine. With that settled, I think we can talk about belief as a phenomenon of human nature, and return to our original subject.

Rick Marken: OK, Bill, speaking as a control theorist, I think I am theorizing that religious phenomena (among others—such as ideological phenomena of various flavors—any experiences which seem to be based on a set of principles) are, in the model, 11th-order control systems. I believe the control-theory model says that different people want to perceive themselves as "Christians" or 'Jews" or 'Nazis" or "Communists" or "Pacifists" because of differences among these people in terms of 11th-order reference signals. One interesting thing about the 11th level (which you brought up) is that the reference levels for these perceptions seem to come from "outside" of the person. I imagine that a person whose reference for "religiousness" has him / her controlling for "Christianity" (as he / she understands it) experiences the source of this reference as being outside—the higher power that is above him / her. This is certainly the way I experience my own reference for religiousness (which is obviously set at "atheism"). It feels less like something I choose than like something I am.

For some reason it is difficult to become conscious of the fact that the reference for a system concept is selected by you, not something "out there" which imposes itself on you. Actually, system-level references *are*, to an extent, imposed on you (from the model's point of view) by reorganization. But it is hard to see that the reference for a religion is something your brain came up with in the same way your brain came up with a reference for a particular sitting position. For some reason, it is possible (though not necessarily easy) to learn that you are the one who selects the references for a particular configuration (like the sitting position), but it is nearly impossible for people to realize that they select the reference for a particular system concept—not to satisfy a higher-level goal, but as a result of fairly random reorganization to satisfy intrinsic needs. This might be an important point for therapists. The 11th order might be the "id" of control theory—the source of one's desires (references) for particular system concepts might be very difficult (if not impossible) to make accessible to consciousness. My hypothesis is that consciousness (whatever that is) can become aware of the source of a reference signal only if it can take a point of view from a level of the control hierarchy which is at the level from which that reference is sent. Thus, it is possible to become conscious of the source of the reference for the sitting-position configuration when you look at configuration perceptions as a means of perceiving a higher level perception—such as a particular relationship between your line of sight and a computer monitor. I suspect that it is difficult (or impossible) to look at system concepts from the point of view of whatever it is which wants to use system-concept perceptions to achieve its goals. Anyway, to the extent it is possible, the hierarchical-control model gives the term "consciousness-raising" a whole new, drug-free meaning.

The bottom line is that, from the control-theory point of view, system concepts (like the ones Joel and Ed and Bill and I are discussing) are perceptions maintained at particular reference levels set there for reasons which are not that well-understood (in terms of the model or in terms of one's own consciousness). I think a person who understands the control model has to accept this fact about the nature of his or her own system concepts.

Problems arise at the system-concept level, not because some system concepts are bad and others are good, but because (according to my understanding of the control model) people tend to assume that the references for their system concepts come from "out there." *That is the problem*. It leads to the conclusion that the level at which you want to keep a particular system concept is the truly right level—forgetting to add that it is just "the right level for you"—because it is *your reference*

signal. There is nothing wrong, really, with any system concept, as long as you can remember that it is just right *for you*, not necessarily for anyone else. This is the message of control theory about all levels of perception The "right" level of a perception is the level matching *your own* reference for that perception The only caveat is that, in controlling your perception, you should do so without interfering with the abilities of other persons to control their own perceptions. This interference is called conflict, and control theorists generally want to find ways to avoid it. Thus, system concepts like "kill the..." can be considered bad if you agree with this principle of conflict avoidance. Obviously, killing is the ultimate way to prevent people from controlling their own perceptions.

If people just could be happy controlling their own system concepts and letting others control their own system concepts, then all would be fine. I could care less what a person believes. My problem comes from the fact that most system concepts have principles involving other people—like the Moslem principle (and Christian, too) of converting the infidel. That stuff scares me; I think principles like that come about because people don't understand that system reference levels (the "right" way to be) are not "out there," they are "in the individual." Why system concepts seem to include edicts about how other people should behave is an interesting question—one that social psychologists, especially those interested in collective behavior, should look at very carefully. I'll leave that discussion for later.

It's easy to see when people are confusing internal references for external references. People who say "we have to do it right" obviously believe that their reference for whatever perception they are controlling is "out there;" so that anyone can control relative to it.

Joel Judd: Bill says: "Religious and political groups keep forming and fragmenting for this very reason: the people develop divergent perceptions and goals, get into conflicts, and split up into smaller groups to eliminate the conflict." This made me think of a couple of things: 1) The adoption of conquerors' religions in history, e.g., the Indians' "acceptances" of Catholicism in Peru. Many of their beliefs were tolerated by priests and have become part of the ritual worship for Andean people, a mixture of Pagan and Christian. A Catholic from New York visiting a chapel in Peru might be astonished or even shocked at the differences in what ostensibly is the same religion. 2) The problems caused by church clergy adopting political stances (e.g., Archbishop Romero). Either one of these would make a very interesting control-theory thesis for some student of political science, anthropology, etc.

I have perceptions of higher levels as possibly having rather long time frames—ditto, reorganization which might involve them. "Christian" might be a perception built up over 10 years, 40 years, or a lifetime. We probably would not consider a newborn a Protestant; the newborn itself almost certainly doesn't either. At what point do we say he / she is? When his / her behavior fits our perception of "acting like a Protestant?" Maybe this was all assumed in the discussion, but I wanted to make sure. The same thing would hold for other concepts, like "language," which develops over years of experience with language. Wouldn't some of the mysterious nature of the origin of higher reference levels be explained if we admit these longer time frames in their development? Then it wouldn't be possible to point to a discrete experience and say, "That's when I developed a reference level for 'family."' That would address the following comment: "But it is hard to see that the reference for a religion is something your brain canes up with in the same way which your brain came up with a reference for a particular sitting position." On the other hand, if reorganization commences to address intrinsic needs, and so much of peoples' reorganizations ends up working with religious ideas / God, what does that suggest about the source / purpose of intrinsic needs?

I can understand that Rick's theory makes no judgments about rightness / wrongness—it is an explanatory tool. That can be as true for development as it is for the description of a mature control hierarchy. But I'm not sure about the idea of negating "right things out there." Isn't there a "right way" of driving a car? That's not the same as saying that there is a right way to do *every little thing every time* I get in the car. Rather, there is a system concept for "right way to drive" which drivers share. We don't sit in the back seat to drive, we don't use our hands to manipulate the pedals, we don't go down the road backwards, though we can do these things. There's a right way to do a lot of things: use the language, pay taxes, get a Ph.D., worship God. For some things, though, there is more than one right way... uh, I just lost my train of thought.

Anyway, developmentally, we have models for developing concepts: parents, God, Michael Jackson In the case of children, we act as if there is a right way (*ours*) and expect them to adopt it. So how do we teach one another system concepts which we can agree on even though each is an individual?

Rick Marken: Here's a quick response to Joel's great questions: "We probably would not consider a newborn a Protestant... At what point do we say he / she is?" When you test for evidence that the person is controlling the appropriate variables. Just apply disturbances and watch for resistance. Acting "like a this or that" is not enough to show that there is control; for example, I can get you to write out a profanity as you move a mouse to counter a two-dimensional disturbance to

the position of a cursor. You are producing a profanity, but you are not controlling it—i.e., you will do nothing to resist my attempts to make your hand write a non-profanity. System concepts are probably not controlled until a person is well into the teens. Lower-level perceptual abilities also develop over time—you must be able to control configuration before you can control transitions. Plooij found clear evidence of this in chimps (who probably cannot perceive, and hence control, system concepts).

"Wouldn't some of the mysterious nature of the origin of higher reference levels be explained if we admit these longer time frames...?" The origin of the higher reference levels is no more mysterious than the origin of lower-level references. They are equally mysterious. The model accounts for the origin of higher-order references differently than that of lower-order references—but there is *no* mystery about how it is done in the model.

"I can understand that Rick's theory makes no judgments about rightness / wrongness—it is an explanatory tool. That can be as true for development as it is for the description of a mature control hierarchy. But I'm not sure about the idea of negating 'right things out there.' Isn't there a 'right way' of driving a car? That's not the same as saying that there is a right way to do *every little thing every time* I get in the car. Rather, there is a system concept for 'right way to drive' which drivers share." Bingo—1 think we have here a place where the content of your personal system concepts might come into conflict with the content of the system concept we call control theory. This might be a job for *Zen and the Art of Motorcycle Maintenance*, but I'll just give you the short answer. No, there is no right way of driving a car. There are just variable (perceptual) aspects of the car's behavior—some of which you can learn to influence in predictable ways—and you can bring these variables to reference levels which you specify in order to satisfy other reference levels. But there is no "right way to drive a car" unless you are talking about the "real-world" constraints on the way you can influence what you perceive. In my car, you can only accelerate forward (when sitting in the drivers seat on a level road) by pushing on a pedal under your right foot (with the ignition on). If you don't do this, it won't go. Same in our tracking experiments: there is only one "right way" to influence the cursor, because we've set up the world that way.

So the real world (the one we know only in terms of our physical models) does impose constraints on how we can influence the perceptual variables we are controlling—but the particular values to which we move our perceptions are right or wrong only in terms of whether they bring higher-order perceptions to their reference levels.

The term "right" implies a standard for comparison—a reference. If you believe that there are standards "out there" for how things should be, then I simply ask, "How do you know them when you see them?" Control theory explains how to determine when a variable is being controlled, and what the standard of reference is for the variable. When we apply this best to objects "out there;" we typically find that they are not maintained at a standard level unless there is a control system around making that happen. I'm afraid that, from a control theory perspective, "right" is defined by the control system—not by anything outside the control system which is not also a control system. This has got to be very disturbing to certain system concepts—but not to mine.

I do plenty of level-five-on-down control studies—and I think we're all convinced that we can demonstrate and account for the phenomenon of control at those levels rather well. It has to be considered one of Powers' most important insights that *all* behavior can, in principle, be handled by control theory. As I said in my foreword to Bill's book, Bill didn't invent control theory, but he noticed the appropriate way to apply it to living systems. He also noticed that all behavior—from tensing muscles to defending principles—is control and, hence, can be accounted for by control theory. What could be more important to promoting the control-theory view of human behavior than to show that system concepts, principles, programs, etc. are controlled perceptual variables?

So, while I think it is certainly nice to have more and more evidence that variables like temperature, chemical concentration, force, or whatever are controlled, it seems to me it could be monumentally more important to show that things like "atheism," or "humanitarianism," or whatever other system concepts such words only point to, are actually controlled variables—and to show how they are controlled, how disturbances are resisted, etc.

I am happy to volunteer myself as a subject for this investigation Perhaps Bill (or anyone else) could start testing for my controlled principles, programs, etc, by introducing carefully selected disturbances.

Gary Cziko: Rick says: "So, while I think it is certainly nice to have more and more evidence that variables like temperature, chemical concentration, force, or whatever are controlled, it seems to me it could be monumentally more important to show that things like 'atheism,' or 'humanitarianism,' or whatever other system concepts such words only point to, are actually controlled variables—and to show how they are controlled, how disturbances are resisted, etc." I agree that this would an important advance for control theory, but there seem to be (to me, anyway) so many problems in demonstrating this convincingly.

One problem is that if we disturb your principles enough, you might change (reorganize) them, and then we won't see you defending them

any more. If we keep telling you how dumb you are, you might at first resist, but after a while you might reorganize your systems concept so that our comments no longer create any error. We can show control at lower levels because we can count on subjects to be nice and adopt the reference levels we give them. But I can't see this working for high-level reference levels, such as belief in God, etc.

There are also serious ethical problems raised by disturbing one's perceptions at the higher levels. Joel Judd thought of giving students disturbing (inaccurate) test scores to see how they would react. Try getting that one pass the research review committees for human subjects!

But with all of the smart people out there in CSGNet-land (except for Rick Marken, of course), I suspect someone will come up with solutions to these problems.

Bill Powers: Rick says: "I believe the control-theory model says that different people want to perceive themselves as 'Christians' or Jews' or 'Nazis' or 'Communists' or 'Pacifists' because of differences among these people in terms of 11th-order reference signals." To be more precise, it's because of differences in 11th-order input functions. At any level, it's the input function which determines what function of which lower-level signals is to amount to a perception. The perceptual signals are just signals which get bigger or smaller. If you stuck an electrode on the signal, it would look like any other neural signal, no matter what it means. Same for reference signals: they just say "this much," not this much of what. The "what" is given by the form of the input function and which control system you're talking about. I am not at all satisfied with this aspect of the model, because it doesn't seem to capture the quality of perceptions. Chi the other hand, when you focus on any one perception very closely, it starts to seem like "just a signal" and to lose a lot of its meaning. Anyway, good or bad, this is how the model is presently designed.

People get a "Christianness signal" from all sorts of different lower-order perceptions, don't they? The perceptions contributing to Jerry Falwell's Christianness are certainly different from those contributing to the Pope's Christianness. It's very confusing when people use the same words for perceptions that are different. But they have to—there are more perceptions than words.

Gary, the test for the controlled variable doesn't require disturbances so large that they destroy control. All you need is a disturbance large enough to call forth an opposing (successful) effort which can be observed. If the opposing effort succeeds, there won't be enough error for long enough to produce significant reorganization. You won't change a person's principles or system concepts by pushing on them a little. But you will find out a lot about what the person will resist and what the person will let pass. “Disturbing” a controlled variable doesn’t mean pushing hard enough to cause it to change. It just means pushing hard enough to elicit an opposing effort which cancels the disturbance as far as the controlling person is concerned. If you use a large enough disturbance to succeed in overcoming the opposition, the result is, as you say, likely to be reorganization, and you won’t be observing the same system any more. But that isn’t how the test is used.

There aren’t any ethical problems in using the test correctly. You don’t actually change anything which matters to the person. The only cost to the person is a little effort to oppose the disturbance. This means, of course, that you must choose your disturbances so they *can* be resisted successfully.

Published by The Control Systems Group

Business Office	*Newsletter*	*Book Publishing*
P.O. Box 2566	*10209 N. 56th St.*	*Rt.1, Box 302*
Durango, CO 81302	*Scottsdale, AZ 85253*	*Gravel Switch, KY 40328*
303-247-7986	*602-991-4860*	*606-332-7606*

Volume 1

Number 4

Published by the Control Systems Group

Closed Loop # 4

Threads from CSGNet

CSGNet, the electronic mail network for individuals with control-theory interests, is a lively forum for sharing ideas, asking questions, and learning more about control theory, its implications, and its problems. The following "thread," stitched together from just one of the Net's many ongoing conversations exemplifies the rich interchanges among Netters.

There are no sign-up or connect time charges for participation on CSGNet. The Bitnet address is "CSG-L@UIUCVMD" (use no quotes in this and the following addresses); "CSG-L@VMD.CSO.UIUC.EDU" is the Internet address. Messages sent to CSGNet via these addresses are forwarded automatically to all participants. Via CompuServe, use the address ">INTERNET: CSG-L@VMD.CSO.UIUC.VMD" to reach the Net. Initially, you should send a note to the network manager, Gary Cziko, at "G-CZIKO@UIUC.EDU" (Internet) or at "CZIKO@UIUCVMD" (Bitnet); Gary's voice phone number is 217-333-4382.

Greg Williams
606-332-7606
October 1991

A (Control) Engineer among the Psychologists

Izhak Bar-Kana: My main field is adaptive control, with some neural networks and robotics. I subscribed to CSGNet mainly due to its "control" name and intended to be a quiet listener. The discussion is very interesting, and to me, quite surprising.

Regarding positive feedback and evolution, isn't "adaptation" the word? A simple time-invariant mechanism can perform that much. If the task becomes more difficult, "tracking errors" (based on the performance index the system tries to minimize) are used to change the gain, for example speed vs. error (as in the case of a pursuer), even in the simplest adaptive mechanisms. After a while, what previously was an extremal situation becomes a normal situation, because the evader is faster. A learning system identifies it as the normal situation, which in my humble opinion is expressed by development of muscles, etc., as when we train. But this is only a pretext to introduce myself, and to explain why I will be mainly a quiet listener, at least for a while, until I get the opportunity to read more of your works and speak, at least, the same language.

Wayne Hershberger: Welcome, Izhak. It sounds to me as if you already talk the CSG "language" very well, if I understand you correctly. I would put it this way: adaptation, including evolution, does not require positive feedback; rather, positive feedback requires (calls for) adaptation! What adaptation requires is *random* (polarity) feedback. That is, adaptation can be viewed as "slow" control (long-loop time) in which elements in the system periodically render the polarity of the feedback loop *random,* in order to "discover" the arrangement of subordinate elements which will restore the level of the system's error signal to within tolerable limits. An *E. coli's* control of its locomotion is the canonical (or prototypical) case of such slow control, complete with a nearly literal, random, "roll of the dice." By varying the rate at which the *E. coli* tumbles (or the rate at which Darwin's blind variations or Thorndike's trials and errors occur), the polarity of the system's feedback can be either (a) maintained when it is negative, or (b) changed when it is positive (and, thus, eventually restored to negative). The adaptation of organisms to overwhelming disturbances is the restoration of control. In his *Design for a Brain,* Ashby called it ultrastability. Bill Powers has called it reorganization. Perhaps it should be called "slow control utilizing random-polarity elements which ensure that feedback is negative in the long run."

Whether evolution (the adaptation of species) is "slow control" is a matter I will leave for others to decide. But the adaptation of organisms is clearly "slow control."

Tom Bourbon: Izhak, welcome, even if you hope to remain a listener. Perhaps we can coax you from that intention. Many of us are in the behavioral and social sciences, so we lack backgrounds in your areas of expertise. I, and probably several others, would appreciate information from you about good general references on the topic of adaptive control. I am presently working on models of human tracking behavior in which two people, or a person and a control-system model, interact. Two people can easily decide to change from one mode of interaction to another, and one person can easily recognize when the other has changed, then adapt to the new mode. I want my modeled person to develop the same capacity as a real one who detects the mode employed by the real person, then adapts. I'll admit I'm in over my head on the topic of adaptive control, but I suspect there might be some basic ideas there that will help me in my work.

I hope you will reconsider your decision to remain silent—yours is

precisely the kind of expertise many of us lack!

Izhak Bar-Kana: Many thanks to Wayne Hershberger and Tom Bourbon for the welcome. It is not easy to keep quiet in such an active environment, though I think I must do a lot of reading and listening to you before I even understand you.

To Bill Powers (and actually to all): I am asking more than claiming, but I am not sure I can agree with the apparent contradiction between engineering control diagrams and living control systems. Or, better, I do not understand it. If the problem is driving a car, the input is the way, the trajectory which must be maintained, and the output is the position of the car. Of course, this difference must be measured, and the control system only receives the output of the sensor which measures this difference. In ideal situations, this measure is exact. In other conditions it has noise, bias, miscalibrations, phase lags ("time constants"), and/or transport lags ("pure delays"). The control system tries to bring the error signal to zero, and the output is the position of the effector ("actuator"). Between the sensor and the effector (motor) there is a controller which transforms the signal in such a way that stable performance of the control system is guaranteed. And this is only the simplest control system. If a "brain" is involved, the signal transmitted to the effectors cart take more sophisticated forms: the brain might know the performance of the control system, might be capable of taking into account its time lags, etc. Furthermore, the brain has stored the final aim of the trip, and might change the route or make other decisions which could not be taken by a simple autopilot whose only purpose is to keep "in line." But I think there must be some separation of the various tasks. And even here, the final point is stored in the brain only because some real final point is in the real world, and this is what we call "input," even though the control system can only affect the output of its own sensors, or its perception of the real world. If the temperature must be maintained, the input can be considered internal, because it starts in the brain. Still, this signal is transmitted to a control system whose function is to execute and reach this temperature, or to annihilate the difference between the desired temperature (registered in the brain?) and the temperature of the body. For this control system (or better, regulation system) the input is external.

I don't understand how the living control system affects its inputs. Maybe only a difference in definition? In a tracking system, the position of a target is the input; the resulting position, of the eye for example, is the output, even if the only physical and measured signal is the difference between these two values. I agree with everything I can claim I understand in Powers' "Manifesto," so maybe I miss the main point here. I would appreciate if you could open my eyes here, because

I am trying to understand, not to prove that I am right.

To Rick Marken: Maybe the engineering control people need other tools because they must *design* the control systems, not only understand them. The "sophisticated" control people use lots of math because of the difficult task of proving that a system is stable. Not because they are crazy about stability, but because it is easy to get an unstable system with an "ingenious" and "intuitive" control method. When control is nonstationary and nonlinear, such as in adaptive systems, the problems and the proofs are even more difficult. The problem is that if you cannot know (prove) that an adaptive system (I mean "engineering" adaptive mechanism) is stable, in general you will discover that it is unstable under some conditions. I don't know how much this group is interested in or how much time it has to spend on this stuff, unless people are interested in the instability mechanisms of pathological cases.

Please see my lines on the car driver above. Of course the organism only receives the signal supplied by the sensors, but that is more or less the measure of the external signal. By the way, besides dealing with theories of systems, I am also an engineer, and I can tell you that no engineer would let a motor run, much less a plane fly, without thousands of simulations, no matter what the theory says, and in fact the theory, the complex functions, differential equations, etc., do not say much when a real, large, complex system is involved. And I would not dare to compare any complex plane with a living organism, not to mention an intelligent creature. So, learning through observation and simulation is a main engineering tool. But when I want to *design* a stable and well-behaved system, I need mathematical tools which express stability and performance, and their dependence on the various parameters I might or might not change. And then things start getting tough, like trying to define pornography: It is hard to define, but it is easy to recognize when you see it. Yet I usually need the differential equations to have reliable simulations, especially if I want to discover when the real plant stops performing satisfactorily. It is not that important whether your simulations are state-of-the-art or not, as long as they are correct and approximate the real thing. I don't know your models, so I hope they are.

To Tom Bourbon: At this stage, I am afraid I can only tell you that the problem is interesting, and that I only have begun studying it. It is not as much an adaptation problem as it is a learning problem. How to guarantee that a mechanism learns while it performs its task and maintains a stable behavior is not an easy task! I will try to be more specific in the future. In fact, part of the new trend in "intelligent (automatic) control" tries to eliminate the differential equations because "the brain does not solve differential equations," and tries to imitate the brain; the algorithms used are just (poor) attempts to reproduce the activities of organisms' neural networks.

Bill Powers: Izhak, I think that the mental model you are using is the one traditionally given in engineering texts, the same one that Norbert Wiener picked up and used in his first book on cybernetics. In that model, "input" means *reference* input. It is shown, usually, entering the comparator as if from the external world. The feedback signal, on the other hand, is just "picked off" the output variable through some feedback transducer.

In the model we use in the Control Systems Group, when we say "input," we mean the sensory feedback input, not the reference signal. That is because the sensory inputs in the organism constitute the "feedback pickoff" which reports, as analogue signals, the states of external controlled variables. The senses do not report the intended or desired state of affairs; only the current actual state of affairs. The reference signal comes not from outside, but from systems superordinate to the control system in question, inside the organism.

So it comes down to how we match the main functions and signals in a generic control system to corresponding functions and signals in a particular control subsystem in the organism. The abstract organization is the same; our model has the same connectivity as the one I believe you are using, so the control-system analysis itself is unchanged. But the meanings and the implications are greatly changed.

Apply this to a model of driving a car. The driver sees the current position of the car relative (laterally) to the road. Out of all of the information in this image, the brain extracts a position signal which varies as the car moves from side to side. Thus the position of the car is the input variable, not the output variable, in the steering control system. The position signal is compared with another signal which specifies the intended state of the position signal: that, of course, is the reference signal. The driver can select any possible perceivable position on (or off) the road as the reference position. The error signal, reference minus sensory signal, actuates the output of the control system, which is the torque applied by the arm muscles to the steering wheel (this requires two phase advances for stability). That torque is the last variable in the output chain which is due strictly to activities in the brain. From there on, we have mechanical linkages and external disturbances coming into play, which alter and add to the effects of the output and are not themselves part of the behaving system. The result is some position of the car on the road, and thus a state of the perceptual signal representing that position. The feedback effect keeps the perceptual signal in a match with the reference signal (give or take dynamic and static errors). It is not necessary for the brain to contain any detailed knowledge of physical properties and events outside itself other than the controlled variable. Variations in output properties have little effect; disturbances are automatically counteracted without any need to anticipate or sense them (except through their effects on the perceptual signal).

With the reference signal moved inside the control system, we can now "parse" complex behaviors in a new way. In order to alter the position of the car on the road, the brain now needs only to alter the reference signal for the steering system that is now in place. To pass a car, higher systems concerned with relationships to other objects change the reference signal enough to move the car to the other lane, keep it there a while, then move it back. Of course there is also a speed control system operating independently, with its own input which senses speed and its own actuator which affects speed (the foot on the accelerator pedal). The "passing-another-car" system alters the reference signal for speed, too, as the driver passes the other car, first increasing it, then decreasing it. So the higher system uses, the lower systems by manipulating their reference signals, just as a human user manipulates an artificial control system by turning the knob which changes its set point.

At the same time this is going on, the driver can use one arm and hand to reach out and change the volume on the car radio, then scratch his neck, all while telling a joke to the passenger. In this model there are many control systems acting concurrently, each controlling just one (perhaps complicated) variable. There are neuroanatomical justifications for breaking down behavior into multiple control systems operating independently and in parallel, and organized into levels of control. And I think this picture also helps us to approach the modeling of complex behavior in an orderly way, solving problems of peripheral control to serve as the foundation for exploring systems at higher levels, more central in the nervous system. We can, of course, pick isolated systems at any level and analyze them as control systems, absorbing lower-level control system properties into their output functions. But the final model must spell out all of the stages of control which exist, while, one hopes, maintaining correspondence to known structures in the nervous system.

The same model applies to human temperature control. The input variable is the temperature of a sensory ending (in the hypothalamus, I think). The reference signal is variable, as temperature can be maintained actively anywhere between 98°F and about 104 or 105°F. I don't know what varies the reference signal, although I know it changes when you get sick. The error signal is translated into shivering and peripheral vasoconstriction if it is positive (sensed temperature lower than reference temperature) and into sweating and peripheral vasodilation if negative: that is the behavior which affects the input, the sensed temperature.

As for simulations: we use them a great deal, where we know how to construct them. They work very well. Stabilization has not yet proven to be a problem, although in the arm model you have seen mentioned, the problem was solved just by introducing known properties of the neuromuscular systems in question (we never set foot on a complex plane). I think that the hierarchical structure simplifies stabilization problems, which might be an indication (and might not) of why the whole system is organized that way. We haven't gotten very far with modeling very complex or high-level behaviors. We're still taking baby steps and learning how to walk. But I think that our approach, probably combined with some of the perceptual models being developed by neural network people, will carry us a good deal further before we have to change the basic structure of the model.

Izhak Bar-Kana: If this is a family with specific definitions and problems, which I do not belong to, I would rather keep quiet or say "excuse me, it was nice meeting you." But if it is a control-theory group, and when the thermostat is an illustration, then some things deserve, at least, clarification. The model you all seem to use, position = position + k x error, is of course, perfectly correct. However, I am confused about what you call output, controlled object, and control objective.

No matter how good or how bad is the function of the organism, or of any control system, the only thing which it can affect is the output. Even if one closes one's eyes, or if one is drunk, he or she still controls the output, which is the name of the controlled variable in my dictionary. A closed-loop system, properly designed or properly organized, will try, in general, to minimize the signal representing the error between the measured reference signal and the measured output (controlled variable). But this is not the only way to do the job.

A fresh driver, on a new route, with a new car, performs very much as was said above. However, a better controller uses all possible prior knowledge to get a "better" quality of control. In our case, the effect of learning (and here we very much try to learn from organisms) seems to become cardinal. After a while, the brain has a sufficiently good model of both the route and, for example, the car, and the control is very much open-loop. Based on some details of the route, the brain predicts both the changes and the rate-of-change of the route, and also the response of the car to various inputs, and passes to the car a very complex signal, mainly open-loop control. No, I do not try to advertise open-loop control! There is still much uncertainty in my model of the route and in my model of the behavior of the car, therefore I keep my eyes open and monitor the error between the desired trajectory and the actual trajectory, but now the gain of this error can be much lower, and my mainly open-loop control is now smoother and "better" than the closed-loop only. In both cases, we control the position of the car. The input is the controlling, not the controlled signal, even though in closed-loop it might be hard to tell.

Now, about the external reference. Indeed, no route can tell me where I want to go, which is a decision. Once I decide where I want to go (and, hopefully, based on the knowledge or valid "representation" of what is going on out there), I must follow a route which exists somewhere (whether I follow it or not). This is the reference input which is measured and transmitted, after "cleaning" (filtering) it of noise and processing it so it fits the needs of the control system (the transfer function with its leads and lags, nonlinearities, etc.). At the same time, the "result" (as named in Bill's answer to me—why not "output"?) of the control, the position of the car on the route, which I call output, is measured and compared with the reference. The difference is then the controlling or one of the controlling signals which now affects the output. The objective of the control can be to cancel this error, or to minimize it without requiring the use of all vital reserves, or to minimize some combination of the error and other variables.

What is wrong here? I don't use this model because it is Wiener's model, or anyone else's model, nor because it fits some complex mathematical formulas, but because that is how I understand the control systems. If a thermostat is designed to maintain a constant temperature, then it is a regulator, and the referenced temperature can be considered as being internal, corresponding, I hope, to some desired but real temperature. Good or bad, the system can only control the output, the heat in the room; either it measures it correctly or not.

Sometimes here I seem to get old names for old things, only shifted. For example, Bill, what does it mean to say that "position is represented as integral of motion?" Position is the integral of velocity. I am not surprised that in the brain, motion is hierarchically higher than position, because motion is obtained by changing positions.

Bill Powers: Izhak, we do have some work ahead if we're to achieve communication. I'm sure that when it comes to control theory itself, you are the expert and I am the amateur. But I don't think we will have any arguments in that area. Where we differ is in how we analyze the behaving system itself, prior to setting down any equations. This leads to some differences in terminology, which we can certainly straighten out. But there are also some fundamental assumptions which do make a difference, and we might have some long discussions about those.

Let's start with the thermostat, honoring tradition. I partition any control system into the system itself and an environment. Where you draw the boundary is arbitrary, as you must have a closed loop in the end, but I think it's worthwhile to give some thought to the natural

boundaries. In the thermostat, the bimetallic element is the system's sensor of temperature, so that is a natural input boundary. The contacts actuate a relay, which turns on the furnace, which converts a supply of oil or gas into a thermal output; that is another natural boundary. I call these "natural" boundaries because they separate signals and functions in the controlling system from processes in the independent environment. The sensory signal (position of the metallic strip) depends only on the temperature of the sensor itself. The thermal output of the furnace depends only on the stored energy and the operation of the furnace. All other processes between the thermal output and the sensor (in the external world) are subject to potential disturbances and changes of parameters, and those changes occur independently of what happens "inside" the system, as I have defined it.

Now, with these definitions, what is it which the thermostat controls? That means, what variable is held most nearly at a specified reference level, in spite of all kinds of changes which can occur in the independent environment? Clearly (to me), that is not the temperature in the rooms of the house. If a bedroom window is open, the bedroom will become cooler. If a window is open in the living-room where the thermostat sensor is, the air near the window will become cooler, but the air immediately around the sensor will be maintained at the reference temperature (plus or minus the dead zone). So we know that the temperature of the air immediately around the sensor is closer to being the true controlled variable than the temperature of the air on the other side of room or in a different room of the house.

Now suppose we set an infrared heater on the floor, so that its radiation passes through the grille of the temperature sensor and falls directly on the bimetallic strip. What will happen? The contacts will remain open and the furnace will remain off until the temperature of the sensor element has dropped once again to the reference temperature. This will happen when the air around the sensor has cooled enough to remove heat at the same rate that the infrared radiation is adding heat. So the room temperature will drop and the temperature of the air at the sensor will drop. What will remain the same? *The signal (position) which stands for temperature.* The temperature of the sensor itself will be held near the reference temperature—but nothing more remote from the sensor will be controlled. By varying the infrared heater's output, you can cause any air temperature you like, and the control system will still experience zero error.

So by thinking of various ways to disturb the temperature outside the system, we show that the only variable reliably stabilized against independent perturbations is the state of the sensing element itself, and of course the signal associated with that state. Variables more remote from the input boundary of the system will be stabilized by the same control action only if they are shielded from extraneous disturbances by something other than the control system, and only if the sensor's temperature depends reliably on those remote variables.

I think that if you analyze any control system carefully, you will see that the same principle applies. The only variable stabilized against independent and unpredictable disturbances *of all kinds* in the environment is the input signal of the control system.

Now, the output. You suggest that the output of the thermostat System is "the heat in the room." But if that is considered the output, then we are in the position of not being able to predict the thermal output of the control system's furnace, because the heat in the room depends on many factors other than what the furnace is doing. The outside temperature, sunlight coming in the windows and falling on outside walls, open windows, fires in a fireplace or a stove, lamps, and blockages of air circulation all contributed to the heat in the room. The furnace's thermal output also contributes to the heat in the room, but it does not *determine* the heat in the room. The only thing which the control system can *determine* is the amount of heat released in the furnace by burning fuel. So I would call the output of the thermostat the quantity of heat released (per hour) by combustion in the furnace. What happens to that heat after it is generated, what other sources and losses of heat might be acting at the same time, is impossible to predict from knowing how the thermostat is designed. The only "output" which depends *strictly on the operation of the control system* is the thermal output of the furnace.

Between the thermal output and the temperature of the sensor, we have a variable and unpredictable environment. The control system is not equipped to sense any of the causes of those variations. Nor does it need to be so equipped. All it needs to "know" is the temperature of its own sensor. By knowing that, and by being able to vary its contribution to the temperature of its own sensor, it can control that temperature.

This is the basis of my general approach to analyzing control systems prior to reducing the relationships to equations or programs. What do you think of it so far?

Izhak Bar-Kana: Bill, thank you for the compliment, but you know and I know that one must learn a lot to discover that no one is expert. In fact, I try to understand your discussions as control theorists in psychology because they might be very relevant to my questions in artificial learning. I am not sure I really understand the difference between my model and yours; maybe what you call "input," I call "measured output."

By the way, the temperature sensor is supposed to be in distant rooms, not in the neighborhood of the furnace. Now, if the window is open near the sensor, the temperature in this room is not affected

by the thermostat system. In this case, the system receives some value of a low temperature, sends out lots of heat, the only thing it can directly control, and has no effect on what it is supposed to affect, the desired output, the temperature in the room, or on its measurement, the feedback input to the system. It does not change the fact that the input signal is used to affect the output or different stages of outputs, some of them measured if needed, to guarantee satisfactory behavior of the system. The basic design would take into account the basic thermal properties of the room, and nominal ambient external and internal temperatures, and the rate of fuel burned and heat supplied would have to maintain this nominal condition "almost" with no other regulation. So, the input is the desired temperature, and the output is the actual temperature. Because I know that uncertainties are always present, I monitor the actual "controlled output" and use the difference between the desired output (by the way, I also call it the reference input, so maybe I just live in another kind of dichotomy) and the real output, and use it (after filtering the measurement noise) to generate supplementary heat that, hopefully, will compensate not only for uncertainties, but also for changes in the ambient temperature, and so on. I might even decide that the closed loop is enough, especially if the desired temperature is fixed. The error between the desired and the measured output is used as input to the controller, amplified and processed, and then sent to fix or change the rate of fuel, the rate of heat, and change the temperature in the room. If the loop is well-designed, it finally brings the room temperature (as it measures it) to, or close to, the desired temperature.

I might repeat myself, but now suppose that the sensor is broken and frozen at a fixed low temperature. Then the control system gets some constant input and sends waves of heat, changing, I call it controlling, this output, whether this is the desired output or not. Anyway, I think I start to understand you, and I am only afraid that it might be difficult having a dialog with the control community at large, if the claim "the control system controls its input" is not understood. In spite of this argument, I think that things are not as distant as they might seem. Some more eavesdropping and more reading from my side will straighten out things even more. I am not used to such a high level of patient discussion about different opinions, and I am honored to participate in it, even if I might introduce noise.

Bill Powers: Izhak, I am very pleased that you are so open-minded and willing to work out these problems of communication. As we go on, I'll try to address selected topics in your communications which might help resolve whatever problems there are. Today's text: "I am not sure I really understand the difference between my model and yours;

maybe what you call 'input; I call 'measured output."' I think you have the key to one of our differences in nomenclature. What I call input is what you call measured output. I call this "output" an input partly because it is the external variable which affects the sensor, an input to the control system, and mostly because it is affected only indirectly by the actual output of the system—that is, the actuator or effector. The state of the controlled variable is not determined by the system's action alone; control is required because there are unpredictable disturbances which also contribute. If we use the term "output" for the controlled variable, then we have the odd circumstance in which we can't define the output of the control system itself—the output depends on independent factors just as much as on the behavior of the effector. I prefer to reserve the term "output" for the effector's action, which is the last thing in the chain of output processes which is completely determined by the control system. Between the effector's output and the controlled variable there are many sources of disturbance—they are the primary reasons control is needed in the first place. We can agree on the term "controlled variable." But I claim that variable is more closely associated with the sensory input than with the effector output.

I don't think it would be practical to design a home thermostat as a basic open-loop system with feedback added to handle details. The "details" are the whole problem. The steady-state thermal output of the furnace which is needed ranges from zero to the most the furnace can produce, depending on unpredictable heat losses and gains in various seasons and at various times of day and night, varying conditions of cloud cover and wind velocity, and various conditions of occupancy. A real home thermostat is simply designed by picking a furnace which can keep the room above the maximum desired temperature on the coldest cloudy day at 100 per cent duty cycle, and then letting the feedback do the rest, as it does. The problem a thermostat has to deal with isn't "uncertainty" in the sense of system noise. It's the fact that there are very large and unpredictable disturbances of the controlled variable. When the main causes of variations in the controlled variable are major disturbances rather than noise, there's no way an open-loop branch in the system can accomplish much.

Human systems, I believe, are in the same situation. Most of the output (motor) activities which take place are there to counteract large disturbances of controlled variables. There is very little random variability in the system itself—only a few percent of the range of action. Living systems got the reputation of being highly variable because the wrong model was applied: psychologists thought that the behaviors are responses to stimuli, whereas they are probably just actions which protect controlled variables from disturbance. The disturbances, of course, were mistaken for stimuli. Because disturbances occur unpre-

dictably, the behavior counteracting their effects is just as unpredictable. But if you know what variable is being controlled, behavior suddenly looks far more regular: it opposes the effects of disturbances in a highly systematic way.

Suppose you have a motor controlling the angular position of a load through a gear train. The angular-position sensor is located at the load and not on the driving shaft, because the shaft can twist and there can be play and runout in the gears. The actual output of the system is not the position of the load, but a torque applied to the armature. If an extraneous force is applied to the load, the torque will immediately rise to counteract it. If the force is too large, the motor still produces maximum output, but the shaft does not turn: the position no longer is affected by the system's output torque. If we now remove the position sensor and substitute a tachometer, without changing anything else (except perhaps the stabilization filter and scaling amplifier), the controlled variable becomes angular velocity instead of angular position. The output torque, in the steady state, equals the sum of all frictional and viscous resistance plus any opposing torques, and the angular velocity matches the reference signal. So *the nature of the sensor determines the nature of the controlled variable.* That's another reason for saying that the controlled variable is an *input* variable.

The reference signal is certainly one input to the system's comparator. The sensor signal is another. But I like to say "reference signal" because in living control systems, reference signals very often are supplied by higher-level systems, not by sensory inputs. In fact, I can't think of any case where a known reference signal comes from the sensory inputs. But we can get into that later. Is my nomenclature beginning to make any more sense to you?

When an engineer designs a control system, there is a control problem defined in advance. The engineer can see what effectors are needed, what sensors are needed, and even (sometimes) how control might be achieved by open-loop means (if there aren't any unpredictable disturbances of the controlled variable). That's because the engineer has an internal reference defining the desired result and the ability to shape a device having known effects on the physical world. The engineer can see all the inner details of the control system, and he or she also has advanced knowledge of the properties of the physical world with which that system will interact.

An evolved organism like a human being, in the process of becoming organized, doesn't know any of that. The environment is known only through sensory inputs and direct physiological effects of the environment on the body (the state of which is also sensed or at least sampled). Actions affecting the environment are produced only by sending signals into muscle systems. There is no a *priori* information in a growing brain concerning physical properties of the environment or the body, or any 'laws of nature" or general principles. The brain can't use any of the engineers knowledge in building up its own control systems as it matures.

Whatever the brain does, it must do on the basis of available information and whatever amount of organization it has and has acquired at a given stage of development. It discovers properties of the external world relevant to control only by acting and sensing the result. No information is available about what happens between action and sensed result: only the result is experienced. Many consequences of acting are insignificant and unrepeatable; only some are consistent and therefore possible to control. The organism learns about the consistent sensory results, selects remembered states of those results as targets to repeat, and by trial and error finds the combinations of output acts which will tend to restore the sensed world to any former state, starting in any other state. That is called control. It does not need to know, and in most cases never does learn, *why* performing a given act results in a given effect on perception. It does not need to know the true nature of Reality, explaining why perceptions behave as they do. Only when higher-level cognitive systems develop does a brain begin to acquire a symbolic understanding, a model, of the external world, so it can explain why acting in a certain way is necessary if control of a certain kind of experience is to be possible. Only at that stage can someone become a control-system engineer—and by that time, the vast majority of the engineer's own personal control systems have been in place for years.

If you ask a child, or for that matter almost any adult, *why* turning a certain knob on a television set makes the picture get brighter, the answer is going to be something like, "Because that's the brightness control." But nobody has trouble with adjusting the brightness to whatever level seems "right." Control is not based on understanding of the physical world. That's lucky for you, if you are a bacterium or a baby.

Here's my nomenclature for a living control system:

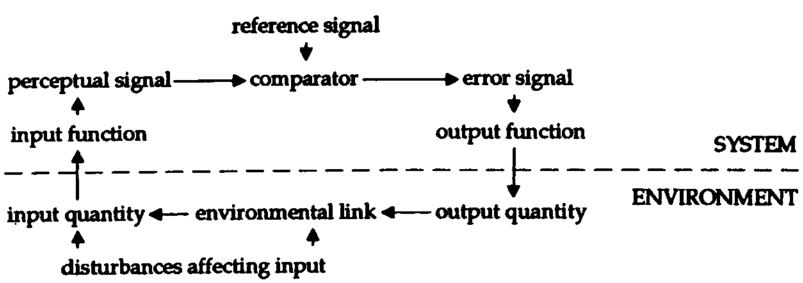

In the diagram, "quantity" = physical variable outside system, "signal" = physical inside system; "output function" = effector, "input function" = sensor + sensory computations, "output quantity" = direct measure of effector action, not subject to external disturbances, "input quantity" = physical variable directly affecting sensor, "disturbances" = independent contributions to state of input quantity, and "environmental link" = path by which effector contributes to state of input quantity. The observable controlled variable is the input quantity. The output quantity varies as disturbances vary, and cancels their effects. Note that the reference signal originates inside the system.

Izhak Bar-Kana: Again, trying to read and understand the discussion which is going on, I have the feeling that I broke into a subject in which I cannot even call myself a novice. I am very interested in your discussion, because I want to use any information related to biological intelligence for our "engineering" intelligent controls. I am pleased and surprised with the nice and patient tone of this discussion, and with the detailed explanations I get.

1. My position in control is actually very similar to yours. I am an engineer, a control one, and mainly, a feedback-control one. In what I call "simple adaptive control," I have lots of fights with some colleagues, *because* I show that some simple feedback loops with "appropriate" adaptive gains can perform much better and can be more dependable than many sophisticated schemes.

2. I cannot afford to mix control with philosophy. For control jobs, I take external "reality" for granted. When my car does not crash into a wall or another car, I deduce that my sensors' approximate interpretation of "reality" was not very bad, which I am not sure of if I am drunk. In both cases I control my sensorial perceptions. I am glad that I am around here and now, even though this phenomenon was not present some time ago and will vanish sometime in the future, and in spite of the fact that I am "mostly void" and from what is left, more than 90% is just plain water. The car which comes toward me might discontinue this experiment too soon, even if it (the car) is also "mostly void."

3. I take for granted that you are there, and that I read a translation of your thoughts, even if my perception might be very poor.

4. As I understand it, the great idea of Bill Powers is the observation or the discovery of *feedback* control in behavior, as opposed to a simple stimulus-response interpretation. I could not agree more.

5. I might make people angry, but because feedback is closed-loop, it is not very important to me what you call what. I only need the dictionary.

6. I do not confuse effect and control. I assume that in order to control, one must be able, first of all, to affect, and in a desired mode, too.

If the thermostat cannot control the temperature in the room, how can it control the measured temperature, which is its input? I hope you do not translate "is affected by" into "controls." I assume the thermostat system is designed to perform a reasonable job, as our behavioral control loops are "designed" or have developed reasonable tasks. The others are not around to testify.

7. One of my problems is that I am, in general, in the middle, and I try to understand all sides, and get the main points of opposite ideas. I am afraid that this group, trying now to explain every behavior as closed-loop, might ignore some very nice and intelligent open-loop controls, based on the splendid property of the brain: learning. Yes, closed-loop is dependable, and when one learns skating, one uses very stiff closed-loop control. With time, after one learns one's own behavior and the response of the skates, one uses much open-loop control, based on the learning process and modeling of this behavior. Closed-loop control is still there, but not alone, and not as stiff as before ("lower gains"). A predictive closed-loop control is also there, comparing the predicted desired position to the actual present position, but this loop is weaker and weaker with training. How much is open and how much is closed is only a question of control gains. It is not a cause-effect interpretation. It is very intelligent open-loop control. In my humble opinion, one misses something if one ignores this aspect of control, in particular when learning (I wish I knew how) is involved.

8. About "controls output" or "controls input": as I said, it is not very important (for me), if I know what we are talking about. Yet, many of your conclusions are based on modeling. Control experience tells me that I cannot derive conclusions about the behavior (or transfer function) of components of a control loop from the behavior of the closed loop, *because* of the nice property of feedback, namely, lack of sensitivity to variations in the parameters of various components. Therefore, to find the transfer function of the *closed* loop, I must first *open* the loop, test the input-output behavior of the *open* loop, when any variation of the components expresses itself one-to-one, and only then close the loop. The math also shows that analysis (at least Gedanken) of an open-loop system provides conclusions relevant to the behavior of the closed loop, but this is only secondary here. Anyway, this might explain why I have an input and an output (or more outputs on the way, but only the controlled variable is the output of interest). Now, I measure the output, and the result of measurement becomes the feedback input, and the other input is the reference input. Along with the closed-loop scheme, which is basically identical to Powers' scheme, I add a function of the desired (reference) input which directly drives the control system, bypassing the comparator.

Bill Powers: Yes, Izhak, it's hard to convince some people that a simple control system can accomplish more than many extremely complicated approaches to the same task. One of the great difficulties is getting people to think in terms of continuous variables, isn't it? The digital revolution really brainwashed everyone. Even in electronics, technicians are happy to learn digital circuits because they're so easy to understand—but they "don't do analogue." They barely understand Ohm's Law, and most of them don't even know what "impedance" means. So much for the great leap forward.

I completely agree that the best attitude in dealing with control-system design is that of the realist. You can't play the piano, either, if you're wondering if the keyboard is really there.

I don't think I agree with you about the gradual progression from closed-loop ("stiff" skating) to open-loop. But we can save a lot of detailed arguments if you will just get my book, *Behavior: The Control of Perception,* and go through it.

Concerning temperature control, I disagree with your statement that a thermostat must "control" room temperature if it is to control the measured temperature, the input. It must be able to increase and decrease the heat content in the room (that is what I mean by "affect" the room temperature), which in turn affects the temperature at the sensor. But the temperature at the sensor is affected by other things, too, not just by the furnace. If something cools air near the sensor, the average room temperature must be raised higher than the set-point in order to bring the sensor back to the set point. If you think in just a little more detail about physical processes in the room, I think you will see that there are temperature differences in various parts of a room, and that there are many variable sources of heat and losses of heat which alter temperature in various places in the room. Only the temperature of the sensor itself is controlled—that is, kept near the set point.

The realization that only the input is controlled by a control system (and not just the idea of feedback) was my "great idea." That's the meaning behind the title of my first book. You already understand control theory—just think about it, you'll get the idea. It's so simple that it took me three years or so to understand it. If you understand it in less time, that will show that you are smarter than I was.

Rick Marken: Izhak, if you know of any example of control by an open-loop system, then I want to know about it! And I would really like to know how it works. I claim that if a variable is demonstrably under control, then that variable is part of a closed-loop negative feedback loop. I think your examples of open-loop control are just higher-order variables under closed-loop control. For example, control is better in pursuit than in compensatory tracking. This is usually explained in terms of open-loop control—the person learns to "predict" the position of the target in pursuit tracking and makes "open loop" movements in anticipation. To the extent that the prediction is correct, then control appears to be better than in the compensatory case, where the temporal course of the target is invisible. I claim that the "prediction" is just a higher-order controlled variable—like a transition or sequence—and, to the extent that prediction works, it is really a result of the subject's ability to detect regularity in that higher-order variable. The subject is then controlling a transition, not just the instantaneous configuration of target-cursor discrepancy. The fact that this is what is occurring can be tested. If the prediction is open-loop, then failure of the prediction should result in no adjustment—after all, open-loop means not monitoring the consequence of the output. There are experiments using regular, predictable patterns of transitions as targets which show clearly that, when the pattern is changed (from a circle to another pattern), there is a clear adjustment to this change. The fact that the subject is controlling the sequence of transition rather than just the target/cursor configuration is evidenced by the rather long time it takes to makes the switch—about 400 milliseconds.

Still, I would be very interested in an example of what you consider to be open-loop control.

Izhak Bar-Kana: Rick, you are so enthusiastic about what you believe in, and work so hard on it! This is the way to results, even if at this stage or another one is wrong, and you are not wrong. Did you miss the start of my letter? I am a feedback-control guy, and I do not agree to give up the closed loop, not even for the sake of this discussion. I only claim that the dosed loop does not explain everything, and, I would dare to say, in particular when brain, learning, and intelligence are involved.

A curiosity. Recently I had an argument with a distinguished colleague, whose argument was: I know the desired output of my plant, and I know the plant perfectly (the transfer function); therefore, I can use this desired output as an input to something which performs exactly the inverse of the plant. (By the way, for everybody, the "plant" or the "process" is what we want to control, but is already there, given. For example, the motor, the ship, the tank. We can only add sensors and the "controller.") This way, the output of the plant will do exactly what I want. And then, I say: My friend, this is too idealistic, one can never know exactly the plant, the disturbances, etc. Furthermore, a transfer function followed by its inverse is not the equivalent of 1. (There is a problem of noncontrollability—I don't want to elaborate.) One cannot control without the closed-loop.

As you see, it is peculiar that I am now in the position to convince

you that one can control without feedback. But now, let us go on from his idea: Assume, for a moment, that we know *exactly* the plant. We can then design a controller such that the plant will perform exactly what we want (not necessarily "exactly," but in a satisfactory way). If we do not know, let us try to learn about it, or to "identify." Once we identify, we can design our controller. Now, I ask, how can one be sure that the identification works correctly? And how can I be sure that the plant will not be destroyed before the identification is finished? And what if the identification process is corrupted by noise? All these thoughts occurred when I wanted to design some controller for a manipulator. It is a nonlinear system, and because of that, the load varies a great deal. If one uses only the high gains needed in difficult situations, one only amplifies noise in other situations. If one uses low gains, then performance cannot be obtained. Then, based on very little prior knowledge about the manipulator, I build a simple adaptive controller *(closed-loop)* such that the gains move up and down as a function of the tracking error, and the performance is quite good. However, since I do not use knowledge, the adaptive gains "work" very hard. Therefore, in parallel with the controlled plant, I use an identifier. The SAC (simple adaptive control) guarantees that no disaster will happen, even if the identification does not work properly. However, when the identification is correct, the controller based on it takes over and my adaptive gains (of SAC) decrease, and may vanish if they are not needed. The closed loop is there, and if the tracking error tries to increase, it will push it back. *Yet,* there is a signal, directly from the input reference to the plant, and for most situations, is now the only signal which controls the position of the manipulator. Is it not enough? Is it not capable of accounting for disturbances? Is it a bad control system? A bad control system is just that, it does not become a non-control system. And now, because my closed-loop gains are adaptive, even when they are called upon to correct for uncertainties, they do it at much lower values, usually, than before the addition of the open-loop control.

I wish I knew how the brain does its modeling and learning. I think that the identifier must use low gains for identification ("slow identification" or 'long-term memory") so that it is not much affected by nonrelevant transients, and only stores relevant knowledge. The SAC must be fast, to get the gain needed when it is needed.

Bill, now the argument is very close. What you call the input we should both call the "controlled variable." If I measure the temperature in the neighborhood of the sensor, then this is the controlled variable. However, this is not the end of the story. A good control system would be careful with its sensors and (as organisms do) would use some redundancy, measuring the temperature in various points, thus maintaining some relevance in the measurements, by averaging the various measurements or even by eliminating unusual inputs. Furthermore, a lot of noise might have been added to the measurement, and a close look would reveal a lot of filtering used to "ignore" this part of the input, and pass only the signal which is relevant to the controlled variable. Only a very primitive control system would just respond to any signal (or sensation) received as input, and respond to any spike of noise which comes from who knows where. There is indeed the danger that the control system might respond to any input signal, even if it is not related to the designed controlled variable, but a good control system takes care of it, sometimes using prior knowledge, sometimes identifying the disturbance and compensating for it, sometimes filtering undesired signals.

Maybe our different backgrounds lead us to put emphasis on various aspects of the same phenomenon. You claim that I control the input, because my output was disturbed by some bias (mainly constant disturbance), and I claim that I do not control the input just like that, because in most cases it is mixed with noise and must be processed before I can be confident (always only to some extent) that it represents the designed controlled variable. About open-loop control, the manipulator example above is also relevant. About open-loop in organisms, I don't have another example than myself, and that is only one sample. I know that before I know my car or my trajectory, I use a very stiff closed-loop system. If I get off the highway, my control loop takes me back immediately, and I might reach the other side or crash into the cars moving in the opposite direction. Experience, or the teacher, teaches me to ignore this signal, keep the same direction, and come back slowly. I see here a combination of closed-loop and open-loop control, but this is only one opinion.

Rick, you are right that the evidence on prediction does not prove anything about the existence of open-loop control. Actually, from an input-output point of view, open-loop and closed-loop schemes are equivalent. They have different properties with regard to sensitivity to uncertainties, stability (including oscillatory or non-oscillatory responses), etc. But the argument is not which concept is better, only whether open-loop control is there. In my case, of the manipulator, I know that the open-loop controller is there, because I put it there, and I see that the closed-loop "disappears" when open-loop is sufficient, because the gains vanish. I have no intention to claim that an open-loop system can deal with drift, disturbances, uncertainties, etc. I only claim that forms of open-loop control might exist along with closed-loop. The relative gains, or weightings, might vary, and in some situations, each one can be zero. If you have evidence that in biological systems there is no open-loop control, then I cannot argue, because I simply don't know. But one can explain some behavior either way,

again, because from the reference-controlled-variable point of view, they are equivalent.

Otherwise, I can add to your argument: a good closed-loop control system does use prediction *in the closed loop*, if it is needed. Any phase lead in the forward (I am afraid to call it feedforward) path, or velocity (tachometer) feedback *along* with position feedback supply information on the future development, in other words, prediction. And this is only an elementary example of prediction.

Bill Powers: Izhak, I believe that one of the important steps one must take in understanding human control systems is to recognize the role of perception. If we were building an artificial device, then as its designers we would be aware of the true state of the controlled variable, and the state of the signal produced by the system's sensors as it represents that variable. We would be able to see what is signal and what is noise; we would know whether the signal is properly calibrated and linearly related to the external variable. But in the human organism, there is no third party who knows all this. I think that as a designer of systems, you can appreciate how the world must look to the system itself. It exists only in the form of the sensor signals, and whatever other signals are derived from sensor signals by computation. How can we discover the world as it is represented by the outputs of human sensors and sensory computers? The answer is so simple that nobody seems to have thought of it: just look around. Feel, taste, hear, and see. There it is. You are already experiencing the sensors' output signals, and you do not experience the world which gives rise to them. I do not think that this insight would come easily to anyone who has not worked with artificial sensors. But most people, in the end, understand it if they persist.

If you look at your experiences as signals in this way, you will realize that there is very little noise in them—they are almost perfectly noise-free. Only in unusual circumstances—near-perfect silence, the threshold of darkness—do we experience our perceptions as behaving in a way which seems at all "noisy." Also, you will realize that linearity and calibration mean next to nothing, because you are looking at the output, not the input, of the perceptual functions. You have nothing to calibrate them against but the outputs of other perceptual functions. You even pick reference signals from previously experienced sensory signals, so the scale of reference settings contains the same nonlinearities. You can say that *this* perception is nonlinear with respect to *that* perception or with respect to a meter reading (another perception), but you cannot say whether all perceptions are nonlinear in a different way with respect to their causes.

I think that this orientation makes a great difference in the way we build models of human behavior. We must realize that however the brain manages to bring external variables under control, it must manage this completely on the basis of information available to it through its senses—its uncalibrated senses. It cannot look at the plant (the universe outside) and see what compensations are necessary in order to represent its variables properly. It knows only the variables, and even then only after they have already been represented as internal signals. The only way it can identify properties of the plant is through experience with sensory representations of the plant's behavioral variables in relation to sensory representations of the organism's own output efforts. It knows something of the inputs to the plant, and something of the outputs from it, but it knows nothing directly about the plant. The world outside is a black box. We who have seen artificial control systems both from the outside and from the inside have some advantage in understanding this situation, because we can appreciate what is lost when you lose that disembodied vantage point from which you can see what is happening on both sides of the sensory barrier and on both sides of the output boundary.

This means that when we try to guess how organisms learn to do things like adjusting their internal part of the loop gain appropriately, as in your simple adaptive control, we must try to see how they can do it on the basis of information available *inside* the controlling system. In your case, you have accomplished that: just use the information in the error signal, which is inside the system. If the average error signal gets too large (which takes in many possible problems, such as oscillations), reduce the gain of the output function, the effector part of the system. I don't know how your identifier works—does it, too, work only on the basis of sensory signals available to the system as a whole? Or does it need external intelligence to tell it what to identify?

When we build models of human control systems, we naturally have to play the part of the "external intelligence" just to set up a plausible system. But our goal must be to learn how the system itself can come to acquire those design features we find necessary, without knowing what we know about our own created system designs. I have felt for a long time that the people trying to reproduce human "pattern recognition" have been on the wrong track, because a "teacher" is an essential part of their approach. Some external intelligence must tell the recognizer if it is right or wrong. Real organisms do not have such a teacher, not when it comes to learning the basic perceptual and control processes themselves. Recent work on neural networks and perceptrons is, I think, a little closer to the right approach, because the system in part creates its own organization. But there is still a teacher who knows the right answer. Real organisms never know if their answers are right, except in terms of how well they serve to control what happens to the organisms.

My reason for total rejection of open-loop control is based on thoughts like these. How can the organism find the feed-forward output signal which will create "almost the right behavior" of the plant, without monitoring the behavior of the plant? You have to imagine setting the direct output signal to just the right value which will keep the car centered on the road without *ever* seeing the relationship of the car to the road. It is impossible without the aid of some third party who knows what the plant is really doing. You could, of course, have sampled control, so that corrections are applied only now and then. But that is still control, and it cannot work without feedback. All of what is commonly called "feedforward" is really explainable only in terms of a hierarchy, of feedback control systems (in many cases, including those in an engineer silently standing by in the background, screwdriver in hand). Only through feedback can the so-called feedforward be properly adjusted. There is no friendly omniscient engineer in the background adjusting our own "feedforwards" for us.

Your example of overcontrolling a car is, I think, only a description of how we learn the right dynamics of response and get the control system stabilized. In the end, we have very fine high-gain control with proper temporal filtering so that oscillations are eliminated. But we do not notice this control because it takes place in our midbrains, brainstems, cerebellums, and spinal cords, where we seldom pay much conscious attention to what is going on. It doesn't seem that we are exerting much control effort when we drive down a straight and level road, but just watch the steering wheel! It moves with every little bump in the road, every slight change in the crosswind, every little tilt of the roadbed. This control system is extremely sensitive to error—but it keeps the error very small, so it does not have to make big efforts. Unless, of course, there's a big disturbance. We habitually observe from a higher level of abstraction, and we don't notice the errors or the corrections because they are happening at lower levels. The car seems to go straight by itself. But just try holding the wheel absolutely still, and you will see that significant disturbances are always present. Their effects are cumulative. If their effects aren't precisely corrected, the car will quickly go off the road.

Rick Marken: These are some comments on the thread related to open/closed-loop control. Bill has looked at the issue "from the inside." Let me try it again "from the outside"—looking at control as the observer of a controlling system, rather than as an example of one of these systems. When we look at a living system, we see that it produces many consequences. These consequences are potentially variable—the temperature at the skin, the position of a limb, etc. The value of the variable at any time depends on many factors—the "causes" of the variable.

Thus, y = f(a, b, c, o), where y is the variable of interest and a, b, c, and o are the variables which "cause" y; f is the function which determines how y varies (over time) as a function of variation in the causal variables. If a, b, c, and o vary over time, then y should vary over time in a manner determined by f. If, however, y remains approximately constant over time, then we might imagine that something fishy is going on. Stability of y could happen by chance—the variations in a, b, c, and o could just happen to produce a constant y. But the longer this goes on, the less likely it becomes that stability is occurring by chance.

Moreover, if we can trace the stability of y to variations in o, which happens to be the causal influence on y exerted by the living system, then there is strong evidence that o is systematically counteracting the effect of a, b and c on y. I take these two pieces of evidence—the stability of y and the fact that system outputs are the sole cause of this stability—as evidence that y is a controlled variable. This is evidence of control "from outside the living system." It says nothing about how this control is achieved. What I am claiming is that the only organization we currently know of which can provide an explanation for control is control theory—that is, the theory that y (the controlled variable) is part of a closed negative-feedback control loop. One other part of that loop must be a reference signal which specifies the particular value at which y is stabilized. If y is stabilized at different values, then this reference signal must be variable. Observation of the environment of the control system reveals no variables "out there" which could possibly function as the reference signal (although people have been fooled into thinking that "targets" in tracking work this way; the simple way to show that they don't function as references is to show that people can reliably keep y stabilized at values different from the target value "when they want to"). Thus, the reference signal must be inside the system itself.

The controlled variable, y, need not be a simple aspect of the system's environment. We see living systems controlling very complex variables, such as their relationships with other organisms. I have seen people keep a variable called "in love" at nearly the same level for periods of years. The ability to control such variables implies an ability of these systems to perceive such variables. Perception, from a control-theory point of view, is not some arcane discipline of only peripheral interest to psychologists. Perception becomes central to understanding human nature. What people do depends on what they perceive and where they want those perceptions to be.

This model of what we see as the behavior of organisms is radically different than other models of behavior currently embraced by most life scientists. It is a model which works, which satisfies the requirements of scientific method, and which provides a comfortably humanistic view of human nature. That is the reason for my enthusiasm,

Izhak. I guess I "believe" in control theory; but not in the usual sense of belief. I am not reverential toward it. If it proves to be wrong, I will happily abandon it for the improved point of view. I believe, based on experimental evidence and matching the behavior of models to that of living systems, that control theory currently gives us the best (and only) model of how people (and other living systems) work. I think this is not only scientifically important, but socially important. I think a case can be made for the notion that people have been screwing around with each other and making life more difficult for each other because they have been looking at each other as a particular kind of object—one which can be controlled from the outside. The control model shows that this is not only false, but also a sure recipe for conflict. And I think most people would agree that conflict between people (classes, religions, nationalities, whatever) has been the continuing obstacle to the possibility of every individual (other than the winners of each conflict) leading a graceful, dignified and satisfying existence.

Gary Cziko: An open question to Bill Powers (or other "serious" modelers): I just finished giving a presentation which involved showing Bill Powers' computer simulation of a simple control system. One person in the audience made the point that because the computer was doing the controlling, it had to be a iterative system. My somewhat lame reply was that, yes, it is iterative on the computer, but that the slowing factor added to the model makes it work like a continuous system.

But I suppose the point the person was making is that iterative control can work, in which case we do have responses which are computed based on the present static state of a number of variables. This is what the computer does, and I suppose all digital control systems as used in engineering do the same. *But* there is nothing in either the data we get from real subjects or in what we know about nervous system and muscle physiology which leads us to believe that control works this way in organisms. So we use the digital computer with slowing as an approximation to the continuous control we get with living control systems.

Somebody let me know if I'm on the right track here.

Bill Powers: Yes, you're on the right track, Gary. The slowing factor is introduced to keep variables (at least one variable in the loop) from jumping instantly from the old value to the next computed value. A real arm obviously can't be in one position at one moment, and in a position 20 degrees away in the next millisecond. The slowing factor is chosen to fit with the assumed physical time represented by one iteration so that the actual amount of movement is similar to the real amount of movement over the least element of time. The less the time represented by one iteration, the more slowly the variables must change. The slowing factor, being in the denominator, must increase as dt (the one-iteration time interval) decreases.

When we run models, we want to run them quickly, so we can try the model over and over while adjusting parameters for best fit with the real data. So we start with a relatively large value of dt. If the interval is too long, we don't get as good a fit as when it is shorter. At some length of time interval, around 1 /20 to 1 /30 second for some of the models I have made, making the interval shorter just slows down the computations without improving the model any more. This shows that over roughly 1 /30 second, the variables in the model vary slowly enough so that the response is essentially the same as if the sampling were infinitely fast. We choose the interval dt so the results are the same as if we were sampling the behavior at an infinite rate.

Even with this explanation, there is still often a problem in getting people to see the difference between an iterative quasi-analogue computation and a sequential computation. In a sequential computation, each variable is calculated in turn, just as in our computer simulations. But the mental image which the listener is thinking of is really cast in terms of *events.* First there is an input event which causes a perceptual signal event. Then the perceptual signal event is compared with the reference signal to yield an error signal event. Then the error signal event causes an output event—a response. And while these events have been taking place, what has been going on at the input? This is the question they overlook; they assume that the input event is finished, so nothing will happen until the next input event occurs, perhaps "triggered" by the response. So each function in the loop takes its turn in acting, and then lies quiescent until it's aroused again. It's never aroused again before all of the other functions have had their turns.

In the real system, of course, the input varies continuously. All of the functions are doing something all of the time. There might be a delay before the next function in the loop receives a given input value, but during the delay the input continues to change. So the next function receives a continuously changing signal, delayed, even while new changes are being introduced at the input. There is a pipeline effect. It's like talking to someone over a satellite link. Your voice vibrations are received at the other end continuously, but delayed by the length of the link. This is very different from thinking about input events and output events.

A truly sequential system would be represented by a feedback loop, digitally calculated, without any slowing factor. We can boil such a loop down to an extremely simple example:

A = B
B = -l0A

If you start with any value for B (except exactly 0), this loop will run away on successive calculations. But suppose we introduce a slowing factor:

A = A+(B-A)/S
B = 10A

Now the loop will converge so that both B and A approach zero, provided that S, the slowing factor, is larger than 5.5. If S is 11, the final state will be reached in one jump. If it is larger than 11, the approach to the final state will be monotonic from any starting value of B.

Even the cybernetician William Ashby fell into the trap of sequential calculation. He concluded that negative feedback systems couldn't have a loop gain of -1 or more negative and still be stable (note that the above system has a loop gain of -10).

The implicit reference signal in the equation above is 0. You can put in a nonzero reference level for A in the second equation:

B = 10(A* - A), where A* is the reference level

Now the system will approach a state with A nearly at the value A*, from any starting condition. You can make A come closer to A* by raising the loop gain:

B = 100(A* - A)

But the system will oscillate unless you increase the slowing factor. If S is made equal to 101, the final state will be reached in one jump. If S is larger than 101 (say, 300), the approach will be monotonic. If G is the loop gain, then S must be greater than or equal to G + 1 in order to get a stable approach to the final state. Note that G is a positive number for negative feedback because we are subtracting A in the above equation.

The final state you reach is predicted by solving the *first* two equations (without the slowing factor) as a simultaneous pair. If there is a non-zero reference value A*, solve this pair of equations:

A = B
B = G(A* - A)

It's not obvious, but introducing the slowing factor converts the pair of equations from a simple algebraic system into a differential equation.

That's why we are able to stabilize its behavior in time, even with loop gains as large as we please.

I realize that you're not going to take a naive audience through all of this in a one-shot lecture. But if you play with these equations enough to get the feel of what is going on, plugging numbers in and running the iterations, you'll probably be able to cope with the misunderstandings a little better.

Izhak Bar-Kana: I assume that Bill's equations with S are as follows (the minus was missing):

A = A+(B-A)/S
B = -l0A

Now, the program takes it as written, but if we want to describe what is going on, the equations should be these:

A(k) = A(k-1) + [B(k-1) - A(k-1)]/S B(k) = -10 A(k)

Substituting B(k-1) in the first equation gives the following:

A(k) = [(S-11)/S]A(k-1)

We get the condition S is greater than 5.5, because we want the loop gain K = (S-11)/S to satisfy the stability condition |K| is less than 1. So, what is Ashby's mistake? What do you call the loop gain?

In general, we get these equations:

A(k) = A(k-1) + [B(k-1)-A(k-1)]/S
B(k) = G[A* -A(k)]

Notice that A(k) and B(k) are Anew and Bnew, while A(k-1) and B(k-1) are Aold and Bold. A* is constant here, but it might be any function of time. The same substitution gives this result:

A(k) = A(k-1) - (G+1)A(k-1)/S + GA*/S

or

A(k) = [1-(G+1) /S]A(k-1) + GA*/S

Now, one selects G and S so the loop gain is |K| is less than 1, where K = 1 -(G+1)/S. If this condition is satisfied, and thus a stable equilibrium

point exists, it is reached when A(k) = A(k-1), and we get the equation:

A(k) = [G/(G+1)]A*

(if I don't have an error of algebra), which tells us that in such a simple system one cannot have perfect following even for a constant input (sorry, I mean reference), unless G is infinite.

Bill Powers: Izhak, your analysis is precisely the same as mine, and you found a shorter way to prove that S = 2/(1 + G) is the minimum value of S for convergence (see my 1978 article in *Psychology Review* for a longer way).

Why can't I learn to get critical signs right when I publish equations? You are correct about the sign of "10B," of course.

Actually, with S greater than 5.5 but less than 11, the approach to the final state is oscillatory, and the oscillations are an artifact of calculation (if you're trying to model an underlying continuous system). The oscillations occur at the iteration frequency and are not tied to physical time. Only when S is greater than 11 can you model the real motions of a physical system.

You have defined the loop gain here a little differently, so that it is the gain allowed by the slowing factor on each iteration. I wish I had thought of that—it's so easy. I would call the loop gain G (or 10 or 100, depending on which equation you read) because that is the gain which predicts the limiting case (infinite integrations)—that is, A[infinity]. In the limit, A = G/(1+G)A*, and S drops out. You arrive at the same result, quite correctly, by specifying that A ceases to change. The same result is given by taking the equations:

A = B
B = G(A* - A)

and solving them simultaneously by substituting B for A in the second equation:

B = G(A* - B), or
B(1 + B) = GA*, or
B = [G/(1+G)]A*

Solving these equations simultaneously is the same as saying that these two relationships hold *at the same time,* so this is a control system with zero time-lag and zero slowing. I use this as a way of showing that a control system which is properly stabilized behaves (in the steady state) just like a system with no lags. Of course, its dynamics will be different, but when you're interested in an overall view of relationships among variables in a control system, dynamics aren't the main subject.

As to where Ashby went wrong, he didn't use any slowing factor in his equations. Of course, when he set the loop gain to any number greater than -1, the system simply went into ever-increasing oscillations. From that, he concluded that negative feedback can't work with loop gains more negative than -1, and therefore that negative-feedback control must be very weak. Maybe that's why he gave up on the negative feedback model and used an open-loop compensation model instead. Ashby was a psychiatrist, after all. He didn't really know much about control theory.

You note that "in such a simple system one cannot have perfect following even for a constant input (sorry, I mean reference), unless G is infinite." Technically, you're correct. But practically, with a G of 100 or 200, the system will keep errors small enough to be ignored in models of behavior. The actual values measured experimentally for subjects in tracking experiments come out in the range from about 50 to 200. So if the model's G is set too high, it will behave too perfectly. With the correct G, the model will make errors similar to those which the subject makes. We have taken to using an integration factor because, with gains that high, there is no significant difference between a pure integrator and a high-gain proportional system with an appropriate slowing factor. I went through a comparative analysis a few months ago and satisfied myself of that. When you're retired, who else do you have to satisfy? And it's all right if you say "input" here, because in the context we will all recognize that it means "reference input" and not "sensory input."

I *hope* you didn't make any algebraic errors, because the derivations looked fine to me. I don't usually bother with the subscripts, but your use of them is the same as mine when I put them in. When one does most calculations through programming, an equal sign comes to be understood as the replacement operation. Bad habit, no doubt.

Izhak Bar-Kana: I cannot become one of the family on the Net, especially when the discussion becomes philosophical. I can smile when you give me the thermostat example as a living illustration for the control of the input, because under the conditions you describe, I would fire the designer. I am not sure I know where I belong, because I try to get something from everybody, so I try, at least, to read your discussions. One thing I do know: I am an engineer—I would say a bloody old engineer—and cannot change overnight.

For me, there is no reason for existence of any control loop or, better, control system, if it does not control the output. I must use measurements to monitor this output, and if I am wrong, I might end up

controlling something else. But, in the same way, the control signal I design is going to be transmitted through sore actuator, and if I am not careful, it might become very different from the desired control signal. If there is a danger that the input I measure does not represent the output in an acceptable way, I will use lots of filtering (estimation) or lots of redundancy. It might be that in organisms, the emphasis is on the other aspect; I don't know. However, if a simple engineering system includes so many redundant loops, I have the feeling that the extraordinary redundancy in living systems has the same role: to avoid controlling the measured feedback input or responding to a measured reference signal which does not reproduce, in a reasonably exact way, the realworld external values.

But I see I am getting philosophical without even trying. All I wanted to say is that G is not the loop gain, once you use the "slowing factor" S. The loop gain is now given by K = 1- (G+1)/S, from this equation:

$$A(k) = [1-(G+1)/S]A(k-1) + GA^*/S$$

and it *must* be less than 1 to get a stable system. Of course, if K is less than 0, A will change signs every interval, and in a first-order system (with only one delay involved) this oscillation can be prevented using K positive. High gain is a solution when noise is not involved, otherwise the difference between integration and high gain becomes evident: while the high gain amplifies any noise, integration averages it.

Gary, yes, the digital computer is a very easy and handy way to approximate and simulate continuous systems. When the continuous system is sufficiently slow and the sampling is sufficiently fast, one can ignore the difference. In more complex cases, there is an entire theory about how to switch from the continuous to the discrete domain and vice versa. This is not a trivial problem. There are phenomena which cannot be exactly reproduced in the discrete simulation (what happens at collisions, etc.). When one wants only to simulate an approximate behavior, especially in closed-loop, many parameters can change without affecting the results, and any discrete approximation will do.

Advantages of discretization? It is so convenient! Try to implement a slow process with a time constant of, say, 10 minutes. One might need the earth for the capacitor which would be required. In discrete form, it is just a line of code. But most important is implementation of timevarying and nonlinear parameters and algorithms which are almost impossible in analog form. And, by the way, delays. Very fast processes, however, cannot, or cannot yet, be implemented digitally, and analog circuitry has made some progress. So, actually I see a combination of both as the future solution for computation. Complex simulators use "hardware-in-the-loop": those parts which are too fast or cannot be simulated with confidence are used directly in the loop, of course using D/A (digital-to-analog) and A/D converters. This brings us to real-time simulation, which is another opera.

I used to simulate very complex systems, such as planes, flexible structures, etc., with large ranges of time constants. There are simulation languages which allow you to write the equations of the continuous systems. The translation to the discrete world is done by the computer, sometimes using different time intervals for different integrals, so the errors are maintained below some admissible value. In these cases, the precision is almost continuous.

The slowing factor does not make it work like a continuous system, it only makes it work. This is also the danger of simulation, especially when presented to inexperienced students. They take the results for granted, because "the computer shows." But the computer shows exactly what we supply it with. As I understand it, you do not have any detailed models of the various components which together form the simulated closed-loop. In this case, one must emphasize the fact that by using a simple model, one manages to reproduce the behavior of the real thing, to some extent. But not vice versa: the real organism does not behave this way, just because the computer shows.

Bill Powers: Izhak, I talk about the output of the control system, which is the input to the plant (the environment), while you talk about the output of the plant, which is the input to the control system. We don't control the input to the plant—that is varied as disturbances require, so the state of the control system's output is just as unpredictable as the disturbances are. The output of the plant is under control, and so is predictable. That is the same as saying that the input to the control system is under control: the only difference between saying "input" and "output" in that case is whether you take the plant's point of view or the sensor's. You see that I am separating the control system from the plant that is controlled; perhaps you draw the boundaries differently.

In artificial control systems, the engineer can see both the sensor signal and the objective variable to which it corresponds—what you call the output (of the plant). In living control systems, the observer (the one which matters) is inside the system and can see only the sensory input. The variable in the plant (the environment) can only be inferred; it is not available to direct inspection by the control system. This makes a great deal of difference when you are talking about systems which, in effect, design themselves.

In speaking of artificial systems, it is optional whether you consider the controlled variable to be an input or an output variable: it is the same variable in any case, just outside the sensor. In speaking of living

control systems, however, where we must account not only for their operation but for the internal organizing processes which bring them into being, we must choose the "input" interpretation. In fact, we must say that the perceptual signal itself is really the controlled variable, for sensors can vary in their properties.

When the sensor's calibration changes, the perceptual signal remains under control in the same state as before, but the external variable on the other side of the sensor is brought to a new value by the control system. If we understand that the perceptual signal is the controlled variable, then we can understand how the behavior of the system changes when its perceptual systems reorganize. If we focus on the external processes alone, we will see only that something has disturbed the control process, thrown it out of kilter. We might even conclude that it has failed, when all it has done is to change its definition of its environment, possibly by mistake, but also possibly for its own purposes.

So I think that we have to think of control as control of input, if we are to grasp what is really meant by saying that we, ourselves, are control systems.

Izhak Bar-Kana: Bill, you seem to repeat some arguments I was trying to use when I started participating in this Net, and I felt that we used the same names for different things, and vice versa. But then I came to understand from you, and even more so from Rick Marken, that things are much more profound.

I am sorry, but for all of my (engineering) life I have been used to "input controls" and "output is controlled." The "control system" includes everything, and, of course, the plant. The part of the control system which controls the plant is the "controller." Now, if the input to the control system (to the controller, and through it, to the plant) is not zero, it will affect the plant. If it is zero, it will not affect the plant. In a closed-loop system, the input is obtained by the difference between the reference input (in tracking systems, control systems with feedback gain one, it is also the desired output) and the measured output. It is clear that the control signal, the signal which affects the controller and the plant, can be only the measured value of the real signal. Similarly, biological systems can only use the sensorial perception as the *control signal,* to affect their control system, and all of the various stages and values, up to the value which is called the "controlled variable." I really don't understand why this language, which you do know, had to be changed in such a way that Rick Marken cannot even talk to someone I call a control guy, because the old-fashioned engineer cannot accept the idea that any control system controls its input.

Now we are in a closed-loop, and you can again change the order.

I regret that it separates you from the general family of control research. More so, since we want to learn about the behavior of organisms from psychologists. I can ignore the linguistic differences, or at least try to, and try to get the ideas, because I don't know a better group and discussion. But why do you have to speak French in the middle of English?

The most intelligent system I might dream of designing does not come even close to the simplest organism. In my humble opinion, again, one of the reasons for the huge and not always motivated (apparently) redundancy in organisms is to prevent an occasionally wrong measurement (or input feedback) from replacing the correct output the control system is meant to control.

Tom Bourbon: Izhak, you have the respect of those of us who labor to understand *living* control systems. I am certain your life would be simpler were you to decide that we are a bunch of misguided nuts who cannot get our control-system diagrams and labels right!

Perhaps I am wrong, but part of the problem when you speak of engineering (designing and building) a mechanical control system and we speak of trying to describe and explain the control created by the living things we find already acting in the world is that we can't design and know all about the living systems. What is more, the variety of control theory we are trying to develop must compete with a host of already established and widely believed theories and disciplines, so we must direct most of our effort to persuading followers of those disciplines that there is even anything out there to notice which is different from what they already know. It is unfortunate that, in the process, some of what we say seems wrong to the part of the engineering community which is probably closest to us.

Living control systems were not designed by us: we found them inhabiting a world which had already buried them in a host of sciences and disciplines—the life sciences, social sciences, and behavioral sciences, recently joined by the neurosciences, cognitive sciences, and many, many more. For the most part, the practitioners of those disciplines and sciences do not recognize that living systems control *anything.* Rather, they speak of the behavior (actions) of living things as *controlled by* antecedents, whether from the environment (e.g., stimuli, contexts, gods, societies) or from somewhere inside (e.g., mind, soul, schema, plan, commands from the motor cortex). They invoke linear causes, and they reject control by living things.

All a living system knows of "the world" are its own sensory experiences of the world, so it follows that all a living control system can control, from its own perspective, are its own sensory experiences. And there is abundant and conclusive evidence that sensory experiences do

not correspond *directly* with the environment. Perceptions as simple as those of brightness, hue, loudness, heaviness, and the like reflect states of the perceptual apparatus of the organism—*directly*—and they *fail to correspond directly* with any unique state of the environment. The state of adaptation of sensory receptors, the surrounding stimulus field, the relative sizes or magnitudes of different elements of the stimulus field, the relative temporal durations of stimulus elements, and many other variables can combine in many different ways to produce the *same* perceptual experience. Hence, a person, like any other organism, can have the same perceptual experience in the presence of a nearly infinite array of different combinations of elements in the environment and in the organism's own physiology.

Because perception does not correspond one-to-one with any unique state of the environment, it follows that an organism which acts to control its own perceptions is not controlling a unique state of the environment, hence is not producing (controlling) a unique state of its actions (output). The specific actions of the organism and the remote environmental consequences of those actions can vary dramatically, yet the organism experiences uniform perceptions. And it is certainly true that an organism which always produces the same actions and remote consequences in the environment will experience variable, not constant and controlled, perceptions.

In your engineering applications, zero input (by that do you mean zero perceived error—a state *internal* to the organism?) leads to zero output. But an organism which adopts a new reference to experience an absent perception experiences zero perceptual input, which creates in the organism a non-zero error, which drives the behavioral actions (output) of the organism to create the desired perceptual experience, which does not, for a perceiving organism, correspond directly with an objective state of the environment. A bird with a reference to sense a not-yet constructed nest experiences zero perceptual input of nest, and it acts until it experiences that perception. And a sculptor who decides to sculpt a bird on a nest experiences zero perceptual input and acts until that experience exists—whether any other person recognizes the finished sculpture as bird-on-nest, or not. To the artist, that is not important (not even if the artist must sell the sculpture to buy food—all that matters is that someone else desires the sculpture and pays what the artist asks).

Those are the kinds of control we find in the world of living control systems. The best we can do is look for situations in which the variables through which the organism or person achieves its control of perception are also sufficiently stable from *our* perspective that we have a due as to what the organism or person is controlling—in its own perceptions. Certainly, the one we observe is not controlling our experience—not as its primary goal.

Interestingly, it is true that the category of humans known as controlsystem engineers do enjoy a privileged position relative to the control systems they design, construct, and study. They do know the references and the "objective" states of the relevant variables in the environments of those systems. In fact, what the control-system engineer intends is that his or her perceptions of the states of those variables in the environment of the artificial control system will match his or her chosen reference. In that context, it is easy to understand why the engineer would speak of the artificial system controlling *its* output—what the artificial system represents is a way for the designer and builder to control *his* or *her* perceptions, relative to her or his references.

I do not know if any of this helps, Izhak. If anything I say violates too many of your ideas about control processes, please tell me.

Rick Marken: Izhak, the difference between input control and control of input is not just a language difference—it's the whole point.

For an engineering psychologist, the organism experiences error due to the discrepancy between an objective reference and input event (the target and cursor in tracking; sometimes the error itself is considered the stimulus). For a Control Systems Group control theorist, both the error and the reference are *inside* the organism. The reference can be adjusted by the organism (by higher-level control systems) so the organism determines what constitutes an error; the organism is in control of the environment, not vice versa—a rather significant difference. The difference accounts for the appearance that organisms can voluntarily change the value at which an environmental input variable is controlled—it's as though the thermostat suddenly decided to keep the room at 65 rather than 72. This is the phenomenon which control theory is trying to point to: voluntarity or, better, purposefulness.

The controlled environmental variable is probably what you call the output which is controlled by the system. That's fine—but, of course, it is this output *as perceived by the organism* which is controlled, not the output itself. With organisms, there is no independent means of checking the validity of the perceptual representation of the environmental variable being controlled—all the organism gets are perceptions of the environment. We cannot look past our perceptions to see if we are controlling what we intend to control (as you do when you design a control system and make sure that it is controlling what you intend for it to be controlling; you can look beyond the sensors, the control system itself cannot). So, for a living control system, reference states of perceptions (not environmental outputs) *are* the intended ends of control actions.

Note, by the way, that the mathematics of the engineering psychology and CSG psychology approaches to control are nearly the same (at least, control works in both cases). The difference is in where you put the variable r (the reference signal). That's all there is to it. Small step for control theory; giant leap for understanding the nature of living systems.

Wayne Hershberger: Izhak, the yoked terms *cause* and *effect (prod* and *product, independent variable* and *dependent variable,* and *controlling* and *controlled),* respectively, have gotten linked to *input* and *output* in engineering, just as they have gotten linked to *stimulus* and *response* in psychology. This linkage goes way back—long before the development of control theory. Relatively recently, Ben Franklin "mislabeled" the polarity of electrical potentials, and just as engineers continue to use Franklin's terminology as an acceptable convention, even though that convention has misleading connotations, so they can and do use the cause/input and effect/output conventions, even though that convention also has misleading connotations.

Consider the following bizarre statement: The output of a furnacethermostat system is an input, not an output, but this input is an output, not an input. Although this grammatical sentence is *not* nonsense, it is certainly gibberish. Deciphered, the sentence reads: The output of the thermostat-furnace system (i.e., what it produces or does or controls) is a particular value of temperature sensed by a thermocouple (receptor input), not a particular amount of heat emitted by the furnace (output), but this sensed temperature (or receptor input) is an output (i.e., the dependent variable controlled by the system), not an input (i.e., it does not cause the temperature being produced).

As you can see, some truths *cannot* be expressed in engineering's input-output terminology without sounding ridiculous. Conversely, some things which are truly ridiculous can sound very true. As long as input means in and output means *out,* it is unseemly to use those same terms (input and output) also to denote cause and effect (or controlling and controlled). As you can well imagine, when one pairs the term input with cause (or controlling) and the term output with effect (or controlled), many of the connotations are as phony as a three dollar bill. I want no truck with that currency.

I suspect that you are not bothered all that much by these mischievous verbal connotations because you think mathematically most of the time (Franklin's convention poses no problem for me when I use Ohm's law, but it is a veritable pain in the ass when I try to think about PNP versus NPN transistors). Perhaps you can understand why those of us who wish to describe control systems in *English* can ill afford to be encumbered by blind (and blinding) anachronistic language habits developed in the context of a technology devoted to *calibrated* "control” systems (wherein the input/cause and output/effect convention is not problematic).

I am not for a moment saying that control engineers do not know what they are talking about when they say that control systems control their output. What I am saying is that engineers do not realize the mischief they are making for the rest of us when they use the terms input and output in this way. I am speaking for myself, but I think the other CSGers would agree.

Izhak Bar-Kana: I think most of you ignore some of my words.

One problem, for example, is that we here must control the position of a satellite up there. So, the position of the satellite is *the* object of the control. I have no direct means to measure it, as I have no direct means to affect it. Yet, everything in my control problem revolves around this control object. I can only use some remote measuring instrument, but I have no intention to control its output. Actually, one must learn how to ignore some of the input signals, because they are disturbances or noise. No one can convince an engineer to accept the idea that he and his control system control an input. You all seem to accept the idea that there is no control unless there is intention to control. Therefore, the position of the satellite will be monitored by multiple sensors, well-filtered and processed (to estimate some of plant’s states, such as velocities, etc.) in such a way that the controlled variable remains at the value I am interested in, namely the position, or more generally, the motion of the satellite.

If the result of the multiple measuring, filtration, and processing is what you would call perception, I have no argument.

Tom, many thanks, but please do not *respect* or *trust* me. The only principles I trust are: 1) never believe in principles, 2) no one here is God, and 3) no one is really dumb. In spite of the fact that you, Rick, and Wayne seem to agree about the contents of your messages, and I think I have no problem with most of them, there is a difference in tone between your explanation and Rick’s.

There is a big difference between Rick Marken when he clearly presents an idea and Rick Marken in an argument. Sometimes, it seems that the second has had *The Revelation*, or even touched *God*, and to hell with the others. Rick, I am afraid that when you talk about control, you have a steady-state image in the back of your mind. I *must* keep a dynamic image in my mind, because some of my sad experiences show that the steady state might be beautiful, but it is never reached. I must use some mathematics because many great, ingenious, and intuitive ideas proved wrong. It is easy to show that “if this is so, the gain must be so and so, and if the error is so, that lets us adapt the gain to be so and so.” These arguments convince, and engineers like

them. However, after a plane crashes, and one analyses (very difficult math, particularly in nonlinear systems) the aftermath, one discovers that things became unstable just because one used non-constant gains, even very carefully, namely within the “admissible” bounds which were tested with constant gains.

When I control the motion of my hand, this is the intention of the control, the object of the control, and I think that the corresponding control system controls the position of the motion of the hand. If you agree that there is intention in control, then this is the only intention. The fact that I must *measure, sense,* or *observe* this motion is a problem, not a principle. As you have observed lately, a closed-loop system is so built (if correctly built) so that the gain of reference-to-output is almost *one,* while the gain of disturbance-to-output is almost zero. We can show that an integration in the forward path makes the corresponding ratios actually 1 and 0, at least for constant reference inputs.

I try to stay aside when psychology is talked here. If you consider that the reference inside the organism is a great idea, showing that life affects the environment, and not vice versa, I am excited by that idea. Unfortunately, I cannot claim the same thing about the artificial control loops. I mean the environment makes a lot of trouble. I only have my own organism to observe, and personally I think that I can decide to drive here or there only after a long period of learning, and that the reference points within are a good mapping of the reference points out there. I might decide just to follow an internal reference with no relation to the outside word, but usually I stop after the right number of glasses. Furthermore, the reference command to be followed I would rather call *decision* than *control.*

If I want to move my hand, or a robot arm, I apply a force. If I meet resistance, I use more force. If there is an egg there, I must behave in a different way, and this is first of all a decision problem, or a detection problem. I don’t call everything control theory. Many control people (engineers or not) do not know detection or decision theory, and they have to rediscover it again and again when it is needed—not the best way.

Wayne: Time flies like an arrow, fruit flies like a banana. I admit that this is not related to our topic, and I admit that what you call gibberish is gibberish indeed, but I cannot see how you relate it to any engineering. Still, I am happy you do not blame the Original Sin on me. I will never tell anyone he is wrong because his arguments remind me .... I can argue about right or wrong, understanding or not understanding, and so on. Too many arguments here blame me for talking like the behaviorists, like Wiener, calibrationists, and who knows what. I think you have better arguments for your position than calling the control of plants a simple misuse of words, even if control of perception is needed and correct under your paradigm.

What an engineer means is that his system controls the position of the plane, and he calls it an output. It has nothing to do with any old-fashioned calibration, as it has nothing to do with the Middle Ages. He will do anything that is needed, and possible, to make this plane follow the desired path. It has nothing to do with grammar, nor with the fact than any input is an output of something else and vice versa. But I think we rotate now in a closed-loop with no reference whatsoever.

I first thought it was worth understanding your terminology and to bring it to some common denominator with the large family of control theory, but it is not very important. When I say "control," I have a plane or a robot in my mind, not a differential equation. My only problem is that this robot should be at a given position at a given time, no matter how I monitor (sense) its motion.

When you see these lines on the display, the desired output of so many control loops, all designed to satisfy your finest perceptions, even if they cannot control (or because they cannot control) your perceptions, if you can claim that an engineer does not care about what is input or output, I can only ask: Who do you call engineer?

Bill Powers: Izhak, it isn't that we ignore your words: it's that we can accept them as truth, but truth of a kind leaving out other important truths, particularly the one we have found the most startling and the most informative: the truth that a control system can control only what it senses.

In the world of engineering, the engineer has full knowledge of both the environment of a control system and the internal design of the control system. So he can point to a consequence of the system's actions and say, "There, that's the output which I want to be controlled, and here, in the system, is the feedback signal which represents that output." In doing this, he does not have to pay any attention to the fact that he must use his own senses to see that output. Literally, however, *for the engineer,* the output being controlled is known only in the form of a perception (whether aided by instruments or not). That is a fact, but it is irrelevant in engineering.

It is not irrelevant in trying to understand how the engineer works. When we look at the design of the engineer himself, according to our best neurological and physiological models, we can see that the engineer's entire world must exist in the form of sensory signals and higher-order functions of them, also represented as signals. In a way, you have given a nice example of this in talking about controlling a satellite: "So, the position of the satellite is *the* object of the control. I have no direct means to measure it, as I have no direct means to affect it. Yet, everything in my control problem revolves around this control object.

I can only use some remote measuring instrument, but I have no intention to control its output." So how does the engineer know of the position of the satellite (other than by looking up)? Only, as you say, by using a remote measuring instrument. He has some moderate amount of faith in the instrument, after calibrating it, but that does not change the fact that *all* he knows of the satellite's position is in the form of this instrument's reading (which consists of numerical digits, not positions). He does not, in fact, know the position itself. He knows only the reading, and he has a complex theoretical structure in his head which converts this reading into a concept called "position." He calls this concept the "output" he wants to control, but in strict literal truth, it is a perception.

The engineer might have no intention of controlling the output of the measuring instrument, but in fact that is all he can control. He has no other way of knowing the position of the satellite except through the use of earthbound measuring instruments (his eyes among them). He trusts that the instrument readings correspond in a regular fashion to the "actual position" of the satellite (with all necessary corrections applied, for example the time-lag of light rays and radar pulses, and the various motions of the earth itself). This trust is an epistemological statement, but its truth or falsity do not matter here: we are talking about practical requirements. The engineer *imagines* that he is controlling the position of a real satellite, up there in the sky, and he can produce all sorts of justifications for accepting this imagined correspondence. But he can't know that position without using the instruments, and he can't know the effect of his remote-control actions until he sees what the instrument readings do. Whatever he believes is actually going on, he is stuck, as a practical matter, with controlling a perceptual representation, not the thing itself. His epistemological beliefs make no practical difference at all.

All animals, and most human beings prior to the age of higher learning, necessarily act from the epistemological position that the perceived world is the world itself. There is, of course, no alternative. Speaking for human beings, the reality we know as solid and real, upon which we act and which we intentionally alter in some regards, is the only world there is. "Perceptions" don't exist except, for some of us, as philosophical abstractions or "signals" in a model. When we forget about models and philosophy and just look around, we see the world, not perceptions. When we look up into the night sky and see that serene and untwinkling point of light moving steadily and silently among the stars, we say, "Look at that! There's a satellite!" We don't say, "This is a perception of a satellite, a signal in my brain." When we point at the satellite, we see our own hands with forefinger extended. We don't say that there is a perception of a forefinger, nor do we pause to wonder about the relationship we call "extended." The relationship is just as much "out there" as the finger is. It would never occur to us to wonder what sort of thing it is, out there, which we call a "relationship."

This simple and self-evident world has conceptual holes in it. The biggest hole lies between the intention of pointing at the satellite and the immediately experienced actuality of pointing at the satellite. How is it that a mere intention, a figment of the mind, actually causes this pointing to occur? Control theory provides a plausible way to fill in this gap, a way which is as self-consistent and as consistent with observation as any finding of physics. In a manner of speaking, it is a finding of physics. To construct this model, however, we are forced to readjust our conception of the whole apparent reality, because the control model can work only if the satellite and the finger exist for the controlling system as signals produced by sensory inputs and subsequent computing functions in the brain.

In your objections to the concept of control of input, you have consistently assumed that the engineer can know the actual state of the output. Within the boundaries of the usual world of engineering, observing is not a brain process: it simply consists of noting what is there, while the role of the engineer's brain in making this possible remains silently in the background. In our explorations of control theory, however, we make this brain-in-the-background explicit. Even in talking about artificial control systems, we habitually take the point of view of the control system, something which few engineers would see any reason to do. We say, "If I were that control system, what world would I be experiencing?" And the answer, of course, is that the world would consist completely and exclusively of the signals coming out of the sensors. We could not know what is causing those signals; we could not even know whether they represented light or magnetism or sound. They are just signals. They get fancier labels only in the context of other signals which are also just signals—or in the mind of the engineer, who occupies an omniscient position in relation to this tiny control system and its surroundings.

When I speak of what "we" think on this subject, I am speaking of those who have internalized this model to the extent of relabeling their own ordinary experiences as "perceptions," at least when thinking in the modeling mode. This relabeling has come to most of "us" in a moment of sudden illumination which forever alters how we understand nature and ourselves. Nothing is changed in ordinary experience: "out there" still seems to be where it has always been. What changes is its meaning in relation to how we interpret the behavior of others and ourselves. This threshold of understanding is either passed or it is not. Once it is passed, the world of experience not only contains new implications, but *it makes a great deal more sense than it made before.* This is what has attracted so many people to the CSG version of control theory in the context of living systems. So many questions are answered, even those we hadn't thought of asking. So many holes are plugged which we hadn't even recognized as holes.

"Reality Therapy" and "Experiential Therapy"

Rick Marken: To Dag Forssell: I have your "Alignment/ Mission Statement" and "Discussion of Issues and Control Theory." The first seems to be a template for a statement of agreed-on higher-order goals for two control systems (people) working as partners in an engineering firm. It looks OK to me. I have a bit of trouble with phrases like "accept responsibility for our lives" and "efficient perception of reality." I also think the statement that control theory views people as controlling themselves misses the point by enough to be misleading.

Your "efficient perception of reality" statement makes me wonder—what' did the guy who developed "reality therapy" see in a model of behavior as the control of *perception?* Is the idea of reality therapy that *reality is perception?* If so, why use the term reality? It suggests a therapy which helps you get in touch with reality, which suggests that the therapist knows what reality is, and you (the therapee) should too. If I understood behavior as the control of perception, and problems requiring therapy as the results of perceiving things in ways which prevent non-conflicting control of those results, I would never have thought of calling my therapy "reality therapy." Maybe "control therapy" or "perceptual reconciliation therapy;" or, best, "conflict resolution therapy." But "reality therapy"? What could be more misleading? Clinicians, could you tell me why William Glasser, who claims to have understood control theory before he even discovered it, called his approach to therapy "reality therapy"? I smell condescension here.

Your discussion of control theory seems reasonable. It does emphasize the control of perception. I would suggest that you make clear the relevance of perceptual control to the problem of conflict and how to resolve it. After all, I think that's what the value of control theory is for effective management: finding ways to perceive the production process so that there will be minimal or no conflict between the cooperating contributors to the process.

Dag Forssell: Rick, thanks for reading my papers. I am glad you did not find any major misstatement on my part. I cannot ask for any thorough critique at this stage, since my presentation is not finished. The particular papers were extracts and summaries, respectively. I am trying to introduce control theory to industry in a fruitful way.

When you came across the word "reality," your configuration references immediately associated with "therapy" and "Glasser." I am *very* glad that I found the Control Systems Group, even if it was by way of Glasser's writings. I am not a student of Glasser any more (I can still see value in many of his writings, both from a medical perspective and in the more recent musings on quality, but he is dangerous because he totally misrepresents—*grossly* oversimplifies, apparently because it reduces his error signals—control theory and what one can learn from it).

Anyhow, upon closer scrutiny, you will note that the word "therapy" is nowhere to be found. I made reference to Abraham Maslow's admittedly unscientific observation that the most outstanding common denominator in people of a high level of mental health is "more efficient perception of reality and more comfortable relations with it."

I am excited about control theory precisely because the model offers "a more efficient perception of reality"—the way the world works, and we with it. In quoting Maslow, I am not trying to adhere religiously to any standard, rigorously defined control-theory terminology, if there is one. It seems to me that in addressing a larger public, I must find a way to use terms they relate to. So far in my attempts, I try to use as many synonyms and analogies as I can find.

Rick Marken: Dag, my comments about reality therapy and Glasser were not directed at you at all—nor were they meant as a criticism of your work. It just jogged a thought in my mind which I wanted to make public—about Glasser's interest in control theory. I've wondered why his therapy is called reality therapy if he is such a fan of control theory. I was asking the Net at large; there are a number of therapists out there, and some are familiar with reality therapy.

Again, I request info on this topic from therapists. I really am curious about it.

Dag Forssell: Rick, now that I understand your comment as a question, I shall attempt an answer.

Very briefly, Bill Glasser used/developed reality therapy 30 years back. His book by the same name is still available in your local book store. He developed an institute and a large following, numbering in the thousands. He has written many books which show his deep interest in matters human (see *Positive Addiction* and *The Identity Society).* He was told about our 1973 "bible" and attempted to write a version more accessible to the public.

I was fascinated by Glasser's *Stations of the Mind,* but then, I am an engineer. The book does a credible job, as far as I remember. He gives

proper credit to Bill Powers. Many of Glasser's senior faculty still go by that book, which is why some are in the Control Systems Group. The book was probably not a hit with the public, and apparently not with most of Bill Powers' followers. Glasser developed a four-color chart to teach by, which is simplified but not bad.

Clearly, reality therapy came first. Control theory failed to support it as Glasser anticipated, since he could not teach it in a way his audience accepted. Problems of organizational control might have contributed to the break with Powers.

It seems to me that Bill Glasser is smart and has made contributions in many ways. Reality therapy is his baby and his dominant systems concept. It comes first. Glasser's book *Control Theory* provides the following definition of control theory: control theory contends that every behavior is a person's best attempt to meet his needs. Perception went out the window because it was confusing to his audience.

This is a quick sketch of my perceptions on this. We all have different contents in our systems concepts. Glasser's priorities are different from ours. Still, he has brought a number of control-theory faithfuls to our fold through his promotion efforts. I am glad that I am one of them.

Rick Marken: Thanks for the thumbnail sketch of the history of reality therapy and its relationship to the Control Systems Group. Actually, I am fairly familiar with that history. I went to a Bill Glasser show in about 1981 when I was in Minnesota. When I found out he was interested in Powers, I went right up to Bemidji or wherever he was. I even had lunch with him in the regal dining room—he invited me in when I told him I had been working on control theory for a couple of years already. He struck me as a consummately self-absorbed individual, not at all intellectually interesting. I still don't really understand the basis for the rift (if that's what it was) between him and Powers, though I would guess that it had much more to do with Glasser's rather shallow grasp of Powers' model than with any conflict over control of any organization (the very notion of Bill Powers trying to control some organization is pretty silly, given what Bill Powers is like).

Anyway, I still don't understand what "reality" therapy has to do with control of perception. I kind of ragged on Ed Ford for not explaining to my satisfaction why it is important to realize that people are controlling perceptions, but I think I react so strongly because it is so important to me that it be made clear. *Nobody* has direct access to reality. We control only representations of what reality might be. To the extent people can see agreement regarding what they perceive, we tend to call those perceptions reality—but they ain't. They are just (somewhat) shared perceptions. Indeed, I think using the term "reality" in a therapeutic situation could actually be dangerous—giving the therapee the impression that there is a right way to perceive things. My pedagogical point is that, when explaining control, just leave out the term perception *at first.* I control the letters on the screen, the position of my hands, etc., etc. Once a person understands that there are variables "out there" which are controlled (many different letters could be typed; many different hand positions are possible) and that these variables are brought to reference levels though the action of lower-order acts which could also counter the effects of disturbances to these variables, the person understands the phenomenon of control. Then you can explain that it is perceptual aspects of experience which are being controlled, and that there are, therefore, different perceptual aspects of the same experience which can be controlled.

Again I ask—really, just out of curiosity, not hostility—why would the person who developed "reality therapy" see perceptual control theory as something which would support his theory? And what is "reality therapy," anyway?

Ed Ford: For 14 years, I was a faculty and certifying member of the Institute for Reality Therapy. I taught at every Intensive and Certifying Week in L.A. (where most were held) for over 10 years. I was very close to Glasser. I left the Institute in 1983. Glasser once remarked (somewhat in jest) that he would have called what he did *The* Therapy, but he might have run into problems with others. He called it reality therapy because it best described what he was trying to do, which was to get people to deal with the reality of their present life. It was the most efficient therapy I knew at the time.

When I was introduced to control theory by Glasser, some of what he said didn't make sense (such as don't deal with perceptions, leave that to the theorists). At the 1989 convention mentioned below, Glasser said, when talking about perceptions, "they say it is a hierarchy and you always start out with this one and then get to this one and this one. I don't recognize 10 [levels]. I don't get involved with it. In terms of them [CSGers], it's a fundamental difference." Thus his total disregard for the hierarchy of perceptions (which he wrote about in *Stations of the Mind,* then obviously abandoned). I then left him and became a pupil of Bill Powers. When, after many years, I had finally begun to understand what control theory was all about, I realized that we control perceptions. Glasser never has.

I heard a tape recording of a workshop on what others are thinking and saying about control theory, given by Diane Gossen at the Institute's 1989 convention, which Bill Glasser attended. During the presentation, Glasser kept making comments and corrections to what Diane was saying. His degree of understanding of control theory was very evident.

When someone at the conference asked Glasser when does output become input, he replied, "the only way that the behaving organism becomes aware of the behavior is through its ability to perceive, which is input .... You can go through all kinds of outputs all the time, but what they [CSGers] are saying is that the only time you become aware of it [Glasser's idea of behavior] is through input." Glasser never has understood the concept of the controlled variable, and that we control for input. He sees "control of perception" to mean that when we perceive what we are doing, we are controlling our perception of our behavior. This is how he understands the title of Bill's book, *Behavior: The Control Of Perception.* He has never gotten away from the fact that we control our behavior. Remember, behavior to him is output, behavior to us is the entire system in operation.

He doesn't understand the levels of perception. Glasser said, "the reason I got rid of the levels of perception is when I started to teach you could adjust the levels, and I don't think you can. I think it is all the way through the top. You adjust to what the ultimate picture is... that's what drives you. If you think you can dissect your behavior... I think that is absolutely impossible." For Glasser, reference conditions are called "pictures in your head." The picture for him is the entire hierarchy of perception, and that is all you can control. Obviously, the entire system is engaged in the operation, but he doesn't believe you can be aware at any one level. Another major problem is that he uses "picture in your head" interchangeably as both perception and reference level. He doesn't see the comparison going on inside the head (between perception and reference condition), but rather between the picture in the head and what he calls behavior (and what we call our actions). At one point, he uses the picture-in-your-head concept as building a perception; at the next, as something you want or a reference condition. The same word is used for two entirely different concepts. The bottom line is that Glasser has never, never gotten away, from controlling output. For him, the comparison is between what we want (which he calls the picture in your head) and what we are presently doing to get what we want (our actions or what we are doing which we call output).

Other areas of misunderstanding by Glasser: He says his idea of needs is what CSGers would call disturbances. Obviously, he doesn't understand disturbance, because he doesn't understand the concept of controlled variable. Glasser sees the reasons for disturbances occurring as the basic needs. And there you have another major problem, the concept of needs. There are basic needs such as the need for water, food, etc. Where it gets tricky is when you get to such needs which Glasser identifies as Power, Fun, Freedom, Belonging. Our genetic system sets the limits on basic needs. But when it comes to those areas through which all of us strive to find satisfaction, they can be seen quite differently by each of us. I really struggled with this idea (with a great deal of help from Bill Powers) in Chapter 7 of my book, *Freedom from Stress.* It seems to me that we set the limits and parameters of satisfaction within our own hierarchy, especially at the higher levels. This setting of limits is really our individual mark or standard for areas of importance to us, what Glasser would call needs. I think Glasser's higher-order "needs" say more about him than anything else. People define their own internal goals and areas of satisfaction, and from my daily reading of this Net, they surely vary a lot within our own Group.

Other areas, which I don't want to dwell on, include the following: He calls feelings (along with doing and thinking) behavior (remember, his definition of behavior is output). His understanding of reorganization is also very confused. He retitled his book *Take Effective Control of Your Life,* calling it *Control Theory.* The sad thing is that he has taken the name control theory and assumed total control of what it means.

Bill Glasser taught me more than anyone else a lot of great techniques for counselling and dealing with others. Unfortunately, when he was exposed to control theory, he changed control theory to suit his own perception of the world, and to suit reality therapy. Over the past eight or 10 years, as I have been learning control theory, I have tried to take my ideas of reality therapy and other ideas in counselling, and adjust them to the new and very different world of control theory. Control theory has opened a whole new understanding of the world to me, and thanks to Bill Powers, Tom, Rick, and the rest of you, I have been able to use control theory to learn and grow as a counselor, father, teacher, husband, and all the other hats that I wear.

The basic tenets of reality therapy are the steps (get involved, ask what do you want?, what are you doing?, and is what you are doing getting you what you want?; then get a plan and commitment). It also says that whatever we are doing is our attempt to satisfy our needs. I have been able to use reality therapy as a jumping board from which to develop a control-theory therapy.

David Goldstein: Many years ago, I attended a workshop which Glasser gave about the time he started to publish books about control theory. I asked him the very question which Rick raises. From his answer, I received the strong impression that as a result of studying control theory, he revised his attitude from emphasizing a more objective view to a more subjective view of perception. Ed Ford and others who started out in the reality therapy camp could probably tell us more about it. It seemed that he found the questioning a little discomforting. I will leave it to Ed and others in the Control Systems Group to tell us how reality therapy compares to control-theory therapy. I am writing

to tell about "experiential therapy."

I have recently read a book by Alvin R Maher, called *Experiential Psychotherapy: Basic Practices.* I would highly recommend this book to other control-theory clinicians. While the theory of human beings behind experiential therapy is not control theory, but rather a form of existential-humanistic theory, I think that much of the practice is very consistent with control-theory therapy. For example, the method of levels is there! And an effort to describe a particular person's control-system hierarchy is there!

The meaning of the word experience in experiential psychotherapy is very much like the meaning of perception in control theory. I am going to focus on the ways in which experiential psychotherapy can contribute to the practice of control-theory therapy, rather than vice versa. I do see control theory as making a contribution to experiential psychotherapy in many ways, but I will save that for another time and place.

To start with, in experiential psychotherapy, the selection of topics is based on the intensity of experience which the patient experiences when attending to the topic. From the point of view of control theory, this makes good sense. Life areas in which a patient experiences a lot of or only a little bit of satisfaction are likely to be ones about which a person has strong experiences.

While they seem to select topics in similar way, experiential psychotherapy and control-theory therapy differ in the way a topic is pursued. It is here, I believe, that experiential psychotherapy can teach us. The therapist attempts to share the patient's experience as much as possible, including the bodily experiences which go with discussing a topic. Control-theory clinicians do not do this, as far as I know; they do not try to generate the patient's experiences within themselves.

The therapist attempts to have the same experiences as the patient is describing. This is called the "working level" of experiencing for the therapist. When this is not occurring, the therapist and patient stop and go back to the point at which the therapist lost touch with the patient. The sameness of experiences includes bodily experiences!

Maher provides some specific methods for helping a person achieve a strong experience, for helping a person become aware of the higher-level experiences behind the one being discussed, for helping a person be/behave in a way consistent with the higher-level experience, and for helping a person extend the changes in being/behaving to the patient's everyday life. These specific methods are very helpful, are based on years of clinical practice, and I see no reason for not using them.

In closing, I wish to emphasize that experiential psychotherapy makes me rethink the way I am using feeling/mood experiences in the control-theory therapy which I do. I no longer simply try to intellectually figure out (imagine) what the patient's blocked desire is, which is generating the feelings/moods. I try to experience the feelings/ moods and the topic being described by the patient. I then use the experiences which are occurring within me, which are not explicitly described by the patient, to experientially figure out what the patient's higher level perceptions are. In short, experiential psychotherapy is like control-theory therapy, but with more feeling for both the patient and the therapist.

Now, back to reality therapy. While I am not a trained reality therapist, I have watched Ed Ford, Diane Gossen, and Perry Good work at the CSG meetings. And I have read their books. It is my impression that feelings/moods, experiences, and expressions also play a secondary role in the therapy which they do. Reality therapy plugs into the last step in experiential psychotherapy.

Rick Marken: Thanks to Ed and David for the info on reality therapy. I think I get why Glasser got interested—just a poor choice of words, that "reality" thing.

Bill Powers: David, I presume that you remember a discussion in which we talked about devising "qualitative models" for kinds of behavior which are difficult to handle in a purely quantitative way on a computer. Your description of Maher's *Experiential Psychotherapy* strongly suggests that it is a way of doing exactly what we were talking about. The "computer" in which you run your simulation of the other person is, of course, yourself. This living computer already contains the capacity to carry out, in imagination, all of the functions of a human being (oddly enough) at exactly those levels of functioning which actually exist. There is no programming problem—we don't have to figure out how relationship-perception works, or program-perception, or system-concept-perception. The computer is guaranteed to be able to run any process at any level which is required. It is also guaranteed to contain exactly the levels which are required, not skipping any and not adding any which don't belong in a model of an adult human being.

As I read your description, I was reminded of the problem of listening to someone who is giving you directions. When I hear the directions, I try to imagine the actual trip being described. Good directions give you a picture vivid enough so you don't have to write anything down—when you actually follow them, it's as if you've already been there. Poor directions, on the other hand, are full of skips and jumps, private associations and incidental anecdotes; they convey a shifting point of view, sometimes from the viewpoint of the one taking the journey, and sometimes as if from an aerial map or the position of a by-stander.

When you try to follow poor directions in your imagination, you get

a picture of a very confused mind. You don't, of course, actually sense the other person's confusion. But by trying to imagine following the directions, you become confused yourself. That is, when you try to run the model the other person is describing, it leads you to see gaps and contradictions and other problems which leave you confused and, by implication, indicate at least a similar kind of confusion in the other person.

I think that this method can be refined by a control theorist into something even more workable than it already is, and also that it can tell us a lot about the role of language in the control-theory model. Language, in the broadest sense of communicating through manipulation of perceptions, is the medium through which one person tries to convey his or her experiences to another. It undoubtedly has limitations—there might be inherent difficulties in trying to communicate principles and system concepts by any means but demonstration, and in trying to communicate very low-level perceptions, such as the way a face looks or the way ice cream tastes. There are problems inherent in private meanings of words. But in an intimate and protracted relationship with one other person, a therapist should be able to cross-check the meanings and put together a quite reasonable model of the other person, through imagination.

This means that the therapist must become an utterly flexible general-purpose simulation device without cultural biases and with no blind spots—a selfless person. At least during the process of therapy. You would not want to simulate every person as if he or she were, for example, a middle-class Jew or an eccentric engineer. You have to allow the properties of the other person, as nearly as they can be communicated, to enter into yourself and to operate as if they were yourself. It seems to me that doing this would amount to a discipline at least as rigorous as that which the Zen masters demand, at least as deep as the analysis which psychoanalysts are required to go through before they are considered ready to treat other people's problems.

Essential books on control-system models of organismic behavior:

LIVING CONTROL SYSTEMS

Selected Papers of William T. Powers

ISBN 0-9624154-0-5, 1989, 300 pages, illustrated, softcover, 5.5" x 8", $16.50 postpaid (KY residents please add sales tax).

Gathered in this volume are 14 previously published papers by William T. Powers, including several which are now difficult to obtain elsewhere. Ranging from two seminal 1960 articles introducing "A General Feedback Theory of Human Behavior" to a recent overview of biological control theory and its relationship to other ideas in cybernetics, the papers in this collection provide a thorough introduction to Powers' models of living control systems.

Introduction to Modern Psychology

The Control-Theory View

Edited by Richard J. Robertson and William T. Powers

ISBN 0-9624154-1-3, 1990, 238 pages, illustrated, softcover, 8.5" x 11", $25.00 postpaid (KY residents please add sales tax).

Here is the first textbook using the control-theory model for *organismic behavior as control of perception via hierarchically arranged negative-feedback loops.* It reviews and reinterprets many facts found by researchers working within the framework of older traditions in psychology, providing what is lacking in other general psychology texts: a unified approach to the entire field, from laboratory studies of animal behavior, through ethology and studies of human behavior, to clinical work.

Order from: CSG Book Publishing, 460 Black Lick Rd., Gravel Switch, KY 40328 U.S.A.

Back cover

Closed Loop

Threads from CSGnet

Winter 1992

Volume 2

Number 1

Front cover

Closed Loop

Threads from CSGNet

Winter 1992 Volume 2 Number 1

Edited by Greg Williams, 460 Black Lick Rd., Gravel Switch, KY 40328

CONTENTS

Members of the Control Systems Group receive *Closed Loop* quarterly. For membership information, contact Ed Ford, 10209 N. 56th St., Scottsdale, AZ 85253; phone (602)991-4860.

CSGNet, the electronic mail network for individuals interested in control theory as applied to living systems, is a lively forum for sharing ideas, asking questions, and learning more about the theory, its implications, and its problems. The "threads" in each issue of *Closed Loop*, stitched together from some of the Net's many ongoing conversations, exemplify the rich interchanges among Netters.

There are no sign-up or connect-time charges for participation on CSGNet. The Bitnet address is "CSG-L@UIUCVMD" (use no quotes in this and the following addresses); "CSG-L@VMD.CSO.UIUC.EDU" is the Internet address. Messages sent to CSGNet via these addresses are forwarded automatically to all participants. Via CompuServe, use the address "> INTERNET :CSG-L@VMD .CSO.UIUC.VMD" to reach the Net. Initially, you should send a note to the network manager, Gary Cziko, at "G-CZIKO@UIUC.EDU" (Internet) or at "c CZIKO@UIUCVMD" (Bitnet); Gary's voice phone number is (217)333-4382.

Social Control

Rick Marken: I want to eat. The company I work for is willing to give me money so I can eat. I understand that my role is to "work" for them—where "work" can mean spending my time doing something which I prefer to do less than other things. I'm willing to make this exchange—the company "controls" what I do, and I control the amount of money I get. It works because, so far, we are both willing to accept a little error—I don't get nearly as much money as I want, and the company probably doesn't get all the work it would like to get out of me. But we're both happy.

Control theory just says that, when you deal with a person, you are dealing with a control system. The result of that "dealing" depends on how you deal with the control system and what the control system's current configuration is. But if you try to control the control system "arbitrarily" (that is, without taking into account its purposes), there is a good chance of conflict. For example, if the company decides that it will only pay me if I work in a certain way, and if it's the only company in town, and I have no alternative means of getting money, then there are likely to be problems if, for some reason, I don't want to work in that particular way. If the company tries to control me—meaning that it will only accept a particular kind of behavior on my part—and if that behavior is something I don't want to do, then there is conflict.

Most people deal with other people as people—they act as though they understand that the other person is a control system, and they respect that. You get into problems with very "purposeful" people, who have to have people behaving in just a certain way—no attempts at cooperation. These people treat people as objects. When I control a hammer, I want it to do exactly what I want it to do. I don't want to compromise and say, "Well, if you want to land closer to my thumb, then it's OK with me; I understand that you have needs, too." I'm not saying that because the hammer has no needs or wants, and I can control it perfectly, we will never have conflicts. But if I act the same way with my daughter, son, or wife, I am probably looking at significant conflict.

Winning a conflict looks like successful control by the person who wins. But I argue that it is a fool's paradise. The winner then imagines that control can always be achieved by force (not true), and the loser never really goes away.

People want control, and, in particular, they want control of other people. It is a lot easier to control a stimulus-response device than a control system. So people are willing to accept a view of themselves as stimulus-response machines (which couldn't possibly do what they want to do: control) as long as other people are also stimulus-response machines which can be controlled by force. Some stimulus-response machines think that they just happen to emit better behavior than others.

Gary Cziko: The recent death of B.F. Skinner, together with my new interest in control theory, motivated me to read and reread some of Skinner's writings. They look quite different since I have gained some understanding of control theory. Despite the assumption in *Walden Two* that human behavior can be engineered and controlled, I discovered almost in spite of myself that I found the community quite an attractive place. Could control theory be used as Skinner used operant conditioning to create such a place? Or does control theory instead show us that such an utterly conflict-free community is an impossibility?

Bill Powers: Most people who watch behavior closely notice that control is going on. Skinner noticed it too. But he would have said that a baby turning red and crying when its bath temperature goes too high is under the stimulus control of the temperature. Control theory says *almost* the same thing: the baby's behavior is driven by the difference between the actual temperature and the temperature the baby wants to experience. But Skinner wouldn't have liked that proposition, because it invokes a causal factor inside the baby: the definition of the right temperature, which is determined by the baby and not by the environment. Control theory says that the baby's internal specification for the right temperature determines the stimulus value of any given temperature. If the specification changes (the baby develops a fever), the same external temperature which was satisfactory before is now "too cold." The baby acts as if the temperature has dropped, and it won't be satisfied until somebody lets it get warmer. That's why we shiver and burrow into the blankets when we develop a fever: the reference temperature has increased.

Skinner described control behavior. He explained it as environmental control. If you just ignore all of Skinner's *explanations* of behavior, I suppose you could say he wasn't a bad observer.

Rick Marken: Any parent knows how difficult it is to "debug" a child; about all you can tell from the child's behavior is that something is "wrong." You try to figure out which variables should be returned to their reference levels. This is by no means an easy process. When the child continues to reorganize (cry, squirm, etc.), you are likely to become frustrated. Skinner made it sound a lot easier than it is to "control a variable" for a child. But he was right about one thing: when you do get all of the variables to their reference states, the baby becomes quiescent. Thus, Skinner did understand the idea that behavior is error-actuated, and that you can determine the reference level of a controlled variable by looking for the level of the variable which produces no efforts to change it.

I have found a couple of Skinner quotes suggesting that he understood something about control. For example, he put a little section on control in *About Behaviorism*, where he actually said something like "to behave is to control." After all, behavior produces consequences (reinforcements), and these often look like the ends towards which behavior is done (they are, but not according to Skinner). Skinner did seem to recognize controlling as a kind of behavior. It is what behaviorists do, for example. In *Beyond Freedom and Dignity*, he talked about the behaviorist who trains a pigeon by doing a behavior called "controlling." "Controlling" is controlled by the behavior of the pigeon (which, I suppose, is doing a behavior called "being trained"). So there is "reciprocal control." Clearly, Skinner's idea of what it means to control is pretty wimpy. When I control something, I know how I want it to be, and, if I can, I get it to be that way. The thing I am controlling has no say in the matter. If it does, then I am in a conflict with it. I lift my glass to precisely the level I want it to be. If the glass is also controlling me, then it is possible that the glass wants me to put it somewhere other than where I want to put it. So far, I have been very successful at placing glasses where I want them, and somewhat less successful at putting control systems (like my cat) where I want them.

Reciprocal control is a crazy notion. Control theory shows that there can be no such thing, except in special cases where the two systems are either actively trying to cooperate, or where they are controlling variables which are not in conflict—as when an experimenter controls the pecking rate of a pigeon while the pigeon controls the amount of food it gets. Either of these special cases could end up in conflict anyway; one member of the cooperating pair might feel that the situation is unfair, and the pigeon might not be able to generate the rates demanded by the trainer and just stop, leading to error (and ugly corrective action by the trainer).

The attractiveness of Walden II comes from the appearance of lack of coercion; everyone gets rewarded for "good behavior," and there is no punishment for "bad behavior." Bill Powers wrote a lovely letter to Skinner which was published in *Science*. In that letter, Bill explained the problem with this "non-coercive" approach: it works as long as the behaviors which the community rewards are the behaviors you want to produce (assuming that all want the rewards—for simplicity, we'll assume all do, but that is another problem). However, what a perso wants to do is determined by his or her internal structure of intrinsic needs and purposes which have been learned to keep those needs satisfied. The problem with Walden H is that nobody can determine what someone else "should" do, even when the "should dos" are for the person's and the community's own good. That's perhaps the downfall of every well-meaning attempt to create a perfect society.

There is no doubt that Walden II might work for those who want it to work, and who are willing to live in the context of the community's rules. Skinner himself didn't choose to live in that society (a community built on Skinner's principles still exists somewhere in the East, but Skinner himself didn't join when asked—he was controlling for other variables). What Walden II shows is that coercion can be masked quite well by good intentions. I find Walden LI a hell of a lot scarier then some repressive dictatorships where the coercion is at least up-front and the hypocrisy is transparent ("this is for your own good").

Bill Powers: Rick said it right: Walden II works because everybody wants the rewards used to keep the society in line, and everyone works (funny thing) exactly as Skinner thinks they will. The real attempt to form a community of this sort didn't run so smoothly: lots of coercion. The problem is that you can't reward somebody who knows how to get the reinforcer without anyone's permission. So you have to make sure you're the only (or at least the easiest) source, and to maintain the behavior, you have to be willing to leap out of bed with a tray full of reinforcers whenever the person you are controlling this way does something right. I'll bet that isn't what Skinner had in mind.

I'm not enthusiastic about demonstration communities. They will work as long as everyone consciously tries to work the way the theory says things should work. Sooner or later human nature breaks up the act. This would be true even of control theorists (especially?). I think the community we need to form is already around us. If we can't help that community to shape up, we wouldn't do much better in an ashram.

Tom Bourbon: Some of the comments on Skinner and Walden II gave what might be a mistaken impression about the community founded on the principles in Skinner's book. The community is Twin Oaks, near Louisa, Virginia. It was founded in 1967. From the early 1980s until as recently as 1986, I corresponded with a young man who lived there. He was the brother of a student here, and we had several opportunities to visit.

From the beginning, I was surprised to learn that Twin Oaks was still there. I had assumed that it died an early death. I was even more surprised to learn that, within the first two or three years, the residents had

abandoned many of the principles in *Walden II*, and in behaviorism in general. They were more devoted to their vision of a free community than to Skinner's utopian ideals, as they understood them. Early on, they decided that the society described in *Walden Two* was unrealistic for them—perhaps for anyone—and that the principles they originally intended to follow stood in the way of their higher goals. So, like intelligent control systems, they began changing anything and everything which seemed to need changing. By the 1980s, the place had a decidedly humanistic quality.

By 1984, I had sent copies of what little was available on control theory to the Twin Oaks community library. In return, I received two books written by residents. I recommend them highly to anyone who is curious about the fate of the Walden II experiment. The books are: *A Walden Two Experiment: The First Five Years of Twin Oaks Community*, by Kathleen "Kit" Kinkade, New York: Quill, 1973, and *Living the Dream: A Documentary Study of the Twin Oaks Community*, by Ingrid Komar, Volume I, Communal Societies and Utopian Studies Book Series, Norwood, Pennsylvania: Norwood Editions, 1983. I assure you that the community described in those sources is anything but a coercive place operating under what they called "Skinner's scientist puppeteers"—the "planners" envisioned in *Walden II*!

Bill Powers: From the viewpoint of the "scientist" or "technologist," the manipulated disturbance is a controlled consequence of action. Action varies to make the disturbance be what the manipulator wants to see happening. Both action and disturbance are *dependent* variables. The disturbance depends on the action, and the action depends on both the current state of the disturbance so far produced, and external influences which interfere with producing the desired disturbance. From the viewpoint of the manipulee, the disturbance comes from outside the loop, arbitrarily altering a controlled variable; hence it is an independent variable. In both cases, a second *independent* variable exists: the manipulator's intention regarding the disturbance which is to be produced, and the manipulee's intention regarding the state of the controlled variable which is being disturbed. Higher levels are involved in both cases.

Gary Cziko: Culture appears to control aspects of human behavior (the system of law is a good example), but it can only do this through the interaction of human beings as autonomous control systems. Nobody outside of me can reach in and change my reference levels. Society cannot control my thought. But growing up in a particular society and culture present problems which might lead me to reorganize my control systems in new and (usually) culturally appropriate ways.

The idea that society can control individuals' thoughts and actions by nonviolent means has been proven wrong many times and remains a dangerous myth. Control theory provides the first real insight into the fallacy of this myth.

Rick Marken: Gary, your comments about societal control were right on target. Let me just say a few related things.

Gary knows that control systems really control their inputs (perceptions). The outputs of a control system depend not only on effects produced by the control system, but also on effects external to the system—these are disturbances. Disturbances can enter a control loop at any point; they can even be added to the neural signals in the control loop. These disturbances can influence every variable in a control loop; but the loop is organized so it always acts to keep the perceptual signal matching the reference. The disturbance might change the amount of output required to keep the perception at the reference; it might change the relationship between external variable and perceptual variable; but it cannot affect the relationship between perception and reference—the closed loop sees to that. So a control system doesn't really control movement or position or reflexes or whatever. It doesn't even control a variable in the outside world. The thermostat doesn't control "heat" in the room—it controls the voltage representing "heat" as represented by the metallic strip. If you change the heat transducer (metallic strip), you get a new voltage (perception) for the same heat, but the control system still keeps the voltage at the reference (which could mean that you experience a hotter or colder room).

The control system controls only one thing: a perceptual input signal. This signal can be a representation of simple or complex variables outside of the control system. When we look at the control system, we see that system influencing our perceptions—perceptions of movements and positions and "instincts" (really, program perceptions). But to know what the system is actually controlling, we must learn how our own perceptions are related to the perceptions being controlled by the control system.

Bill Powers: If A controls B, then for any disturbance acting on B, A alters its action so as to prevent B from changing significantly. A also, at the same time, determines the state in which B will be maintained (the state can be dynamic). I don't see how this applies to "society controls thought." If I change my thought, does society act on me to restore my thought to its original state? I don't see how society can even know the state of my thoughts, much less affect them in some way.

Joel Judd: The one overwhelming thought I've had as I've been coming around to a control-theory point of view is "how powerful this makes a human being!" That is, one is truly free to act/think as one wishes. We can be persecuted, persuaded, tricked, forced, etc., in order to make us act or think as another would have us, but, ultimately, we are responsible for ourselves! My religious beliefs have held this to be the truth for me for quite awhile, but it's refreshing to find a sort of secular basis to hang it on. The insidiousness of stimulus-response really serves to demean the human (or any) organism, doesn't it?

Bill Powers: There are all sorts of feedback relationships *between organisms*—anything imaginable, because there is no superordinate system regulating the interactions. Nor is there any control. There are only limits. And these interactions can be experienced by individual organisms which, being control systems, will modify their behavior to cope with them. This leads to an enormously interesting kind of study: the study of phenomena which emerge from interactions between true self-contained negative-feedback control systems. Because there is no superordinate system, no supervisor, the outcome is not governed by the same laws applying within the individual organism. Social laws are not simply a higher level of the laws of individual behavior. They are not analogous to the laws of individual behavior in any but the most superficial ways. A real understanding of how organisms interact is going to tax our capabilities for modeling for a very long time to come, and can't really get started until we have brought our models of individual organization to a much higher level of competence.

I think that negative feedback totally dominates all processes inside an *individual* organism, including the processes of growth and learning. I think that positive feedback could well be important in the realm of interorganism relationships. Those relationships, I think, are the major source of evolutionary pressure: the passive physical environment is, comparatively, a pushover. If it weren't for all those other organisms, there would be plenty to eat, plenty of shelter, plenty of safe places to mate and rear young. Of course, in that case we'd all still be at the bottom of the food chain, so perhaps I shouldn't complain.

Dennis Delprato: It might be appropriate for me, as a representative from the fair State of Michigan, to point out that control theory has recently won a "victory" of sorts. A district court judge here dismissed a first-degree murder charge against Dr. Jack Kevorkian, who was alleged to have supplied toxins, an apparatus, and instructions which enabled a 54-year-old Alzheimer's-disease patient to kill herself. The basis of the judge's decision to dismiss the charges apparently was simply that Michigan has no law applying to assisted suicide. The dark side of all this is that the forces of antiquity are stirring more than ever to have the legislature write a law to make assisting suicide a prosecutable offense. I'm sure that if they could get their hands on the nonspatiotemporal "soul," they would seek to prosecute this aspect of those who elect to kill themselves, as well.

I suggest that control theory (by whatever name) is extremely unappreciated in terms of its respect for individual liberty. As Bill Powers put it previously, control theory ethics (or a fundamental ethical dictum of control theory) is that others are control systems, too. In other words, keep your cotton-pickin' hands out of others' business (i.e., lives). Along these lines, I get very discouraged when I note enormous inconsistencies among individuals' positions on various issues. Most prominent in line with the present topic is when the same individual touts "civil libertarian" positions, yet holds to one of the many extant versions of one-way determinism. It certainly is difficult to refrain from strong-arm tactics in interpersonal and other social relations, given certain "positions" in which we find ourselves and the various ways in which such tactics are encouraged. It is all the easier to use authoritarian techniques when one assumes that others are simply subject to wind-weathervane operation, anyway.

You all better agree with me on this—or else!

Mark Olson: Dennis, I agree that an ethical system derived from control theory would basically have the form of "keep your hands out of others people's business" as you stated. But I do think that if the output of one person is going to result in a large (maybe irreconcilable) error, then laws (or something) need to be there to ensure that it doesn't happen. If they are not there, then everyone will reorganize and reorganize until they have made their worlds small enough to protect themselves from the outputs of others. Go to New York or anywhere people are screwing around with other people's inputs and try to communicate with those people.

Rick Marken: Dennis, I agree that individual liberty is fundamental to a control-theory ethics (if there is such a thing), but I also think that one of the most interesting (and calming) things which control theory does is to help us understand that the behavior of the "forces of antiquity" (an excellent appellation) is an expected consequence of the fact that they are control systems, trying to control variables which they feel are very important. (I guess any variable we are trying to control is important to us.) Control theory gives us some sympathy for the devil (and these forces of antiquity can be quite devilish if you are controlling for a variable which they think should be kept at another level or not controlled at all). You cannot control the controllers any more than they can, in the long run, control you. I don't know the solution to this

mess, but the ones who think they have a solution are usually the most obnoxious and dangerous controllers of all. I guess there really is no solution other than education, and, after seeing what some educated people think it is important to control, I'm not even really optimistic about that. For some reason, all this doesn't depress me—I think that, at a personal level, understanding the nature of control systems has helped me get along with other control systems; and, of course, there is great "spiritual" satisfaction to be had from learning a little bit about how living systems work. But I would be surprised if learning about control helps reduce the amount of unnecessary conflict in the world. "People just want to control" (sung to the tune of "Girls Just Want to Have Fun").

Bill Powers: There's a step beyond the "keep your cotton-pickin' hands off my control systems" prickle. We do have to live together, and therefore we have to respect each others' wills, as well as defend our own. This often means letting other people control our behavior.

Cooperation is harder to achieve than individual action, especially under the axiom that each person's will is entitled to respect.

Regarding the war in the Middle East: The only way to control a control system is through the application of overwhelming physical force.

It doesn't do much good to urge people to be wiser or cleverer than they are. We need to understand how system concepts come to be invented and accepted, and why they can become so compelling as to enlist the support of millions of people and throw them at each others' throats. Given that system concepts are perceived at a level higher than that of logic and rationality, how are we to discuss them, compare them, or teach them? From what standpoint can we even knowingly examine them?

I was thinking this morning that the thought of going against an enemy by using invincible force is comforting. War is exhilarating when you are sure you can win. You feel safe knowing that nobody else is stronger. If you are stronger than everyone else, there is no need to be wise or clever. You can even afford to be compassionate, within cost-benefit limits.

War is caused by fear, not bravery.

Rick Marken: I think that social organizations exist to help everyone involved control what they need to control a bit better than any individual could control by him/herself. This requires cooperation, which means everyone can't necessarily have things exactly the way they want—they must defer gratification or settle for a little less than they might get if they didn't have to cooperate (take the requirements of others into account). It would be nice if people could pay better attention to this aspect of social organization (mutual benefits through compromise and cooperation). I think one way to tell when people are *not* taking this into consideration is when they start talking about "legal rights" and "historical precedents" and other verbalizations which are used to justify screwing people up. In the Gulf War, one side talks about "legitimate claims and grievances" which Iraq had with respect to Kuwait. I say, who cares about such claims—they are just words and phrases. What I care about is that a very heavily armed group of people came in and very forcefully prevented another group of people from having any chance of being able to control the variables they needed to control. I think this violates the fundamental sense of justice which Hugh Gibbon talks about with respect to law; coercion is perceived as just when it is used to stop someone from intefering with the agency of another person (who is not, through their agency, interfering with anyone else). I think it is difficult to see what is happening over there as anything other than Iraq forcefully and brutally depriving Kuwaiti people of their agency. This was not done justly—to prevent Kuwaitis from brutally surpressing another group. There is no set of symbol manipulations which can make Iraq's actions seem just. So coercion was exerted by the US.—since coercion can only be exercised by an agency physically able to exert it.

I guess I'm saying that coercion is just when it prevents some person or other agency from depriving another person or agency of their ability to control. Verbalisms about "legitimate rights," "manifest destiny," "God-given rights," and "a legacy of imperialist domination" seem to me to be most often used as smokescreens to justify unjust coercion: depriving people of their ability to control for no reason other than unwillingness to take the time to look for cooperative solutions.

I will say that many of the U.S.'s verbal justifications for the war are also irrelevant—probably an old habit left over from our earlier commitment to using coercion to suppress rather than expand people's ability to control. I think the reason for overwhelming support for the war is that many people see this war, at a non-symbolic level, as a just use of coercion to prevent deprivation of agency—the same conditions under which we recognize the use of coercion as just in this society. The U.S. is acting as the policeman: as the policeman exerting legitimate coercion (police coercion can be seen as quite illegitimate if it doesn't seem to be used to protect agency).

Gary Cziko: Rick mentioned the need to give up a little when living in a society for the sake of cooperation with others. This reminds me of A.S. Neill's school called Summerhill, and his book by the same name, which I am now rereading.

Summerhill was remarkable for the total lack of authority in the

school. Classes were optional. Students could do anything they wished, as long as they did not infringe on the rights of others. Violations of others' rights were dealt with at a weekly meeting run by the students. Students were made to repay for goods or services stolen or damaged, but there was no real punishment, and absolutely no moralizing about good or bad.

A.S. Neill's acceptance and approval of each student was absolutely unconditional. He did not withold his love and support so he could "reinforce" desired behaviors. In fact, he often "reinforced" undesired behaviors, as when a student was rewarded for stealing. This reward was seen by the student as a sign that he was approved of as a person by Neill, no matter what he did, and made him feel to be part of the school community. And when this happened, the stealing stopped.

I wonder if anybody else out there is familiar with A.S. Neill's non-authoritarian method of child raising and would care to comment on its relationship to control theory. It seems to me that Summerhill is the closest any community has ever gotten to the type of community which control theory would lead us to have.

Rick Marken: On permissiveness and Summerhill: I don't think that control theory suggests that "Walden III" would be a place like Summerhill (as Gary described it; I've heard of A.S. Neill, but know little of the details of his community). I do think control theory makes it clear that people are autonomous control systems. But that means *all* people—students *and* educators. Problems arise when people start trying to control other control systems—but how can they keep from doing it? Control systems control. If another control system disturbs a variable you are controlling, you react, possibly affecting the other control system's ability to control. A.S. Neill might be perceived as more permissive than Skinner, but he is still a control system. If he really has a community where he just lets other control systems control, even if this influences the things he is controlling, then he is not alive any more. As long as there is more than one control system around, there will be some degree of mutual influence and, possibly, control. This does not mean that things will necessarily go to hell. All control theory does is draw our attention to the fact (and theory) of control and interacting control systems. The "solution" to whatever problems might arise because of this fact is not provided by the theory itself. I do agree, however, that efforts like Skinner's to control behavior will likely, but not necessarily, lead to enormous conflict. But then, complete "permissiveness" in a world of limited resources is likely to lead to the development of some pretty problematic control systems itself.

That's one of the problems of control theory—we don't sell well, because we can't honestly sell utopia. All we can sell is quality.

Hugh Petrie: Rick: Generally, it seems to me that your approach to the justification of coercion from a control-theory perspective is correct do, however, have one question with respect to your application of it to the U.S. actions in the Gulf. Remember "The Test"? If what the U.S. is doing is justifiable coercion, we should be able to ask about other disturbances to the world order, and whether the U.S. always acts to protect the agency of other societies. It appears not, to many of us, e.g., Lithuania. Thus, although justifiable coercion is a plausible candidate for what the U.S. is doing, it does not seem to survive "The Test." So what are we "really" doing? This is what worries some of us.

Gary Cziko: Rick, yes, we all control, but we can also control what we control. And control theory shows us what we can control (ourselves), and what we cannot (others).

This is the one thing I never understood about Skinner's behaviorism. If he was right that all behavior is completely determined by the environment (plus biology), then how can one have a technology of behavior? How does the behavioral technologist get outside the deterministic system to make things better? I can't imagine that Skinner didn't consider this problem somewhere, but I have not yet been able to find him writing about it.

I would hope that control theory avoids this problem by its hierarchical system of levels of control. What we think at a higher level *does* make a difference in how we behave. If control theory suggests that the only way to avoid violence is to respect the freedom of others, and if we want to avoid violence, then we might begin to respect the freedom of others. The thought is not something we induce using our senses, and we do not need to be "rewarded" by the environment for such a thought. It just has to make sense at a higher level—and it can change your life (and others' lives). In this sense, we are all "outside" of Skinner's deterministic system. And if this isn't a useful psychological theory, I don't know what is. Skinner's seems useless by definition. Control theory does seem to have the potential to make a difference. If not, I might have to pull the plug and start looking elsewhere.

Bill Powers: Rick, you old warmonger, you. I think we have a chicken-and-egg problem here, just like the one between the Israelis and the Palestinians. How far back do you want to keep score on who provoked whom to do what? The British screwed Iraq; Kuwait screwed Iraq; Iraq screwed Kuwait; we're in process of screwing Iraq; now Jordan and Morocco and Lybia want to screw us, etc., etc., etc. It's been a nonstop international tag match for as long as anyone can remember. There isn't any Gulf Crisis. It's just another episode. And everybody, of course, is completely justified. Just ask. We need some social and political scientists who can step outside this endless circle of words and show clearly how this mess is being caused by the people on all sides doing *exactly the same things* to each other. We have to go up a level, not take sides.

Rick Marken: Behaviorists like Skinner (and most other psychologists as well), claim that behavior is *controlled* by the environment; in Skinner's case, this control is exerted by reinforcers selecting the behaviors which produce them, but the mechanism is not important. If behavior is controlled by the environment, then the behaviorist can control behavior if he or she can control the environment of the behaving organism. *But* the behaviorist him/herself must also be controlled by the environment. So how can a person who is controlled be *in control?* Skinner has spoken to this problem. He talked about "reciprocal control": the animal controls the behaviorist as much as the behaviorist controls the animal. So the behaviorist gives a reinforcement as a result of seeing the animal do the desired behavior. The reinforcement makes the desired behavior more probable, making it more probable that the behaviorist will give a reinforcemnt. There are obvious problems with this analysis (it seems to predict that the animal and behaviorist will accelerate into a frenzy of behaving and reinforcing, which is not what we observe). But the real problem is that the behaviorist is not really in *control*. A small disturbance to the animal's behavior could lead to a very different end result produced by the behaviorist. The behaviorist cannot intend to have the animal "make a figure-eight"—this result cannot be expected on each occasion, because small changes in the animal's behavior lead to small changes in the behaviorist's behavior which might end up with the bird making a zero rather than an eight. Control implies purpose: making something happen even if circumstances are working against that end. This kind of purpose is what the behaviorist claims to have with respect to the behavior of others ("I can make you do what I want"), while denying such purposiveness to those very others. But both the behaviorist and the organisms studied by the behaviorist are supposed to work according to the same principles. I think this is the inherent paradox of behaviorism—if behaviorists can control, then they can't be controlled. But if they can't be controlled, then neither can the objects of their control, and so the behaviorist can't control if he/she can control. It's like the man who says he is from Crete where all men are liars. If the statement is true, it's false; if false, it's true.

I know that control theory does not suffer from this paradox. Control theory has no problem explaining the behavior of the control theorist with respect to the behavior of the objects of his/her theorizing.

Bill, I don't mean to be that much of a warmonger. I don't justify this war to myself (as I said) in terms of keeping score on who screwed whom the most in the Gulf. I would be happy to "go up a level" to find a solution. But how do you do this? What's up there? How do you "go up a level" when someone has just robbed you at gunpoint? Trashed you apartment in order to get you to move? I have been involved in conflicts where I have made every effort to be conciliatory and look at things from new perspectives, only to find that my "opponent" was perfectly happy *not* to "go up a level," but, rather, to take the simple expedient of threatening or using violence. I believe that people can and should try to get along, and I respect the fact that everybody is just trying to control what they feel they need to control. But what do you do with people who would rather kill than talk? I am not interested in past wrongs or justice. I just want to know how you deal with people (and there are such people) who consider no other option than force. I agree that it's best not to get into situations where that kind of confrontation emerges. But what's past is past. There was a rape of a country (regardless of the cause or justification). What do you do? Ignore it? Mind your own business? Non-violence and reason work if the people you're dealing with actually respect life and thought. What do you think Hitler would have done to non-violent Jewish protesters suggesting that it was inappropriate to gas German citizens? My guess: laugh his head off before shooting them all.

Don't get me wrong—I'm not pessimistic about the value of control theory as a basis for solving human problems. Part of my interest in the theory is motivated by its optimistic, humanistic perspective. I hope that the understanding we get from control theory will help us keep from getting into Gulf-type situations. But I have a feeling that there will always be people like Hitler and Saddam and many others of their ilk. I wish it were possible to wish them away, or love them away, or non-violence them away, or go up a level and make everything better—but I doubt it. Still, I'd love to see some concrete proposals for alternate approaches to the current and possible future world problems based on control theory.

Mary Powers: Summerhill! What a liberating read that was! Really the opening gun of the '60s—all you need is love, etc.

But raising children needs a whole lot more. Two thoughts:

First, Summerhill was an isolated place. All hell could break loose and did. But we have to raise kids in a cultural context, and we are their cultural context. They have to learn to get along in our society with its range of ways to be, just as they have to learn our language with its range of sounds. And some of those ways are not things children are going to want to do spontaneously. Fortunately, kids want to please, and are insatiably curious, and what they *must* do can be

offered in an enticing and interesting manner. But that takes a hell of a lot more work than just love—love's an attitude, not a curriculum.

Second, it's all very well to say that wonders occur when you give a child (or anyone else) unconditional positive regard. But who can do that all the time? It takes a very unusual person. I bet even Neill got fed up occasionally. Carl Rogers felt this was the key to therapy, but he only had to see each client for 50 minutes at a time. And I left the counseling center convinced that there was more to what was going on than UPR. The idea that you could simply reflect what the client was saying was almost a joke—it certainly lent itself to parody. But that's not the point here, which is that when you're a parent of little kids, you're with them morning, noon, and night, and sooner or later you are dealing with them on a gut level, drawing on a lot of unconscious stuff you learned from the people who raised you. If that was kind of screwy, you're going to be screwy too, however much you swear things will be different. Either you'll do the same things again, or you'll compensate and do the exact opposite. And in "doing" I include a lot of talking and acting which is probably very conditional indeed, which you might not even recognize as such, because it is so automatic.

Rick Marken: The simplest case of conflict occurs when two control systems have different reference specifications for the same controlled variable. The control systems can be in the same physical system (like a person) or in two separate physical systems (like two different people): in the first case, we have intrapersonal conflict (the person is in conflict with himself or herself); in the second, we have interpersonal conflict (two people in conflict with each other). War is an example of interpersonal conflict involving many persons.

It is pretty easy to model a conflict between two control systems. For example, the outputs of the two systems could be as follows:

$$O1(t) = K1(R1 - P)\ S1(O(t - 1))$$

$$O2(t) = K2(R2 - P)\ S2(O(t - 1))$$

Time is t; the perceptual input P is also changing over time. Assume that for both systems:

$$P\ O1(t) + O2(t) + D$$

Thus, each system influences the perceptual input to the other (since it is the same input). Each system responds to the discrepancy between this perceptual input and its reference for the input (R1, R2). If R1 = R2, there is no conflict—both systems want the same perception. If R1 differs from R2, then there is conflict. The outputs generated by one system will be a disturbance to the input to the other. The result of this conflict depends on the relative strengths of the two systems, as represented by K1 and K2 (output generated per unit error). If K1 = K2, then there is a stand-off. The systems match outputs until they are each producing the maximum they can physically produce. If one system can produce more output than the other, then that system will dominate, but not necessarily "win" the conflict unless the residual output can completely compensate for the output produced by the other system. If K1 and K2 differ, one system might dominate the other, but, again, winning depends on the maximum output which can be generated by each system. The intensity of the conflict between control systems depends on the relative values of K1 and K2, the maximum values of 01 and 02, and the difference between R1 and R2.

There are several ways to solve a conflict like this, where "solve" means that all systems get their perceptions to match their reference states. The simplest approach is to simply let the conflict go and hope that the output limits of your system are greater than those of the other system—much greater. Then one system (the stronger one) can get the perception it wants, and the other system gets massive error. This is the solution called war. There are obvious problems with this solution: 1) you can't be sure that you are the system which is going to "win," and 2) unless you completely eliminate the other system, it will never stop trying to get its perception to match its reference, so you will always be generating some output to prevent this (rather than devoting this energy to controlling other variables). There are other problems, but that's enough for now.

The solution to conflict which is "best" requires that one or the other party to the conflict "change their reference" for the mutually controlled variable. That is what "going up a level" is about, if you believe that the reference signals are set by higher-order control systems. The higher-order system could then see the lower-level perception as part of a higher-order controlled variable (like "being a big hero"). If this system could find, say, other lower-order perceptions which satisfy this perception, then maybe R1 could be eliminated (so one system no longer has a reference for this perception), and a new perception could be substituted. The problem is that, when the conflicted control systems are in different physical systems, it is hard to get *both* systems to solve the conflict by changing references. If one system is always willing to change its reference in order to prevent conflict, then there is the possibility that the other systems will notice this and rely on it. It could get to the point where the accommodating system becomes a doormat (which is certainly OK if it really never has any interest in controlling any variables at levels which might cause conflict with other systems).

There is no "morality" in this view of conflict. Conflict just happens because control systems control, and there are limited degrees of freedom (apparently) available in the perceptual world of variables with which all these systems interact. Conflict seems to me to be unavoidable. But people who understand the nature of these conflicts will probably be better at dealing with them than those who are paid to do it (the politicians). Still, if one system in the conflict just refuses to change a reference for a perception which, if kept at that reference, will cause intrinsic error to another system, what can you do?

It might be interesting to try to model interacting control systems which can get out of these conflicts—and do so in ways which do not destroy the physical integrity of either system. Instead of praying for peace, we could be *modeling for peace*.

Bill Powers: Therapists, like everyone else, want to be in control. Some of them want to have the power to cure people, like a doctor. Many of them dream of saying *just the right thing* so the patient's jaw drops, the patient's eyes bug out, and the patient cries, "Oh, thank you, Doctor, that's exactly what's wrong, I understand everything now! You're so smart!" If a therapist doesn't provide insight, diagnose problems, give people good advice, administer treatments, and cure the patient, what's the point in being a therapist? That's what it's like to want to be in control of the client.

Did anyone ask the patients what they want?

From the control-theory viewpoint, the goal is for the client to be in control, isn't it?

Rick Marken: This is for the social scientists. I would like information about the following: What is the history of the concept of control of human behavior? (Particularly the idea that people can be controlled by non-coercive means.) People have known that animal behavior can be controlled fairly non-coercively for some time. Rulers have known how to control people coercively for some time; they have understood the effectiveness of one contingency—if you do this, you'll die. Machiavelli apparently was an early writer on controlling people, the how and why, but I never read him. Is he a good one to include in such a history? Is the idea of non-coercive control really as modem as I think it is—beginning about 1913, with J.B. Watson? Didn't people always believe that kids could and should be controlled?

Also, I wonder why psychologists don't talk much about the control of behavior any more. Any ideas? After all, if cognitive or connectionist or whatever models are right (successful), then they should make it possible for people to control what is being modeled. Why is there no more concern about behavior control and brainwashing? Is it because it hasn't worked? And if it hasn't, why haven't people abandoned the causal framework which suggests that such control is possible?

Tom Bourbon: Rick, I think I understand your request for citations about non-coercive control of people. But I am not sure I understand completely, due to one example you cited, namely J.B. Watson, the original American version of a pure environmental determinist. Watson? Non-coercive? I'm not really sure what you mean by the word, if Watson is an example.

A possible source for information on control, both coercive and allegedly non-coercive, is Harvey Wheeler, editor, *Beyond the Punitive Society: Operant Conditioning, Social and Political Aspects*, San Francisco: Freeman, 1973. The book came out in the wake of Skinner's *Beyond Freedom and Dignity*, and it includes arguments pro and con on whether operant conditioning and its then-fashionable applied wing, "behavior modification," represented the leap beyond punishment which Skinner claimed. Of course, those who agreed with Skinner conveniently overlooked the fact that their "positive reinforcers" worked only if the recipients of this "non-punitive" therapy were first denied something they previously had, and were not allowed access to the denied substance, item, or action *unless* they did what the "non-coercive" therapist required.

In China, where Taoism certainly embraced a non-coercive model of nature and of society, Confucianism, the philosophy of the "practical and applied" side of society, was almost a polar opposite of Taoism. That was the idea—a balance, within society as a whole, between the restrictive, coercive practices needed to keep the society running, and the free, childlike Way of Tao to which people were encouraged to return—after they had fulfilled their obligations to state, family, and all of the rest. Precisely that same balance between coercion and freedom existed in traditional Hindu culture, where the free and enlightened path of Buddhism came into being as a counterpart to the mandatory rigors of organized society, and people were encouraged to recapture some of the freedom and spontaneity of youth, *after* meeting their social obligations.

The modern West does not deserve credit for discovering coercion.

Chuck Tucker: In my lectures in Introductory Sociology, I tell the students that there are only three ideas which have developed in the history of Western civilization regarding the concern human beings have had throughout *recorded* history about control. I claim (correctly or not) that since the time of the Greeks (our beginning of records for Western civilization), there has been concern about the "forces" which make us do what we do, individually and collectively. The ideas are

Nature, God, and Society (or Man). The introduction of the last idea (Society) did not occur until about the 16th century. The idea of Society as a force is in opposition to the other two ideas, but all of these ideas (and some from the non-Western world) are used by people today to answer the question: Why do I (or we) do what I (or we) do? (Although I don't use it directly, the book *The Day the Universe Changed*, by James Burke, makes this point much better than I do.) *But* throughout the history of Western civilization, the idea of control has been *coercive*. And the control-theory idea of control is non-coercive. *That is our problem*: we are presenting a view which, although consistent with the idea of Society (as compared with Nature), calls for a departure from "outside forces."

I have just begun to read Jack Gibbs' book *Control: Sociology's Central Notion* (1989), which shows how the idea of control has been used only indirectly in the social sciences. Gibbs claims that control is a central idea and that if it were used explicitly, it would improve our understanding of social life. (He has no references to cybernetic control theory in his book.)

Kent McClelland: As a sociologist, I've been interested in control theory ever since I read *Behavior: The Control of Perception* several years ago, but I didn't get time to dig into it deeply until a sabbatical finally came along last spring. When my sabbatical ran out last summer, I had an unfinished, rather sprawling manuscript on connections between control theory and sociology.

The draft contains an introduction focusing on Jack Gibbs' book, *Control: Sociology's Central Notion*, saying that Gibbs missed the boat by ignoring Powers' control theory; a very brief but fairly comprehensive review of the multi-disciplinary literature on control theory (now a year out of date); a section attempting to explain the basics of control theory to an audience of sociologists; and a final section applying control theory to a discussion of interpersonal power.

Bill Powers: People often use terms like "live in harmony" and 'learn to cooperate," which sound like unbiased and fair prescriptions for social life. But they seldom mean it that way—just think of a teacher Who puts on a report card "needs to learn cooperation with others." What it means is that Little Johnny had better get in line or there will be trouble for Little Johnny. Cooperation definitely doesn't mean that the social group is going to cooperate with the individual; it's a one-way street the other way.

The person with power (personal, financial, or political) uses it to assure that those who want the gold (food, shelter, health) behave as that person wants, whether or not it conflicts with what they want.

This method creates a certain limited range of harmony within a group of limited size, but at the expense of harmony within all but one of the individuals: outward harmony, inward conflict. It does not work across groups.

The opposite is no cure: inward harmony, outward conflict. Think of Donald Trump enjoying his triumph over that little businessman on whom he stomped. He bragged about it in a book! No inner conflict there. But who would want to live in a world of Donald Trumps? How long would such a world last?

True harmony means inward and outward harmony. It means that in finding ways to avoid conflict with others, individually or as a society, the individual is also able to avoid inner conflict. Our world is not set up this way at present. It's set up on the basis of controlling others and winning conflicts. It's set up so that most people must be losers, because winning is organized like a pyramid. Kids are taught that in this country *anybody* can grow up to be President. Somehow, all these millions of kids fail to be advised that in their lifetimes there will only be perhaps ten Presidents. Some opportunity: there's a better chance of getting into the NBA. The ladder of success is not designed to let everyone climb it without knocking someone else off.

I don't think that the world is going to be either saved or destroyed by any particular set of proposals as to how we should run our affairs. Specific proposals are at too low a level. So are specific principles—moral standards, economic principles. The ideas which stick around and have a long-term (if slow) effect are the system concepts (or whatever that level of conceptualization is). The question is always "What kind of world do I want to make and live in?" That question is even more important than "What kind of person do I want to be?" Living in a world of limited resources with other people who are just as autonomous as you are is a difficult problem, an extremely complex problem. We will arrive at successive approximations to solutions by trying different solutions and seeing how they work. Gurus and saviors come and go; they leave their traces, and we choose which traces to retain. Blind variation, but selective retention.

My point is that when we think at the system-concept level, we are far more likely to be helping to provide a choice of viable futures than when we simply propose clever sets of principles and rules which look as if they might achieve some immediate semblance of order—even a New World Order.

Ed Ford: Bill said, "... when we think at the system-concept level, we are far more likely to be helping to provide a choice of viable futures than when we simply propose clever sets of principles and rules which look as if they might achieve some immediate semblance of order...."

I just don't think you can separate the two levels. They have to be in sync with each other. You can think at the system-concept level, but ultimately that thinking has to be translated into some kind of trial and error process which tests the validity of the system-concept level. That means you have to set standards, then make choices based on those standards. I think the harmony within us—the real, continuous, long-term, peaceful harmony—has to exist between the levels and within the individual levels. This harmony can exist to some extent even in trying times in the external world (Viktor Frankl's *Man's Search for Meaning* is an example). I agree that dealing only at one level doesn't offer a "viable future." The key is to maintain harmony throughout *all* levels as the system continually interacts with the environment within which it finds itself in order to satisfy the demands it makes on itself and the demands made upon it. As I work with clients (who often are locked into marriages, children, and/or jobs), I am trying to help them establish some peaceful order within their systems which will help them to find as much peace as possible (if this is what they want) in a very trying and stressful environment or set of circumstances. Is there a set of system concepts (and subsequent and corresponding lower levels such as standards and choices) more efficient at achieving these goals than others? For me, I think so. That is my search. For others, my job is to help them search for what might help them. I have known too many people at peace in very conflicting situations (my wife's handling of eight children and her husband when the youngest was still a baby and the older ones where in their teens).

I see problems arising when people set very different standards for the "same" system concept. The recent differences within the Presbyterian and Episcopal Churches are examples. Thus, the need to follow up on an established set of system concepts with standards which will make consequent choices reflecting what is wanted. The ultimate test of a set of system concepts within a living control system is its ability to deal with the present and future environmental situation in which it finds itself and the subsequent sense of satisfaction (peace, harmony, whatever) which follows within that system. Ideas just have to be tested in the marketplace to determine their validity, that's all. And to do that, standards will give specific direction for the choices we make. The ideal might be to have both internal and external harmony, unfortunately we don't live in that kind of world.

Rick Marken: Ed says, "Is there a set of system concepts (and subsequent and corresponding lower levels such as standards and choices) more efficient at achieving these goals than others? For me, I think so. That is my search."

The same set in all situations? For all people? If one takes the control model seriously (as an approach to understanding human nature) then system concepts are perceptions set to particular values to maintain other variables at particular values. The model implies that even at the highest level of the control hierarchy, there is no absolute "right" set of references (absolute across people and environmental situations) which can achieve control.

The only possible variables which might qualify as "absolute" in the control model are the intrinsic variables—things like carbon dioxide and oxygen concentrations in blood and tissue, etc.—that must be at particular values or the physical system itself stops being able to function, and there is death. Looking for a best set of system concepts, principles, or whatever has been, in my opinion, the main cause of problems among humans. After all, if there really were a best set of system concepts, then the only right thing to do would be to teach them to others. But there is always the annoying possibility that other people won't buy into these concepts the way they should. This leads to ostracism, prejudice, and, of course, genocide. I think it's better to look for the right model of systems—and forget about the right system concepts which systems should have.

Ed also said, "The ideal might be to have both internal and external harmony, unfortunately we don't live in that kind of world."

We certainly don't, and we never will if the only test of a set of system concepts is the extent to which they give the system the ability to deal with present and future situations (i.e., internal harmony). As Bill pointed out, there have been people with lots of internal harmony (as far as anyone could tell) who created enormous external conflict. Slavery made it in the marketplace for years. System concepts, values, standards, and whatever have been changing over the years as the demands of the marketplace have changed—human sacrifice used to be a very big item in the marketplace of values.

I think people are frightened to realize that system concepts, values, and standards are not absolute—never were, never will be—because they feel it means that things will quickly get out of control with no absolute, correct standards. The control model shows that this is precisely the opposite of the truth. Changes in these variables indicate that control is going on—and that the principles, standards, and values are simply part of the means of controlling some other variable—something we can name and experience, but not very easily describe—what we have been calling system concepts. But even system concepts can vary to control something even more basic. I argue that if these standards and values were absolute, then things would definitely be *out of control*. The "things" I mean are the things which are most basic (and elusive) about human nature. Again, I note that trying to keep your standards, values, principles, or whatever at one absolute level puts

you as out of control of whatever is defined by those variables as if you decided to keep your hand in only one fixed position while you are playing tennis. Variability of means is as important a part of control as is consistency of the ends.

Absolute (or fixed) references at any level of the hierarchy mean the end of control and the beginning of intra- and/or interpersonal conflict. Maybe.

Joel Judd: Rick, replying to Ed, says, "If one takes the control model seriously (as an approach to understanding human nature) then system concepts are perceptions set to particular values to maintain other variables at particular values. The model implies that even at the highest level of the control hierarchy, there is no absolute "right" set of references... I think it's better to look for the right model of systems—and forget about the right system concepts which systems should have."

I keep wondering if I'm thinking about system concepts in the same way others are. "Mechanistically," I can see how there wouldn't be a specific, unchanging value for every level in each control system. But by the time you reach higher levels, the very reference itself, while we give it a name, is "variable," isn't it? I mean I could argue that a certain definition of "family' (e.g., mother and father and children) is the best. But of course, every single instance of family would not be exactly the same. In one, both parents might work; in another, only the father. One might have three children; another six. But a "family" of mother, father, and children could be the "best" social organization for having and raising kids, continuing the species, whatever. A single parent is not. Orphanages are not. Living with your aunt is not. That doesn't mean those things don't happen. At intrinsic levels, you can say that certain oxygen levels are best (even necessary). At higher levels, why can't you say similar things? The difference is in the variability (degrees of freedom?) allowed by something like "integrity." The things I do and say are going to be different than the things you do and say, but wouldn't you rather deal with someone who has integrity than someone who is untrustworthy? Is this idea of greater latitude as one goes up levels accurate? Is there a better terminology for it?

Just because *humans* can't always do things right doesn't mean there isn't a right way to do things (with the individual variability alluded to above).

Rick again: "I think people are frightened to realize that system concepts, values, and standards are not absolute—never were, never will be—because they feel it means that things will quickly get out of control with no absolute, correct standards."

It's the meaning of that word *absolute* which I'm asking about. Rick's comments remind me of some comments made by a visiting behaviorist. The seminar was on education, and he was asked about his views regarding the model he used. Some of these questions led into aspects of curriculum decision-making. Whenever this happened, he deflected the questions by saying something like, "That's a political question. I leave those decisions up to parents, school administrators, politicians. If you want to ask *me*, as a person, I think I can give an answer. But my [behaviorist] model simply describes/explains learning, decision-making, whatever. It doesn't imply *what* learning, decisions, etc., would be 'good' or 'best!"

That's the kind of message I get sometimes from this discussion. I can understand it. But I have to wonder at what point (and others might say "if ever") we allow those "political" questions back in. Much of the attraction for me of control theory is the implication that there is reason to argue for better ways of doing things. Bill Powers mentioned once that this starts to leave the realm of modeling, inasmuch as *what* a system concept *is* isn't necessary to an understanding of how that level and others might work. But people in counseling, law, education, etc. have to work with real systems every day. That teacher is deciding the right thing to learn, as well as the right way to do it. What kind of help do we provide them with?

Bill Powers: Joel, the "variable" aspect of a system-concept perception, as I've imagined it, is simply the degree to which the perceived situation (principles, programs, etc.) qualifies as an example of the given system concept. In other words, I assume the "pandemonium" model, in which there are various system-concept recognizers all working in parallel, and they all receive lower-order information. They all respond to some degree by producing perceptual signals, but some hardly respond at all, while others respond maximally. The alternative would be to say that there is *one* system-concept recognizer, which responds to inputs by producing a perceptual signal somehow encoded to indicate the presence of one system concept or another one—but only one at a time.

I think the first model, although probably too simplistic, is closer to the way real perception works. I can say that a strange animal is a little like an elephant and a little like a snake, but perhaps more like an aardvark. What I can't do is say it is like some animal I have never seen. I don't perceive a single thing which is somewhere on the scale going aardvard... snake... elephant. I have to figure out what it is by looking at simultaneous responses from a number of recognizers, each set to recognize something I've experienced often enough before to recognize again. Same for system concepts: I can say that a particular parent-society-child system is a little like a nuclear family, more like an extended family, and very little like a state-controlled family (take

them away at age 5 and raise them in an institution). I don't see a given family arrangement as a single point on a scale of different kinds of family arrangements.

Just what makes these different perceptions different can't be seen at the system-concept level. You have to look at the different principles employed, the different strategies of rearing, and so on down the levels. The hierarchical model does with levels what the multipurpose single-signal model does with a lot of internal complexity and memory in a single system—less informatively, I think.

So the *kind* of system concept is fixed in any one control system at the system-concept level. What is variable is the degree to which a given environment exemplifies that system concept (this agrees, I think, with your proposal). This means that we judge the environment at this level in terms of several, even many, different system concepts, all at the same time, in parallel.

The reference signal *can* be a constant. If you want to see a "nuclear family," you choose the degree to which this perception is to be sensed. Do you want a "pure" nuclear family which excludes teachers, friends, honorary uncles, and so on? Or is some degree of nuclearity less than the maximum more preferable?

In general, different system concepts can be derived from overlapping subsets of principles. For example, in the nuclear, extended, and state-controlled families, one principle in common might be that of keeping the child safe; another might be that of educating the child; another might be that of giving the child a sense of success and approval. Other principles might not be shared: giving the child a strong sense of self; providing experiences of equal love and trust with many adults; teaching the child to subordinate self to society. Various principles are chosen to be consistent with each other under a particular system concept; different system concepts are built from different subsets of the principles one knows how to perceive.

Contrary to what Ed Ford said a few posts ago, I don't believe that we choose a system concept *in order to* promote principles. That would make principles into a higher level than system concepts. I think we select principles so as to fit a given system concept. Of course we entertain more than one system concept, and the ones we choose to defend can easily require selecting contradictory principles. Christian businessmen have problems like this all the time, whether they ever reflect on the contradictions or not. I go along with Ed to the extent of saying that we have to revise our system concepts to eliminate such contradictions, but we do so to eliminate conflict, not to preserve any particular principles.

Joel, you say, "... wouldn't you rather deal with someone who has integrity than one who is untrustworthy?" I don't think that words like "integrity" and "trustworthy" can serve as system concepts. They have to do with principles which are necessary to make system concepts (particular social ones) work, but they say nothing about the system concept itself under which they are applied. Hitler wanted trustworthy aides to be in charge of getting rid of the Jews. The interrogators of the Spanish Inquisition might well have shown integrity in not pretending to have obtained a confession which was not actually obtained before the subject died. No matter whom you ask about system concepts, you will find those concepts defended in terms of uniformly noble principles, principles which most people would agree with. For a long time, the United States government hesitated to extend the right to vote to black people, for fear of violating states' rights and overextending the reach of the central government. Opponents of gun control do not argue that they should have the right to shoot anyone they please; they talk about the Constitution, a man's right to defend his home, the need to retain the ability to resist dictatorships, the right of self-defense. They cite all the principles that people with other system concepts are likely to share, thus making it difficult for others to say that the NRA is wrong about something.

People are pretty fuzzy about system concepts; they get them mixed up with principles and often get the order reversed, as if the principles were more important than the system which makes sense of them and selects them. When people come right out and describe their real system concepts ("This is a white Christian nation"), they tend to leave their opponents discombobulated—it's hard to say what you don't like about that (if you're white and Christian), other than that you just don't like it.

System concepts aren't justified by principles; they determine what principles you will employ. I think we sense that when we come across a bigot. The bigot's problem, from our point of view, is in the basic premise. You can't argue anyone out of a basic premise because it isn't controlled by something at a higher level (as far as I know). A system concept is part of a world view, and world views are very hard to budge. They determine what looks like Truth and Right to you. So everyone, even the KKK, thinks that Truth and Right are their property. We say they're doing Bad things; they say they're doing Good things. They even quote from the same *Bible*.

Principles of justice, honesty, and faithfulness are not sufficient to define a "good" system concept. They can be subordinated to concepts we might approve of, and to others we would abhor. And what "we" means depends on which patch of earth you happen to occupy, an accident of birth.

So what is the right system concept? I agree with Rick. There isn't one which can be proven objectively right. If human beings don't

know the right one, then nobody does. The rest of the universe is not designed to "know" anything. And I don't think that anyone right now is in a position to say which one or ones are empirically right. The whole picture is just too muddled; as I say, few people even discern a difference between principles and system concepts. Before any concerted effort to revise and improve our system concepts can be made, people have to acquire at least some notion of when they are talking about system concepts and when they are talking about the means of implementing them.

Religion has preserved an interest in questions like these which science abandoned long ago. So I'm glad that religion is still around. I can even see merit in some of the system concepts implicit in various religious beliefs. Love thy neighbor is a pretty good principle, especially if the neighbor is me. I'm even willing to take it on as my own principle, within reason, because it seems to fit with a workable system concept of a society of human beings. But I don't think it's going to do anyone much good if it's taken as a command from God. If you take it that way, you will never try to work out *why* it's a good idea to love your neighbor. So you'll never grasp the system concept within which this principle makes sense. You might even conclude that in order to love your neighbor, you had better stay in the right neighborhood and not let inferior unlovable people move next door.

Rick Marken: Joel, Bill Powers said much of what I would have said in reply to you—only better. So I will just make some general observations.

My reluctance to recommend specific reference levels for system concepts, principles, programs, etc. as being the ones which people should control for is not based on "political considerations." It is due to my current understanding of human nature, which leads me to believe that they do not exist. The "right" reference level for any controlled variable depends on 1) the context of disturbances in which higher-order variables are being controlled (and in which those variables are among the means used to control other variables), and 2) the context of other variables being controlled by the system.

Claiming that some principles are better than others is as meaningless as saying that some postures are better than others. (By the way, this can all be made more tangible by watching the behavior of my spreadsheet hierarchy. It really helps you get a picture of how a multilevel hierarchy of control systems, with many systems at each level, works. The behavior of the model is really quite amazing.)

It is possible, in principle, to say things about the result of controlling a variable at a particular level in a particular context. For example, I could say, "If you take a step forward when you are standing on a cliff, you will fall." Does this mean it is now possible to say, "Never take a step when standing next to a cliff"? Of course not, because the person might *want* to fall off the cliff—like the divers in Acapulco. Saying you know the "right" references for system concepts, principles, etc. presumes that you know everything about a person's entire hierarchy of goals and, more importantly, the current and future state of the world in which they live. I don't think anyone imagines that such knowledge will ever be possible, even in principle. So, the hierarchical control model implies that it is only the system itself, not anyone outside it, which can determine the right setting for all of the perceptual inputs it is controlling.

I suggest that this implication of the control model is one reason people will always find it hard to accept (just as the implications of the evolutionary model make it hard to accept). People (well, most of them) seem to want nice rules to live by. And they have them—in the reference signals to the program level from the principle level of their own hierarchy (I see rules like "thou shalt not kill" as programs, the particular instances of which are selected to instantiate principles, like "life is to be valued"). But people tend to imagine that these references for principles come from "out there"—and they *do*, with respect to the lower levels of one's own hierarchy. Moreover, people tend to think of them as "right"—because they *are* right for that person. But somehow people go on to assume that these references for principles *must be* right for others too. Part of this results from the fact that most people understand that they must cooperate with one another to some extent in order to succeed individually. So there is always the fear that if everyone sets their own references for system concepts, principles, etc. there would be chaos—everyone would run around killing each other and stealing stuff. There is no question that people must agree on some high-level variables which "must be" kept at certain levels, or cooperation will fail. But that's the problem control theorists are talking about, and there is no magic solution to, the problem, no set of clever rules from on high which will result in everyone getting along. People have done pretty well at cooperating for quite some time. The control theorist has "faith" that an understanding of what kinds of variables people control and why they control these variables could lead to approaches to personal and interpersonal interaction which will produce better results from everyone. But I am sure that solutions can only be defined from the point of view of the participants themselves, who are living in an ever-changing environment.

So I am sure that improvements in personal and interpersonal control will not result from the discovery of the "right way to behave." I'm afraid it's a bit more complex than that—whether we like it or not.

Ed Ford: Rick says: "Looking for the best set of system concepts, principles, or whatever has been, in my opinion, the main cause of problems among humans. After all, if there really were a best set of system concepts, then the only right thing to do would be to teach them to others. But there is always the annoying possibility that other people won't buy into these concepts the way they should. This leads to ostracism, prejudice, and, of course, genocide.... System concepts, values, standards, and whatever have been changing over the years as the demands of the marketplace have changed...."

I don't believe people are any different today than they were 200 or 2000 years ago. I don't think the demands of the marketplace change people, I think people themselves create their own demands and are responsible for them.

But to the heart of the problem: I spend most of my time counseling others and working as a consultant in various social service facilities, especially treatment centers and schools. I've been married 41 years, I've raised eight children, and I work with couples and families who are trying to establish or restore harmony in their lives. My total experience leads me to believe that there are certain values and standards from which people make choices, and upon which people base their lives, which provide them with a great deal of peace within their family and within the community in which they live.

I use control theory ideas daily with my clients to help them reflect on their created system concepts, standards, and choices. I don't ask people to buy into my concepts. Frankly, most people really don't care what I believe, but whether I can teach them how to rebuild their own lives.

I teach them to reflect on what their present values are and how they have prioritized them, on whether their standards reflect their values, and on the current choices they are making. I deal solely with their internal living control systems. My system concepts are not dealt with. Their values represent their present blueprint for how they believe their lives should be lived. Rick also says, "So, the hierarchical control model implies that it is only the system itself, not anyone outside it, which can determine the right setting for all of the perceptual inputs it is controlling." I couldn't agree more. Anyone who tells people they are wrong, tries to convince them to do such-and-such, tries to make them follow certain external rules, etc. is doing irreparable harm.

However, if their lives are not going well, and there is conflict within their systems, then my job is to teach them how to review and then evaluate their system in light of their own hierarchy. I don't believe it is possible to force my ideas on anyone (anyone with children should know that). For example, a man might be having an affair (program level) and have a belief system which says it is wrong (principles level).

He has put himself into conflict. Or, I had a man who was trying to work at his marriage, and his wife's priorities were work, alone time, children, husband, in that order. Guess what happened to that marriage? Rick, theorizing is one thing, but taking control theory into the marketplace and trying to apply it there is quite another thing. And what does that involve? I think it involves teaching clients how to deal more efficiently with their systems as they presently have created them so the conflict from which they are suffering can be reduced.

Rick, you quoted me as saying, "Is there a set of system concepts... more efficient at achieving these goals than others? For me, I think so. That is my search." I am not talking about my specific religious convictions, nor am I trying to force anyone to conform to my specific religious beliefs. I am talking about the system concepts, the values and beliefs, the priorities, the standards, and the choices of the hundreds of people I see yearly and whose lives are a mess. They are looking for help. I believe from my experience of working with families and individuals over the past 25 years that there are certain principles which work much better than others. I don't force my specific values on others. My experience with others shows me which values seem to work at restoring harmony, and which don't. I watch people struggle, and I teach them how to rebuild their lives. From this experience, I can only say this: you bet your sweet life there are values which really work well—such values as respect for one's spouse, seeing value in one's children, having respect for the integrity and worth of another human being (read: living control system). What I do is to teach clients to evaluate whether the implementation of their concepts and principles is getting them what they want (peace, happiness, whatever).

Last night a woman called me asking for help on dealing with her husband whom she had just learned was having an affair. Ultimately, her husband is going to have to come to terms with his system concepts, his standards, his choices, and all those things with which all of us have to deal. That's what I am talking about when I say there are certain values which seem to be universal, which work well for most people. I'm not on a crusade to get Rick to conform to my standards, I'm just trying to figure out how to help those in need more efficiently by using control theory, and, in the process, look for universally accepted standards.

Bill Powers says, "Contrary to what Ed Ford said a few posts ago, I don't believe that we choose a system concept *in order* to promote principles." If I said that, I was certainly wrong. I've always felt that principles should reflect the higher order. But, when we are building an understanding of a system concept, don't we move from a lower to a higher order?

Bruce Nevin: Part of my checkered career has included a two-year training program in family therapy. My perspective there has been that family systems appear to be living organisms. Probably it is best to phrase this in terms of the unconscious participation of individuals in family processes which continually recreate and sustain the family system. The actions which are matters of conscious individual choice pertain to the individuality of family members; the actions and inactions making up the fabric of "being a member" and the fabric of that of which one is a member involve distinctions which don't make any conscious difference to the ordinary individual (dialect, body language, posture, voice qualities). I have slipped here from family to broader constructs of social class, community, ethnos, and culture, so let me explicitly say that I believe we as individuals participate in the same sorts of processes continually to reconstruct and sustain our social reality in all its aspects.

The point is that this participation is out of conscious awareness, except for individuals specially trained or adapted to control some aspects of these processes consciously, such as salesmen, politicians, and actors. As Gregory Bateson pointed out, this is why we distrust salesmen, actors, and such. He was referring specifically to how the body language expressing a given interpersonal relationship cannot be subject to conscious control without thereby losing its ability to convey that relationship sincerely.

This sets up a dilemma for study of higher levels of control. Ask a fish about water. You can experimentally ask the neuron or the muscle what it is controlling, because the neuron or the muscle is not itself framing the experiment. But in asking about perceptions controlled by yourself, the experimenter, or by your peers, fellow humans, you require awareness of differences-which-make-a-difference of which you must not be aware if you are to continue to control them appropriately.

In family therapy, the perspective is to prescribe actions to the individual family members which don't make sense to them as individuals, or which might seem paradoxical, because the only way you can address the family system as patient is through the individual members of that system. As a family, they understand and learn, even while to them as individual persons the prescriptions continue not to make sense, except that their relations and communications with one another improve. Something of this addressing of human systems through their constituent members is, I think, required for experimental work with higher levels.

The following survives from a dialog in another context about the virtue of competitiveness. My interlocutor challengingly asked how there could be any success in the world without competition and without the dynamic of victor and vanquished. I offer it as a contribution to the discussion of the evaluation of values and system concepts. In my view, the root of ethics is this: that which tends to unity is preferable to that which disintegrates—a dictum to be interpreted in terms of systems and levels if it is to be sensible.

Some of this is preaching to the choir in this forum, but the information about anthropology and Ruth Benedict's work is, I think, news here.

The Hopi don't fight, they have lost surprisingly little of their resources given the vicious history of their territory, and the members of their communities are not suffering in any obvious ways because of their pacifism. They don't kill because it is obvious that it is not an appropriate thing to do, as you would not eat feces.

Judgment is a tricky matter. There can be no judgment without a point of view. There can be no "objectivity." But taking a point of view imposes a perspective in light of which some things appear more favorable, others less. The usual way of approximating objectivity is a process of consensus. You have your perspective, I have mine. If by reciprocal explorations we discover some commonalities, they are more likely to be "true" than points on which we disagree. To accomplish this, one's own perspective must become an object open to inspection and potential falsification, on an equal footing with the perspectives taken by other participants. Something of this is what is called scientific method, honored, alas, more by verbalism than by practice.

Evaluation of cultural differences is especially tricky. It is extremely difficult to bring one's own cultural perspective to a conscious level, where it is open to inspection and potential "falsification" on an equal footing with other cultural perspectives. It requires enormous effort, and that effort, in my experience, can only be mounted if one is motivated by a commensurately enormous desire for a greater grasp of truth, at whatever cost.

The costs are great, partly because co-members of one's own culture might not take this distancing and "objectification" of the givens of their world lightly or even kindly, but mostly because it bucks the stream of one's own desire, as a mammal, to belong, to be in proper relation with one's peers. All mammals share this very deep requirement for relationship. (I refer you here to some of Bateson's writings on the cybernetics of human and cetacean social systems, for starters.)

If you really want answers to your questions—can there be success without competition, what metrics for success can there be other than dominance over one's peers—I suggest you become acquainted with some of the varieties of culture and begin the struggle to understand, first that alternative perspectives are possible, then, beyond that, perhaps that they can have genuine validity, and maybe even that the alternative perspectives are *not* in competition: one does not have to be

proven "best." (Note that this judgment of "best" cannot possibly be bestowed without first taking a point of view, and that amounts to a pre-judgment that one's point-of-view-for-the-sake-of-judgment is in fact the best. One might believe that this "neutral" point of view is in some way set apart from the set of perspectives being adjudicated, but that is only the gesture which cements the prejudice.)

What one can achieve is not "objectivity" (one of the illusions spawned by the conviction that one is/has an independent, separated ego), but rather the ability to recognize ambiguity and work constructively with it. Think of the now-familiar gestalt-shift images, like the black vase which turns into a pair of white faces nose to nose and then back into faces, or a 3-D drawing of a cube, or Escher's work. A useful initial hypothesis is that *everything* is ambiguous, that is, capable of alternative interpretations from alternative perspectives.

The place to start is becoming better acquainted with the work of those who have tried to understand other cultures, workers in anthropology for the most part, in subfields like ethnology and the ethnography of speaking.

What might an alternative metric for the relative "goodness" of different cultures look like? We have to clear some confusions out of the way first. Recall that the unit of survival in biological evolution is not the individual, but, minimally, the mating pair. Among mammals, survival of a more extended group is the focus. (The mammalian emphasis on relationship I noted earlier is both an outcome and a contributor to this—what the Buddhists call mutual causation.) "Survival of the fittest" very definitely does not reduce to survival of the fittest individual. Indeed, individual fitness as measured by likelihood of mating with progeny which survive is well correlated with the individual's contribution to survival of the extended social group which provides a matrix supporting survival of the mating pair plus progeny. Darwin emphasizes in his *Origin* the importance of cooperation as being at least as important as competition, probably more important.

Nonetheless, "social Darwinism" followed the publication of the *Origin* essentially as justification for conservative social and political agendas which included racism and sexism as unexamined tenets, as justification for destruction and forcible assimilation of "primitive" peoples for their own good. When most of us hear the expression "survival of the fittest," we assume this social and political analogy to an erroneous view of nature, "red of tooth and claw." Survivors of abusive parenting, in particular, have a strong emotional attachment to this perspective as a means of reconciling hatred and rage at their abusers with the ineluctable love of parents and family which comes with the package when you are born as a mammal.

Again: what might an alternative metric for the relative "goodness" of different cultures look like?

In 1941, the anthropologist Ruth Benedict gave a series of lectures calling attention to the correlation between social structure and character structure, especially aggressiveness. She compared cultures for their differing capacities to support or humiliate the individual, to render the individual secure or anxious, or to minimize or maximize aggression. She borrowed the term "synergy" (independently of the somewhat divergent borrowing by R. Buckminster Fuller) from medicine, where it had long referred to combined action. (See "Synergy: Some Notes of Ruth Benedict," *American Anthropologist* 72,1970, 320-333.)

It is clear that U.S. culture, like many of its most influential tributary cultures, is toward the low end of the synergy spectrum (though not so low as the aptly named Ik, whose dreadful degeneracy was documented by Turnbull). For us, self-interest is clearly opposed to altruism, and accounts of cultural realities for which these notions are so closely identified that there can be no distinct vocabulary for them strike many of us as the wishful thinking we might associate with fairy tales.

A simple example: Hopi and Navajo children do "poorly" in school in part because, e.g., when the teacher sends a group to the blackboard to do a math problem, with instructions to turn around when finished, they wait until all have finished and then turn around together. How can you grade on a curve when the aim of their game is to present a flat profile, and the only way they can do that within the schooling framework is by the smart ones staying back with the slowest?

Through Benedict's students (notably Herbert Marcuse), the term "synergy" has become an icon of the "human potential movement." Most centrally what the human potential movement is about, I think, is bringing about, through change of individuals' values and cybernetic patterns of evaluation and reaction, an amelioration of our culture toward greater synergy, a change increasingly seen as crucial for our collective (and therefore individual) survival.

People stuck deeply in being isolated react with deep distrust to exercises they see as "touchy-feely" (no matter if there is no physical contact between participants). The reason, I think, is that in a low-synergy culture, group membership is perceivable only in hierarchical terms, as subordination to superiors and dominance over inferiors. The expectation of abuse following vulnerability is too strong. That a high-synergy situation *empowers* you as an individual precisely through your participation in a team or other group is beyond comprehension.

We understand power, in our culture, only in terms of dominance and submission, and competition and violence are necessary consequences of that low-synergy way of framing situations and events, not antecedent causes. Beating criminals into submission is both expression and reinforcement of a low-synergy frame of interpretation. But many

people stuck in a low-synergy perspective take it as "obvious" that human character is not susceptible of change. Given that assumption, steps to remediation (of criminals, kids in schools, employees) appear to be patent nonsense. It is a bitter and bleak reality which these folks project onto our shared situations and events. When by ruthlessness they have achieved positions of relative dominance, it is hard to keep clear of the same cognitive traps in dealing with them, harder still to be compassionate for their genuinely tragic plight, like the king in the old tale who went mad and insisted on living in the basement of his palace. ("Better the devil you know," he said, "than possible ones you don't.")

But people are capable of sometimes radical change. Even those of us who are most stuck. Were this not so, I would be a sad and forlorn man today, were I alive at all.

Rick Marken: Ed Ford says, in answer to my claim that there are no "right" principles or systems concepts: "My total experience leads me to believe that there are certain values and standards from which people make choices, and upon which people base their lives, which provide them with a great deal of peace within their family and within the community in which they live."

But are these always the *same* values and standards? Are you saying that only a particular set of values and standards leads to inner and community peace? If so, why keep them a secret? Why not tell what they are—for the sake of those people (probably nearly everybody) who seek those ends (inner and community peace).

Ed also says, "Rick, theorizing is one thing, but taking control theory into the marketplace and trying to apply it there is quite another thing."

I think we all live in the same "marketplace" (which I take as a synonym for the "real world"). We all operate in the marketplace based on theories of how it works. You imply that my theorizing is not tested against the realities of dealing with adulterers, murderers, or whomever it is you deal with whom you consider "the marketplace." From my point of view, controlling a line on a screen is as real as controlling the number of extramarital affairs one has. If the theory of control doesn't apply to everything purposeful which people do—from controlling lines to controlling crimes—then the theory must be fixed to handle it. But I don't believe that control theory is all well and good for understanding computer experiments, but inapplicable to the big mean world outside the lab. Your statement implies that there are very important phenomena which occur in your therapy sessions which control theory can't handle. What are they?

Ed also says, "I'm not on a crusade to get Rick to conform to my standards...."

I know. I don't feel that you are. You wouldn't need to, anyway—most of the values you mention sound much like what I would think of as mine, also. The question is whether *any* particular values of any controlled variables can ever be considered absolutely *right* from a control-theory perspective. Variables (in theory) are always controlled in order to control other variables. The only absolute, fixed references for variables in the model are intrinsic references—and those are references for variables which reflect the viability of the organism itself. I could accept the idea of "right" references for system-level variables (but not principles or anything lower used to control system-level variables) if you could convince me that a particular level of a system concept is required for survival of the organism.

The principles you list could be seen as a reflection of a system concept which could be described as "belief that other humans have the right to control their own perceptual variables, as long as this does not deprive me of the ability to control my own perceptual variables." I guess I agree that, when you take a group-survival perspective, there could be "right" references for non-physiological controlled variables—variables which don't have to do with individual survival. But I do think that 1) these "right" references must be for variables at the top of the hierarchy (system concepts) and 2) the consequences of selecting "not-right" values of these references is not necessarily a problem for the systems adopting these "wrong" references. I think this is what we have in the so-called "psychopath" or "sociopath." This is a person who is perfectly well-organized to control system concepts relative to references which are set at the "wrong" level. These individuals experience little internal conflict—but create enormous external conflict—by pushing strongly and effectively against the variables which others are trying to control.

But remember, in order to keep perceptions of system concepts at the "right" levels, it will be necessary to vary references at the lower levels, and this means changing *principles*, if necessary. Bill Powers said it well: "I can even see merit in some of the system concepts implicit in various religious beliefs. Love thy neighbor is a pretty good principle, especially if the neighbor is me. I'm even willing to take it on as my own principle, within reason, because it seems to fit with a workable system concept of a society of human beings. But I don't think it's going to do anyone much good if it's taken as a command from God. If you take it that way, you will never try to work out *why* it's a good idea to love your neighbor. So you'll never grasp the system concept within which this principle makes sense. You might even conclude that in order to love your neighbor, you had better stay in the right neighborhood and not let inferior unlovable people move next door." The last part here is the important one—principles do vary in order to preserve system

concepts. Look at what happened to some of the nicer principles (what I thought were the principles) of early Christianity; things like live a simple life, the meek shall inherit the earth, it's easier for a rich man to get through the eye of a needle than through the gates of Heaven. Well, there were some system concepts which demanded some conflicting principles. We now live in a Christian, capitalist country where it's a positive virtue to work hard to get rich.

If there really are certain system concepts which are better than others (for group survival), then it might be helpful to try to articulate what they are, rather than claiming that certain principles (which are used to control these concepts) are absolutely correct. The latter could prove problematic for individuals. I bet that most of us who are in this discussion are controlling for the same level of one very important system concept—the "mutual respect" concept which I tried to articulate above. There are likely to be slight differences in the levels of certain principles which we all set in order to control that concept; for example, I believe it is perfectly possible to control that system concept by controlling the principle "trust in the Lord thy God" at many different reference levels. For whatever reason, the level at which I control that principle is different than the level at which Ed or Joel (I think) controls it—but I bet we all end up perceiving about the same intended level of the "mutual respect" system concept.

Joel Judd: Rick says, "The question is whether *any* particular values of any controlled variables can ever be considered absolutely *right* from a control-theory perspective." This is the point (I think). For a model of a control system, the answer is no. For a control system cum human being, I'm not so sure.

Rick again: "I could accept the idea of 'right' references for system-level variables (but not principles or anything lower used to control system-level variables) if you could convince me that a particular level of a system concept is required for survival of the organism." Isn't there more to (human) existence than just surviving, though? The remarks about *group* existence which followed the above make sense. We are creatures of society, not individuals.

Rick: "But remember, in order to keep perceptions of system concepts at the 'right' levels, it will be necessary to vary references at the lower levels, and this means changing *principles*, if necessary." This is why I asked about "variability" with regard to system concepts, as well as the origins of system concepts, and their developmental time frame. I can see principles and other levels varying around unified system concepts.

Bill Powers says, "But I don't think it's going to do anyone much good if it's taken as a command from God. If you take it that way, you will never try to work out *why* it's a good idea to love your neighbor. So you'll never grasp the system concept within which this principle makes sense." This sounds like one of the main objections to many religious practices: unquestioning compliance. As adults, we like to analyze (well, at least most of the cultures with which most of us are familiar do) the things we do. But for the unquestioning, naive, "suckers," and children among us, it would seem that principled action, generally directed by someone more mature, is one of the ways, if not *the* way, to develop system concepts. And so churches and schools and TV and friends and politicians and families all try to instill in us their standards.

Rick again: "If there really are certain system concepts which are better than others (for group survival), then it might be helpful to try to articulate what they are, rather than claiming that certain principles (which are used to control these concepts) are absolutely correct." I think this would be fruitful, for two reasons. One, as I asked before, teachers and others are doing this all the time anyway; are we all satisfied with such influences? How can this issue be addressed (if not providing specifics, then increased awareness of the mechanisms at work)? Two, it seems like these could provide testable hypotheses.

Bill Powers: Joel says, "As adults, we like to analyze (well, at least most of the cultures with which most of us are familiar do) the things we do. But for the unquestioning, naive, 'suckers,' and children among us, it would seem that principled action, generally directed by someone more mature, is one of the ways, if not *the* way, to develop system concepts." This, too, is the way in which "someone more mature" gains converts and exercises power, regardless of the merits of that someone's system concepts. Someone has to take responsibility for what is taught. I don't think that the development of one's own system concepts is optional. Without them, principles are chosen at random or at the whim of any persuasive person. Autonomy requires not only that you *have* system concepts, but that you have the ability to modify them and acquire new ones which enhance your prospects for controlling what happens to you. Nobody else knows how a given system concept will interact with your other system concepts. The ultimate criterion for a "right" system concept is one which fits internally with all other system concepts, both directly and in terms of the required lower-level goals and actions. I believe that there are natural physical and logical constraints on which system concepts will prove best. In a society composed of autonomous control systems, only certain ways of living together will enable individuals to seek their own conceptions of the good without acting on other people in ways which frustrate that very seeking of the good. There's a lot of latitude—it's probably easier to talk of ways which don't work and the reasons why they don't work. Control theory gives us a pretty good idea of what those reasons are, particularly if we assume that people will normally try to reach agreement on system concepts (the most obvious way to avoid conflict). Lying, for example, gives other people an incorrect picture of the effects of their actions (when they must rely on communication). A society which accepts lying under its system concepts will weaken or destroy everyone's capacity to control cooperatively.

All of the deadly sins imply principles which, if allowed under a common system concept, destroy the organization endorsing them. The reasons are neither subtle nor complicated. All of the commonly recognized sins create conflict with others, and others' attempts to prevail in their own processes of control will counteract one's attempts to reach the misguided goal. From the greedy, it will be taken away. Who lives by the sword will die by the sword. Give Caesar what he wants, and he will stop bugging you. If someone compels you to walk one mile with him, go cheerfully and chattily for two miles, or however far it takes for him to be sick of your company and order you to go where you wanted to go in the first place. All good control-system advice, for someone who understands the concept of a control system.

System concepts can, of course, be proposed and taught. But someone has to accept the proposal and the teaching, convert them into a real internal way of perceiving and acting, and test the result against direct experience to see if it actually works as advertised. Unfortunately, we can't pass system concepts directly from one brain to another. What is understood is never, at first, what is meant. As adults, we always begin with an organization which works under *different* system concepts and controls *different* perceptions. The new always hooks up to something familiar at first. The greater the novelty of the new idea, the more unhooking has to be done; the more radically will the initial understanding change before the learner finally feels the lightning bolt and says, "Oh, my God, is that what you meant?" (And answers, "Yes, of course it is.") At that point, of course, it doesn't matter any more how the system concept got in there. Or where you think it originated.

Ed Ford: Rick Marken asks, "... always the *same* values and standards? Are you saying that only a particular set of values and standards leads to inner and community peace?"

As I understand control theory, all concepts are created from lower-level experiences. Since our individual experiences vary as we grow and continually create perceptions, when we create concepts, not only do we create them according to our own individual personal goals and from our own created memories, but we create similarly named concepts from a variety of differing experiences. Thus the problem of trying to create a similar "understanding" of commonly understood system concepts. An additional problem is that when I attempt to describe my concepts at principles and program level, I assume that the thoughts I generate and the words which flow from me are going to be the same as those perceived and created in the receiving living control system. Obviously, they aren't.

With this in mind, I address Rick's question. I think there are values and standards which lead to individual and community peace. The problem is taking those system concepts and setting them to standards and criteria which are universally understood and applicable. I am not, I repeat, not talking about revealed truth. I am talking about my attempt to arrive at some system concepts, priorities, and standards from which actions can be taken such that people can live a more satisfying way. For example, my ideas of commitment and quality time as I've defined them seem to work well universally with couples and in parent-child relationships. The standards and criteria I've set seem to lead couples to an experience of intimacy which provides the kind of satisfaction which satisfies their internal idea of happiness with another. Thus, I've been able to help others achieve what seems to be a goal common among the variety of people with whom I work. Obviously, describing that experience is like your wife trying to explain to you what it is like to have a baby. To those who've had the experience, no explanation is necessary, to those who haven't, or who can't, no explanation is possible.

A recent workshop participant told me, "Having read your books and listened to you speak, I get a certain sense of where you're coming from." And that's my problem. It's hard to describe a system of ideas (system-concept level) in lower-order terms and have it adequately understood, not because of the listener, but because of the way we're designed, especially due to the variety of experiences (or lack of) we've had from which we have created similar words. "Love your neighbor" and "respect the rights of others" are great ideas. I shudder to think of the millions who have suffered from the hands of those who have claimed to live by those ideals. But in my own way, if I can help people achieve a similar experience which brings lasting satisfaction, I think I have broken ground toward finding universally acceptable system concepts.

Rick says, "From my point of view, controlling a line on a screen is as real as controlling the number of extramarital affairs one has."

I agree that the theory is the same in both instances, but humans deal with each other primarily at the highest orders, and their purpose for controlling perceptions not only varies, but is far more difficult to define and understand, and a lot more complicated to deal with. I have found control theory and the perception of humans as living control

systems to be *the* single most important tool for helping people deal with their conflicts and finding satisfaction in their lives. The strategies I've derived from this theory boggle my mind. In fact, I no longer see myself as a reality therapist. Reality therapy is only a small piece of the control-theory pie. As one teacher said after a two-day workshop, "Control theory gives the counselor such a broad understanding of the client. It gives me so many more options and allows me to explore so many more ways to help people deal with their problems." It is hard for someone who doesn't do what I do and isn't faced with the complex human problems with which I deal (experiences) to perceive how control theory is so useful in the area in which I deal. And, I might add, it is hard for this social worker to understand the complex world of ideas and concepts with which you theorists deal.

I think the test for system concepts is the harmonious cooperation they provide, regardless of the environment (the last four words were added to deal with Bill's concern about the application of the principle of loving thy neighbor).

The bottom line in all of this is that when you deal with system concepts, you are dealing with an area which, by its very nature, isn't easily understood. The variety of experiences which define this area vary from one living system to another. And how another's system concepts are prioritized, how their standards and criteria define their limits, and the variety of actions all make this an area which is easily confused and hard to deal with, much less understood. The best example for those who are married is trying to understand one's spouse. (I gave up trying to understand my spouse 17 years into my marriage and things improved remarkably. Now, after almost 41 years, I'm still very happy.)

Rick Marken: Ed, you're right. I do have a continuing agenda with respect to any system concept—whether you call it a religion, a science, or an attitude. That is *falsifiability*. I think there is one thing which distinguishes the people I have admired in my life: the willingness to propose a brilliant (and usually unpleasant) thesis *and* the willingness to subject that thesis to test. I, personally, think that control theory (as articulated by Bill Powers) is such a theory. I (obviously) enthusiastically support that theory—and work hard to promulgate it *and* test it. But I am prepared to see it falsified (in the sense that a better theory is needed to account for data that control theory, as presently articulated, cannot handle). I don't know *any* person, for whom religious theories are part of their understanding of the world, who is really willing to *give up* his or her theory based on evidence (whatever that might be; evolution seems like a pretty strong rejection of a large part of the Judeo-Christian model of genesis). I don't mean to sound anti-religious; I'm just anti-dogma. And religion (usually; I can't think of an exception off hand) is institutionalized dogma.

Would Ed or Joel be willing to abandon, say, the theory of "transubstantiation" if it were proved to you (based on tests you or others adhering to this theory accepted as tests) that the predictions of this theory don't pan out? That's not what religion is about—I think.

I am a bit fed up with the idea that religious principles don't conflict with principles derived from other system concepts. They do. That doesn't mean that religious people aren't nice people (often), but their little logic traps can be a real pain for the soft of brain.

I don't know what the best system concepts and principles might be (though I agree that we have to have tacit agreement on some to make it as a species, probably). But it's time to admit that unfalsifiable beliefs are internal conflicts, and as such are "software" cancers as deadly (for many hosts) as the hardware versions. I hope that grace will eventually be seen as the acceptance of the fact that you *might be wrong* (anyone—myself heartily included).

Just one more little point. I believe that control theory, like evolution, treads in realms which were once the exclusive province of religion. I know that there have been (and are) good religious psychologists and neurophysiologists (MacKay, Sperry, and Eccles are examples). But my contention is that their religious preconceptions prevented them from making any really deep contributions to the field. Their latter days were spent rationalizing away the religious implications of their own work. Control theory, which gets really deep, is not (I argue) going to be taken very far by one who finds implications of the theory untenable—in principle.

Bill Powers: The question control theory should try to answer is what belief (firmly held or otherwise) is and how it works, not which belief (or non-belief) is best. It is the same sort of question which applies to control of limb position: what position control is and how it works, not whether we should use our arms to hit people or pat them on the back.

When people are controlling for the truth (high reference level) or falsity (low reference level) of any belief, they will resist disturbances which suggest a change in the level of that belief. If two people who maintain different reference levels for the same belief try to alter each other's perceptions in this regard, a conflict results which is evident in their communications (if not in their face-to-face interactions).

Rick Marken: I'm sorry to have given some participants in this discussion the impression that I was arguing for the superiority of control theory over other beliefs. Obviously, I did not make myself clear, or what I said was a disturbance to some of your beliefs (which is informative in itself about the nature of one's own high-level control systems). I have no interest at all in proving that one belief is better than another—unless those beliefs are models and I can "prove" them by testing their predictions experimentally (where prove is an inductive test, not deductive proof).

What I was trying to get at is the question of whether people can really maintain apparently conflicting beliefs. Do the beliefs really conflict? I am also trying to probe around, testing which beliefs people are controlling for. I don't care what those beliefs are (I only care what I believe, of course) but I am interested in trying to show that these beliefs *are* controlled variables.

These kinds of controlled variables are hard to study. We have not done much work on them. So I use CSGNet sometimes to play with ideas about these "high-level" controlled variables. The problem with this, of course, is that people have "strong feelings" about these beliefs. So it is hard to talk about them as just controlled variables. There is the perception that, when a person discusses a particular belief, he or she is trying to prove its merits. I do like falsifiability; I don't like religion. But that is me. I am not trying to convert anyone. My own belief in control theory makes me realize that my beliefs can only work for me—not for others. Unfortunately, beliefs (and everything else which is human) are part of the control model—and people have trouble treating these aspects of the model as just another set of controlled variables (like arm position).

One last point. I think that the control-theory model itself does have implications for certain belief systems. That's just the way I see it. That does not mean that I think, therefore, that the control-theory model is superior to these beliefs. No one is forced to do or understand control theory. But I believe that, if one chooses to try to understand life in the context of the control-theory model, certain other beliefs are, indeed, impacted (I won't say which or how in the hopes of staying out of trouble). But it's like astronomy. If you want to play by that model, then you have to give up belief in, among other things, a flat earth at the center of the solar system. Astronomy doesn't prove that it is superior to a belief in a flat earth (under most circumstances), but if you accept the assumptions and rules of the model, then a belief in a flat earth along with it is difficult. Control theory does have implications for certain cherished beliefs about the nature of life, but that doesn't prove that control theory is superior to those cherished beliefs. It depends on what *you* want to control for.

Bill Powers: Control theory tells us that all rules, conventions, laws, and so on (by which people actually live) must, in order to be effective, exist as reference signals at the appropriate level inside each person. They do not reside outside people, even when they are written down or present as physical constructions. They are not implemented by any mechanism outside individual human beings. There are no natural control systems outside the individual human being. Not even in a society.

At any given time, a society is made of adult members and of young people getting familiar with it. What they have to get familiar with is not some external structure, but the other people in it, and the way the other people construe and use those external structures. The adults teach the children by example and by explicit instruction how to use language, how to use a knife and fork or chopsticks, which side of the road to drive on, how to get money in an acceptable manner, and so on. Each adult teaches these things out of a single person's understanding of them.

This teaching includes teaching what to perceive as well as the reference levels for the perceptions. If a child construes the world in some novel way, the adults will not see any sense in the child's control actions. There will be both active and passive pressure to see the world in the conventional way and to learn to control conventional perceptions.

There are, however, variations from person to person. The child doesn't get the same story from everyone. Also, children come up with novel ways of saying and doing things, and adults pick them up because they're funny, insightful, and refreshing. The children don't always get corrected. Sometimes they are allowed to introduce variations of their own. I still love "far out!"

All of this goes to show that there are no "social reference levels." If there were, there would be control actions which always brought the social variables back to the same form. What happens instead is that all pressures to change the social forms are resisted (because they create errors in individual people), but at the same time the perceptions in individuals gradually change, and the reference levels chosen from among them also gradually change. In the long term, there is no resistance at all to social change; that is how we know that there are no external social control systems. There is inertia, but no control.

In the short term, people learn and retain ways of perceiving and controlling. Each person comes to an understanding of what is worth perceiving and what is worth controlling. The main teacher is conflict. Conflict frustrates control and causes a waste of energy. So people naturally modify their own goals and perceptions to minimize conflict with those around them. When they try to deviate too far from social norms, they create errors in many other people. Each other person, in opposing the disturbance, pushes back in some fashion against the deviant behavior. The deviant person feels the sum of a thousand mild resistances as if it were one powerful sanction against the change. A thousand points of light make a searchlight.

This is what creates the inertia. In order to minimize conflict and maximize freedom to control, society-wide changes must always be gradual so that, in effect, everyone changes at once. No short-term deviation can escape what appears to be coordinated social pressure against the deviation. But the only coordination necessary to achieve this effect is that each person resist what that person perceives to be an error. This resistance does not even have to be exerted directly against an individual's attempt to reach a goal. Others are affected only by side-effects of control behavior. All that is required is for one of those side-effects to disturb some variable which is important to another person.

To this natural appearance of coordination of opposition, we can add, of course, deliberate coordinations of opposition to deviants such as carried out by police forces, schools, and scientific disciplines (appropriate word!). This more organized way of resisting deviations, however, works exactly the same way: one person at a time. There is simply a more conscious attempt to reach alignment of goals among the enforcers. The result is also a narrower definition of what amounts to a deviation. I suspect, too, that the time-scale of change is shortened rather than lengthened by this sort of deliberate coordination. The reason is that when people try to define their goals very narrowly, and to resist strongly the slightest deviation from them, the attempted coordination is more likely to turn into dissension and eventually into fragmentation. Fanatics necessarily end up as loners.

Language can appear to be a thing, a universal force or rule, without actually being that. Of course the same argument applies to any apparent social ordering influence which seems too long-lived to be associated with individuals. My argument is that individuals are entirely responsible for such things, but that in their need to avoid direct conflict and in their natural resistance to disturbances, they seem to be under the control of something larger than themselves. In fact, they are: they have no choice but to go on being control systems.

Bruce Nevin: Bill, I have to say that, ideologically, I find your position most congenial. I have been an anarchist for as long as I can remember. But the tendency to hypostatize the constructs we make of Family, Society, the State, etc. is pervasive, and not to be dismissed, I think, without plumbing its depths. And it is precisely those agreements we have no memory of making which are problematic for our coming to consensus.

I believe my response will use an analogy between the relations among people as control systems and the relations among control systems within people. The top level in both cases is reference values, and it is the reference values which are socially inherited.

There is no forest, only the trees, eh? There is no society, only the people (control systems) in it... there is no person, only the control systems in him/her.... So long as we don't shift from one kind of thing (control systems) to another (neurons), you might be able to get away with this reductionism.

Rick Marken: Bill Powers says, "There are no natural control systems outside the individual human being. Not even in a society." Except, of course, other people. But that was the whole point of his post. Social rules are the result of the mutual interaction of hierarchical control systems. Bill also says: "In the long term, there is no resistance at all to social change; that is how we know that there are no external social control systems. There is inertia, but no control." A big question is whether the drift in references and perceptions which we see happening historically is basically random or whether it is constrained, to some extent, by intrinsic references common to *all* people. My impression is that there are some general constraints on the inertial change in reference levels. I also think that technical developments have made certain directions of drift more likely, as an accidental side-effect. Birth control pills and safe surgical procedures have surely made it easier for references to change regarding sexual mores, gender roles, and abortion. The resistance to these changes produced by other control systems is obvious. But, nevertheless, a drift has occurred and, I think, will continue toward more "liberal" sexual and sex role references (AIDS notwithstanding). And I think this change is being eased (or exacerbated, depending on your reference setting) by the technical developments. Of course, this is also producing more strenuous resistance by those with "conventional" references. So maybe the "push" by these technological "lubricants" to change is offset by the efforts of the control systems with conventional references.

I think there are "natural" disturbances which contribute to the difficulty of controlling relative to "deviant" references. For example, societal references regarding acceptable levels of sexual activity are surely enforced, in part, by the unpleasant consequences of adopting "deviant" levels—i.e., you get venereal disease or become pregnant. To the extent that inventions like penicillin and the pill reduce the chances of such consequences, more people will be willing to test new references for sexual activity. They will still get resistance from the majority, but more and more of the new generations of control systems can try the new references with not only fewer natural consequences, but less resistance from the fewer control systems around trying to defend the currently accepted references. So this might be the way technological change can "push" social references in new directions.

I guess that similar kinds of developments ease changes in references for language rules. Groups relying on written-language communication technologies will experience an "easing up" in certain directions of language-reference inertial change. Groups that rely on auditory language (street communication, TV, radio) should experience easing up in other directions. I think this is what we observe (though I think it would have been difficult to predict). Certain usages seem to be accepted in auditory communications which are not accepted in writing. I argue that this is a simple example of a technology (writing) influencing the ease and direction of inertial reference drift.

My gosh, I think I just argued for a dynamic attractor model of social rule drift. Yikes!

Ed Ford: Regarding Rick's remarks: Having spent the last 25 years in a counseling office (among other places), I would say that presently there is emerging a growing change in references and a perception of "natural consequences" not anticipated by those seeking "liberal" sexual and sex-role references. The harsh reality is that all this new sexual activity has made creating satisfying relationships more and more difficult. As one physician client remarked recently, "I hate dealing with post-orgasmic depression." It is my experience that humans learn more from their self-created internal conflicts, that is from the consequences (other than social pressures) of their attempts to control a desired perceptual variable to a set reference signal, than from having to deal with social pressures (disturbances), whether from home, cultures, organizations, or whatever. Sexual activity involves one living control system dealing with another, but that specific activity involves only one very narrow and restricted aspect of our many relationships. Many are finding that kind of activity detracts from, rather than enhances, relationships. And there are many reference signals which go into making relationships. In human relationships, it is the value you see in others and their perception of value in you which brings the greatest enhancement. The question is: Are the reference signals we set for building relationships really bringing long-term satisfaction? My experience in working with young people is that there is a growing trend toward a more conservative view of sexual activity, in spite of all of the great scientific advances.

Perhaps what really pushes social references is whether they bring continuing satisfaction over a long period of time. Occasionally, cultures test those references. Over the past 30 years, we seem to have done a lot of extensive testing in this country. Are we as a people a lot happier and more satisfied with our lives, and especially with our relationships, than we were 30 years ago?

Bill Powers: To Bruce Nevin and Rick Marken, with an implied challenge to the social scientists.

Reductionism results when you ignore higher levels of organization: Bach spent his life drawing little slanting dashes or dots with vertical stems on pieces of lined paper. I'm not sure what the name would be for the sin of extrapolating a theory to the point where it turns into a metaphor: synthecism? Synectady (in New York)? Or argument by analogy?

The control systems inside a person consist of specialized input and output functions, with comparison processes variously achieved. These structures appear to exist independently at many levels in the brain. One of the levels organizes a person in an encompassing way we experience as being "a person" (and recognize in others through what we perceive of organization in them). Each person is a structure of interdependent systems and many levels.

In a society, there are no people who devote themselves to one level of function only, or to one specialized function in a single control system. It is impossible for a person to behave in such a way and live. So the control hierarchy in a single person stops at the highest level in that person; there is no way to continue it to a higher level outside the person. It probably continues downward through the biochemical rather than behavioral branch, however (starting roughly at the level of the hypothalamus), all the way to the inner working of the genome.

People have attempted to form societies organized as artificial control systems. As a society is envisioned by many people (including some in the White House), there are social mechanisms for monitoring the actions of individuals (informers, covert and overt investigative agencies, panels of experts, news media); comparison processes for detecting deviant behavior (definitions of disease, insanity, torts, crimes, obligations, duties); and procedures for correcting deviant behavior (penalties prescribed by law for each wrong or crime and each omission of duty, treatments indicated for each deviant mental condition, illness, or incipient departure from health). This system is supposed to operate automatically because the specifications for all of its parts are written down—and fairly, because it is automatic and applies uniformly to everyone for the benefit of society as a whole. Like any control system, it is supposed to control through opposition to disturbances, the opposition adjusting from mild to overwhelming as befits the size of the disturbance.

This concept of a society is a natural mistake born of each person's need to have control of the experienced world. This mistake has been made over and over. Some people have tried to devise utopias and anarchies to get away from the flaws of the social-control design, without remarkable success except perhaps on a very small and localized

scale. But most people are persuaded that we need law and government and medical treatment and the like: social control for the good of the many.

The greatest problem with this concept of an artificial social-control system is that it comes into direct conflict with the basic nature of the individual, which is to control himself, herself, or (if living) itself. So each individual breaks the laws and flouts the rules of health in small and large ways every day, and devises means of not getting caught. The voters vote for control of other people and against control of themselves. The powerful maneuver to obtain maximum freedom for themselves and minimum freedom for the rest (particularly for those who would also like some power). The wealthy try to free themselves from restrictions on how to spend what they have and how to accumulate more, and they try to set conditions to prevent others from taking back some of the wealth.

Each person wants to use this vast automatic machine as a means of controlling what happens to himself or herself. Thus, individual freedom is in constant conflict with the social-control system which has been set up for the good of society. The greatest flaw in this concept of an artificial social-control system is that it is not and cannot be automatic, running independently of any individual's whim. In fact, it is run by individuals and is constantly subject to individual whims.

There is, in fact, no System. I said this in my 1973 book, and I still believe it. You can walk into any bureaucrat's office, and all you will encounter is a person. When you stand before a judge, you do not stand before the law, but before a person who listens to you (and, too bad for you, others) and tries to make sense of everything in terms of what the judge remembers and understands and wants of the written law. A different judge (or jury) will hear differently, understand differently, want differently—and decide differently. The clerk at the driver's license desk can make it easy for you or endlessly difficult. The county assessor can be reasonable or implacable. The System consists of people, all of whom are different. You will never encounter anything but the people and their individual wants and desires.

The worst nightmare of anyone who has grown up in a free society is to lose that freedom, that independence from external control. As examples of threats to freedom, to what do we point? To dictatorships, whether of the proletariat, the armed and dangerous, the religious, the politically ambitious, or the deranged. And what is a dictatorship? It is a system devised so as to exert social control exactly through the kind of automatic control machinery described above. What we fear most is law applied blindly and without regard to circumstances, by the book; force applied without regard to our wishes; goals imposed on us without our inner acceptance; duties demanded of us without consideration of what will satisfy us; loss of control over our very bodies, our very Selves, our very lives.

Even the freest nations in the world are still hanging onto the old forms, the old conviction that we need an automatic social-control system which is not just human interaction but something larger, more protective, more powerful. Yet the freest nations are what they are precisely because the individual's need for autonomy has prevailed to some degree over the very system which people are convinced is needed to protect their freedom—and which, in fact, might be needed to protect them against other people who would impose their rules even more strictly and thus go even more harshly against autonomy. But this is not where we are headed—toward the perfect social-control machine. We are headed inevitably toward something else. I can't say what it will be—we have yet to work it out. Understanding that social-control systems are illusions and threats to freedom is the first step in working it out.

Rick Marken: Ed Ford says, "Perhaps what really pushes social references is whether they bring continuing satisfaction over a long period of time." He might be right. I was just suggesting that technologies might create a groove ("push" was probably the wrong word) making it easier for certain references to change in one direction rather than another. I used the example of sexual activity not because I am in favor of a particular direction, but because it seems that there has been a change in the majority "reference" for, say, "women's role in society" which seems to have been made particularly feasible by certain technologies. Perhaps there are fairly universal "intrinsic references" which prevent the inertial drift in references from straying too far. This seems like a reasonable possibility—societies have tried lots of different sexual mores (references for sexual principles), but none that I know of settled on an average norm which encouraged, say, incest (except among a select group of individuals, as in the royal families of Hawaii). So I am just suggesting that the "inertial reference drift" discussed by Bill could tend in one direction rather than another at particular points in history as the accidental side-effect of the development of certain tools. I think James Burke was making this point in his marvelously entertaining "Connections" series.

I do not believe that these technical developments act as some kind of "invisible hand" acting as a reference signal outside of people which specifies how they should change. I believe, as Bill Powers said in his latest post, that the only references for how things "should" be in society exist in the individual members of that society. Actually, this concept once lost me a job. I was interviewed for a position with a law firm many years ago. I guess they wanted a psychologist as an expert

witness or something. Anyway, one thing they asked was whether I believed that "companies" are entities in themselves or just the sum of the people that make them up. The goal was to see whether I believed that people (like the company presidents, vice-presidents, etc.) are liable when the policies of the company lead to harm. I said the latter—since "the company" is defined by the understandings, goals, and perceptions of its individual members—and never heard from them again.

Bruce Nevin: Bill, what you are resisting is a notion of suprapersonal control systems. But you don't have to assert that to talk about structures of social convention.

I assert that language has structure which can be observed and studied not only in the outputs of language users (speech, writing), but more especially in the results of testing for what it is they are controlling for in their use of language. Assuredly, they can control for reference values of any kind only after having assimilated them into their own control systems. My only claim is that there is something there to be assimilated, pre-existent in the linguistic outputs of other language users and in their resistance to perceived error.

This structure is there because people cooperate to learn it, assimilate it as their own individually, and maintain it as their own collectively. By this last, I refer to the fact that control of language and dialect is one very important means by which people identify the membership or nonmembership of people in groups to which they refer as "us" and "them" and "we" and "you."

Yes, this structure exists in the language-learner's world of experience only by virtue of other people's individual control of perceptions. I do not deny that. Nor do I assert that there is some suprapersonal control system governing it. I only assert that it is there. It is present not just in that individual who is currently teaching the child by precept and by example, but in a number of individuals on many occasions, so that the example is not isolated but rather is an example precisely of agreement and communal synchrony. Individual idiosyncrasies are also interesting, and children learn from them, but it is the fact that they are shared and indeed must be shared to function which gives special appeal and importance to structures like those of language. The structural facts of a language are not rare, they are expectable and expected.

It is there not because some superordinate control system sets reference values to which individual people are compelled to conform, in the way a control system for a certain kinesthetic sequence must, if stimulated, control for repeated nodding of the head and cannot do otherwise. Such compulsion is inimical to our nature as autonomous control systems. The coordination among individuals must be a voluntary agreement. We agree to refer to that blue, sometimes cloudy expanse above as "sky" and not as "aseH'la." We do not remember having agreed to it, nor do we remember what must amount to many hundreds of thousands of other agreements by which we came to be persons recognizable as members of our families, various groups of friends and cohorts to which we have belonged, and other social groups and systems. Where with our fellows occasionally we perhaps forged new agreements, we did so by adapting what was already there, not by creating anew like the mythical Adam assigning names and attributes in the Garden.

Nor do we feel free to undo such agreements. We could invent new words for things only at the expense of dismooring ourselves from the linguistic continent of English and all its inhabitants, and that, for a great many reasons, we choose not to do. Having made that greater choice, we find ourselves not free with respect to the lesser ones which make it up. It is precisely so for the child learning the language. Given the commitment to participate in the ways of being human which are normal for his or her family and friends and community, it is as though all those others could reach in and set reference values within his or her control hierarchy for "sky" and myriad other matters of arbitrary but (crucially) shared convention. It is as though he or she actively offers up these comparators within him or her to be set by others around him or her. He or she is alert for evidence of disturbance, acutely observant and mimetic, and during the early years of most active language acquisition is quite amenable to explicit correction, especially by example.

In the process, over many, many generations of many, many people *individually* "avoiding direct conflict and resisting disturbances," they have *collectively* created structures which are not mere dissipative systems like a vortex or a sandpile, precisely because each participant (unlike grains of sand) has agreed to participate and controls for participation. These systems of agreements, in various aspects of language and culture, are of great complexity, elegance, and beauty, and are most worthy of study and appreciation. As Sapir observed, they are like collective works of art, which some individuals are more able to display and use than others, and which no individual holds entire. Through them, individuals not only make known to their fellows their membership, but what their contribution in membership might be.

Control. theory provides a crucial moiety which has been missing from the study of what human beings are and do, but it still must be seen as incomplete, as intersecting another perspective. This other perspective is concerned with what human beings externalize among themselves for the sake of relationship with one another.

Bill first said: (A) "Control theory tells us that all rules, conventions,

laws, and so on (by which people actually live) must, in order to be effective, exist as reference signals at the appropriate level inside each person." Then he said: (B) "They do not reside outside people, even when they're written down or present as physical constructions. They are not implemented by any mechanism outside individual human beings. There are no natural control systems outside the individual human being. Not even in a society." (B) does not follow from (A). It is simply asserted. I say: (C) They do reside outside people, because that is where people put them, and they are implemented not by any mechanism outside individual human beings but precisely by those individuals, as autonomous (not independent) control systems *voluntarily* conforming to them. They do this for the sake of cooperation with other human beings. They do it because if feels good to belong—because there exist control systems which they have in common with all mammals (said Bateson) which control for relationship.

The study of these structures to which people agree to conform is not merely the study of behavioral outputs. Bill's critique of stimulus-response theories makes it clear that we can only learn about them by testing for control, and that is precisely what the techniques of linguistics do. It does not present anything like a statistical average as "results," because the objects and relations studied are precisely defined as reference values for individual control. It is different from the study of control of, say, locomotion because people have placed arbitrary constraints on the degrees of freedom normally available for control. They have done this stylizing and conventionalizing so as to differentiate membership from non-membership, relationship in the social sense from relationship merely in the sense of physics. This range of choices is in, the rest is out.

These things have been abused as matters of coercion, and will be, but they are not inherently so. They can be matters of play and mutual enjoyment, and often are.

What is lacking for there to be suprapersonal control systems is a means for *setting* reference values from outside the person. The agreements I mentioned above depend upon means for communicating or transmitting or advertising reference values, but the setting of these values is a matter of voluntary (or coerced) choice in each individual.

As Camus put it, we are condemned to freedom. We act as though we don't like our freedom; we seem to give it away as quickly as possible. For the most part, we do so for the sake of participating with others in some unity larger than any of its participants. As we grow, we become more discriminating.

Rick Marken: OK, Bruce, language structure might be a controlled variable. But that controlled variable cannot be seen just by looking at the language. That's the essential point. It might be "out there"—but there are many possible structures out there. The goal is to find out what people are controlling. So looking for structure in the language itself is like looking for "affordances" in the environment or for the "reinforcing" properties of food. You might perceive interesting structures in language, but they are likely to be side-effects of what the language user is actually controlling for, just as the three-dimensional movement of the *E. coli* bacterium is a side-effect of its control of a unidimensional quantity (*E. coli* cannot perceive in more than one dimension).

Bill Powers: Bruce, I said that control systems (other than the class of devices called servomechanisms—artificial control systems) do not reside outside people. You say: "They do reside outside people, because that is where people put them..." Give me an example of putting a (social) control system outside of people, and tell me where all of the functions and signals are. Or let's make it easier: just describe the comparator to me. (I cheat. This means you have to describe the perceptual signal, the reference signal, and the error signal, too.)

What comes in through the senses? Not reference signals, but perceptions. Perceptions are reports on the (purported) current state of affairs. They are not prescriptive. Aha, says the perceptual system, I discern that this car with me in it is headed toward a tree. The reference signals, oblivious to the current situation, say that the car is centered in its lane. The comparators in the brain must take in the reference signal's specification and the perceptual systems' report and make of them an error signal which leads to action which tends to reduce the error. Without reference signals, perceptual signals imply no behavior. If you wish to crash into the tree, you can actively maintain the perceptual signals as they are. The perceptual systems will continue faithfully reporting the current situation until the moment of impact.

Furthermore, what comes into the brain must always begin as a collection of elemental stimuli which excite sensory receptors to produce trains of impulses representing intensity. The rest of the nervous system lives, therefore, in a world comprised of intensity signals. Out of the behaviors of these signals and all of the relationships which the brain can construct on them, the rest of the world comes. As we gain experience with this world (even in the act of constructing it), we record enough of it to be able to select previous states and use them to create reference signals defining intended states. Structures higher in the brain select and set reference signals for structures lower in the brain, as required for higher control processes, old and new.

When you get to the top of this hierarchical structure, you are as far as it is possible to get from the sensory periphery. The highest reference signals can be derived only from recorded states of the highest perceptual signals, or from fixed genetic information, or from the random trial-and-error of reorganization. The only way for any higher entity to insert a reference signal into the comparators at the highest level would be to drill a hole into the skull and stick an electrode through it (or to reach in through the fourth dimension or a theological loophole). It is physically not possible for the environment to adjust reference signals at the highest level. It is therefore not possible for the environment to *determine* reference signals at any lower level.

There is only one way in which a reference signal can depend on an external event. That is for the external event to disturb a variable under control at some level in the hierarchy. When this happens, the corresponding control system at that level will alter the reference signals sent to lower systems in such a way as to counteract the effect of the disturbance on the controlled variable. Those lower-level reference signals will therefore appear to depend on the external event as long as the higher-level reference signal stays constant. However, it is generally not possible for someone in the external world to know just what *other* controlled variables have been disturbed by the same event, and thus to understand all of the adjustments which are being made internally to the brain. We can predict that the disturbance will be counteracted by *some* act of the system, but whenever there is more than one act which would serve (and there usually is), we can't predict *which* act or combination of acts will be employed. Whatever act is chosen by the brain must satisfy the requirements of many control systems at many levels. Unless you have a complete map of another person's goals at all levels, you can't predict how a given disturbance will be resisted—unless, like Skinner, you arrange the environment so that only one act can have the requisite effect. Of course, all such predictions depend on the constancy of reference signals at levels higher than those involved in counteracting the disturbance.

So my objection to the idea of social-control systems has nothing to do with abstract principles or philosophy or activism. It is simply a deduction from the apparent facts of our physical construction, coupled to a model of how the brain manages behavior. Human beings can act on each other only through the exchange of chemicals and physical forces and through altering the patterns of intensity signals at the periphery of their respective nervous systems. They provide each other with experiences, but not with reference signals. I can *describe* a reference condition to you ("go jump off a cliff"), but I can neither interpret the description to you in terms of specific target-experiences nor cause you to accept the meanings in the description as your own active reference signals.

This is, I presume, how all people work—even those who work for

"The System." Each person lives inside one brain. In this brain are that person's perceptions and that person's goals. Some of these perceptions represent the output acts of other people—but never their perceptions or goals. So each person lives in a purposive system and is surrounded by other people known only through their shapes and their acts, and only inferred to be purposive.

From interacting with others, one comes to form concepts of systematic entities, system concepts. Each person does this independently and alone. As a result, the inner organization of each person takes into account the properties of others as they are visible through the acts of others. The concepts thus formed embody theories of human nature, theories about human interaction, concepts of what you're allowed to do and what you're forced to do. These concepts might have nothing to do with real human nature; they might be completely erroneous. Nevertheless, they determine what goals you will pursue in relation to other people, and what means you will employ in pursuing them. They also determine the properties you will exhibit from the standpoint of other people.

The interactions which develop among people organized in this way can be of any conceivable type. There can be negative feedback and positive feedback and open-loop relations. The entire social system can oscillate or run away or lapse into quiescence. There can be direct physical conflict. There can be loners who shun company. People can develop different customs, languages, means of livelihood, attitudes toward law and religion, definitions of fun, and styles of family living. Anything is possible: there are no overriding rules, and there is no overriding entity capable of enforcing any particular style of being.

Each person, of course, has needs and requirements. These must be met, and they play a large part in determining when a person will reorganize and stop reorganizing. Everyone has to eat, breathe, stay warm, play, think, and experience Good. So there are inner forces which are similar in all of us. But these forces are inside, not outside. The constraints they introduce work through reorganization, not though external direction.

The physical world also introduces constraints, but not purposive constraints. It is apparently true that energy and momentum are conserved, and so on. It is true that two bodies can't occupy the same space, at least if they are human. It's true that if there is less food than is required, only some people get to eat enough. And so on. Physics, chemistry, and biology create constraints within which all learning and interaction have to take place. But these constraints exist without purpose and they apply equally to all.

There is and can be no social-control system because there is no place for it to exist, and no organization external to human beings capable of carrying out its functions. Even people who think they are part of a social entity have different concepts of what it is, what its goals are, what it should be perceiving, and how it should act in specific circumstances. The cop peering in through your car window could be a liberal or a Nazi. He could be following the book, interpreting the book, or looking for a contribution to a worthy cause. He might cite you for speeding or for not having an emissions sticker, or both, or neither. That's up to him, not to the System. Only he can decide, and that decision comes out of the way he is personally, individually organized inside.

This is true of every single individual you will ever encounter in the process of interacting with the social system, no matter how impressive the building in which the individual works or the equipment he or she chooses to bring to bear on you. It is true even when people use force on you, even when they gang up on you. What they do comes out of themselves; they are responsible for doing what they do. Just as you are. Just as we all are. People can use a mythical concept of a System as an excuse, as a way of attributing cause elsewhere, as a way of unloading responsibility. But the responsibility for how you move your arms and legs, for the way you move your mouth and face to shape the sounds you utter, how you mobilize yourself for action, is yours and nobody else's. It is your responsibility not for any moral reason, but simply because your purposes determine all these things and therefore you, as a whole behaving system, are causing them.

If no individual can correctly blame the external world for the purposes presently being effected by that individual, then there is no System, because the system is manned by individuals (and womanned). It is simply the way they interact in the physical world.

Rick Marken: It strikes me that the idea of social control seems rather ideological, since there is virtually no evidence for the existence of control organizations outside of the individuals participating in society. The idea of external social control seems to me equivalent to the idea of environmental control—except that now the control is somehow exerted by collections of living things rather than by inanimate objects (like reinforcers). I think it will be as hard to convince people that there are no social-control systems as it is to convince them that there are no environmental-control systems (like the reinforcing contingencies of the behaviorists).

I guess one step toward convincing me that there are social-control systems would be to point to what you think is an example of the phenomenon of social control—then model it and see if the model behaves as expected. We already have models which show apparent social control (organized crowd behavior) "emerging" from the behavior of interacting control systems. The models have no control systems outside of (or made up of) groups of individuals.

Why do people want to believe in social-control systems, anyway? I suspect it's another surrogate "higher-level" control system up there in Heaven checking to see who's been naughty or nice. Maybe when we find that external control system, we can finally tell which group was right about who's up there.

Bill Powers: The platoon leader says "Advance!" This creates a situation requiring me either to expose myself to enemy fire or to explain to the platoon leader (and eventually the Provost Marshall) why I have concluded that it would be wiser to go the other way. I must also deal with my own goals regarding patriotism, cowardice, hesitation to do harm to others, organizational consequences of disobeying an order, helping my co-dogfaces, and so on. It's a problem—but it's not control from outside. Even the Army admits that obeying orders is controlled by the individual. Otherwise there wouldn't be any mechanism for dealing with disobedience. In general, the law considers intent a necessary component of committing a crime. Intent without control means nothing.

Suppose there were mechanisms for transmission of reference values between individuals? If such means existed, the external agency would have a problem fighting the other systems in the brain already contributing to the same reference signals. The goal structure in an individual has evolved through a lifetime of learning and interacting with the world; everything interacts with everything. Your goals serve *your* needs, not those of others. Even your altruisms have been structured to satisfy your concept of the "right" way to help and accommodate others. You can't change just *one* reference signal in the brain and expect anything but massive resistance to the change. This is a *system*, not a collection of reactions.

But individuals *can* voluntarily set certain reference signals within themselves to socially agreed values as their only or best means of controlling certain other perceptions which have higher priority to them. The setting is done by the individual. I have never said that it is done without regard to happenings in the perceived external world or without regard to other reference settings in the same individual. I am only saying that there is no way for an external agency to reach inside an intact individual and physically alter reference signals. Or no way that would work in a significant number of cases (psychosurgery?). Even if you could do this, a higher-level system would immediately restore that reference signal to its former setting, or a conflict would be generated, destroying control. Unless you broke something in trying to effect the change.

I hope I'm not being dogmatic. Dogma is stating conclusions without justification. I justify all of my statements as clearly as I can, referring to the publicly defined model from which I deduce them. As far as I can see, a "social reality" which has the same common existence for all people is inconsistent with the control-theoretic model (as well as its epistemology). If this concept is consistent with some other model, then I suggest that the other model be presented, and its properties be laid out. It would also be nice to see some tests, even if they are very simple, which the other model would have to pass to be accepted.

Bruce Nevin: In control theory as Bill has articulated it, an elemental control system (nice term!) can get its reference signal only from some other elemental control system.

The carrying out of familiar sequences and programs exemplifies perceptual control very well. The occasion for initiating one or another program or sequence is not always so clear. It appears to come out of a realm which is much more wet and leaky than the control hierarchy, a realm with which we associate emotion and empathy, intuition and impulse. The implications for interpersonal influence are considerable.

Even so, this is influence and not control. Control is compulsory. Given a reference signal with a certain value (rate of neural firing), an elemental control system has no choice but to calculate the difference between its reference signal and its sensory input. Unless some other control system has changed the connections, it has no choice but to output this error signal to the reference-signal input of one or more other control systems. One control system sets the reference signal of another.

Within a "bag of skin" you have hierarchical control, perhaps made a bit more mushy than some would like by mechanisms which can render some reference signals subject to influence. Between "bags of skin" you have influence. Interpersonal and social influence is sometimes made more hierarchically controlling than is appropriate by interpersonal coercion and manipulation. These techniques always run into problems because they result in conflict within the "bag of skin" being coerced or manipulated. Chapter 17 of *Behavior: The Control of Perception* describes this dilemma clearly.

Influence works best by suggestion. As long as it is not deceitful (which is a form of manipulation, and does not work when experience eventually gives rise to conflict), suggestion works very well. As Albert Schweitzer said, "There are three ways to teach a child: The first is by example. The second is by example. And the third is by example." Has no one in this group looked at hypnosis?

Rick Marken: What you might see in a cell is control occurring—maintaining a certain level of chemical concentration. You then imagine a control model which might produce such control—even assigning functional roles to likely cell components. But in a society you see what appear to be the components of the model used to explain control—this person functions as a perceiver, that person as a comparator, and that person as an effector—but you don't see any control being effected by these individuals as a group, do you? If so, where is it? What is being controlled by the hierarchical control system made out of individuals? And, to the extent you can identify control, is it control which *cannot* be explained in terms of the operation of each individual control system making up the group?

Bill and I are just saying (I think) that there is no phenomenon of control which is explained by imagining that groups of people are an operating control system. We also see no evidence of a social-control system external to groups of people—whatever that system might consist of. If there were evidence of control carried out by groups of people, and if this control could not be explained in terms of the individual control actions of people, then we would be happy to entertain the proposal of a "multi-animal" model of control. Control theory is a model of a phenomenon. We (I) like the model because it explains the phenomenon of control, a phenomenon which has been completely neglected by psychologists. What phenomenon of control does a proposed "multi-animal" control system explain?

Tom Bourbon: Just a brief note to echo the remarks of Bruce Nevin, Rick Marken, and, especially, Bill Powers, concerning alleged "insertion" of reference signals into individuals by "social systems."

Bill was on the mark when he said the only way that could occur would be for the individuals who comprise the "system" to bore a hole into the skull of the unwitting controlee and somehow manage to stimulate all of the proper channels which could eventuate in a perceptual reference signal, which is, after all, a "request" for a perception. That is all the "brain" ever provides to the pathways which eventually reach muscles—the brain does not send commands to the muscles. In spite of the massive literature to the contrary, there is no convincing evidence that the brain commands anything, so it is a poor analogy for a social system which commands the behavior of individuals. Behavior is not the end result of a linear chain of command, wherever that chain is alleged to begin, whether in a "stimulus," a neural "command," or a social edict.

In a cooperating group, each individual adopts reference signals for his or her own perceptions. Each individual acts on the environment to achieve the perceptions requested in those reference signals. Living systems cooperate (a) when doing so allows each of them to achieve control of perceptions which neither system could control when acting

alone, or (b) when they decide to do so for the sake of doing so—which allows them to control for doing so. In no way does the "cooperating group" put reference signals into the head of any member of the group. All any member experiences are perceptions. Whether the perceptions are even disturbances depends on the reference signals already adopted by the members.

If it were possible for a group to insert reference signals into the heads of others, do you really think control theorists would miss that trick? It would be infinitely more simple than all of this pounding of keys and flapping of tongues we go through!

Bruce Nevin: If A and B on level n + 1 both contribute to the reference signal of C on level n, that reduces the "compulsory" correspondence of either A with C or B with C. On the one hand, that explains the refractory nature of living control systems under coercion (by way of the limiting case, conflict). On the other hand, it corresponds in an interesting way to interpersonal relationships in which people seem to have their reference levels set by other people.

Consider a military hierarchy, or "authority" experiments. One could surmise that there are multiple sources providing input to the reference signal for certain high-level control systems concerned with interpersonal relations, governing who is judged credible, whose midlevel requests or commands for action are taken as setting reference signals for action, etc. Not all of these persons and other contributors need be physically present to provide that input; most of them are in fact present in memory and imagination. But all that is physically present are environmental events interpreted by the subordinates as intensities, sensations, transitions, configurations, etc. All of them probably evoke memories and initiate imaginative processes, in particular the processes we experience as understanding gesture, language, etc. Having worked its way up the hierarchy to a fairly high level, this input might contribute to the reference signal for control systems governing other hierarchies down to effectors and action. There is no direct input of signals from control systems in one person to the reference-signal "wires" of control systems in the other. The entire depth of the control hierarchy literally stands between, and it does so in each person.

Bill's familiar box diagram helps me to see the disanalogy. We might suppose that a single control loop implemented in neurons is internally complex, but we ignore all but the inputs and outputs identified in the box diagram. Why not have a single box diagram for a living (human) control system? Then we could think about social-level control systems. In such a diagram for one human, there would not only be innumerably many more inputs and outputs at the extremities (sensors and effectors at the bottom, supposed interpersonal reference-signal inputs at the top), but also the "comparator" box in the middle would be enormously complex—the whole intervening control hierarchy, in fact. This makes clear to me that we are engaged in an error of logical type when we do this.

Bill Powers: Hurray, Bruce, you get my point.

Bruce Nevin: I got your point a few months ago. Influence of one hierarchy on another seems to be horizontal rather than vertical, and on any corresponding level of the two hierarchies. Cf. my suggestions about the necessarily indirect character of divine intervention: If you're God or a messenger thereof, how do you influence a world full of autonomous hierarchical control systems (free will)? By suggestion. Cf. also how hypnosis works. All of which is precisely not germane to understanding hierarchical control, only to understanding relations among hierarchical control systems. I'm not putting this down here for the sake of provocation, but only to contextualize: I have not needed persuading that social "control" is not hierarchical and that it is necessarily illusory.

Kent McClelland: As a sociologist, I've been quite interested in the thread on social control. Overall, I find the position outlined by Bill Powers, Rick Marken, and Bruce Nevin to be generally persuasive, particularly their skepticism about the existence of social "control systems" which operate in the same way as the control systems in an individual. On the other hand, my sociological training gives me some sympathy for the opposing point of view, the notion that social conventions have a "reality" external to individuals.

In spite of my general agreement, I wonder if Bill is perhaps stating his case too strongly when he almost makes it sound as if the social environment allows people unlimited degrees of freedom to do whatever they please in any situation. No doubt, the highway patrolman who pulls you over might decide to have a nice chat with you about the weather, or decide to beat you half to death with his nightstick, or do anything in between, entirely as the spirit moves him. Nevertheless, I feel quite confident in predicting that no highway patrolman will ever pull you over to give you a big kiss on the cheek. Some things, I argue, are truly out of bounds in given situations.

In other words, while the social environment is surely not one big negative-feedback system, some or all of the people who constitute a person's social environment cooperate to impose organized disturbances which then place limits on the range of reference values the person can bring under control. This social constraint happens in

much the same way as, to quote Bill, "The physical environment also introduces constraints, but not purposive constraints." Kissing cops are nearly as improbable as pigs with wings. One important difference, however, between the social and physical environments is that some social constraints *are* purposively imposed by at least some of the people participating in the social environment. This constituent purposefulness tends to make the social environment more complicated to describe than the physical, and it could also be the source of our illusion of social control.

Bill exhorts people interested in these issues to work on devising plausible models. With that goal in mind, I have a modeling question. First let me sketch in some background. We know from numerous tracking demonstrations by Bill, Rick, Tom, and others, that the actions of a complex hierarchically organized set of control systems (a human being) can be modeled with great accuracy as a single control system when the task is as simple as keeping a cursor in line with a target on the computer screen.

One of the demonstrations Tom set up at Durango allowed two people to work together on the same tracking task, and my impression from that demonstration was that the joint actions of the two people could also be modeled with great accuracy as a single control system, at least as long as the two people were in agreement on the reference level for the task. To an outside observer, the movement of the cursor on the screen seemed about the same, only a little more precise (higher gain?) when two people were working together on the task than when one was working alone. When I was one of the people involved, the task seemed not to change, just get a little easier.

My question is this: Under what conditions can two (or more) independent control systems, working in parallel in the same environment, be modeled as a single system? How much discrepancy in reference levels, disturbances, system gains, speed of response, and the like are possible before the outside observer would need to posit two (or more) control systems at work instead of one in order to model their joint behavior? How would you devise a test for whether two independent simultaneously operating control systems had the same or different reference levels?

I suspect that if we could specify the conditions under which independent control systems can "cooperate" to produce behavior indistinguishable from one "super" control system, we would make a start toward resolving the "social control" issues.

Am I on the right track?

Bill Powers: Kent, in Colorado it's apparently possible to be pulled over and be given some sort of good driving citation. I don't know if you get an actual kiss.

Maybe it would be useful to distinguish between "social control" and "concerted control." When 20 people decide that an ocean-going lifeboat should be launched off the beach, each person adopts the reference signal "boat in water," grabs the boat, and drags it into the water. Of course, if one person alone tried that, the boat wouldn't move. If 10 of the people adopted the goal "boat 50 feet further from the water," the boat wouldn't move, either. Concerted control is something like distributed processing. If the goals are aligned and the perceptions commensurable, you get the effect of a single control system with much greater output than any one system alone has (and higher loop gain).

A related kind of control is "coordinated control." Now all 20 people together are unable to move the boat using a steady pull. However, if one (any one) of these people says "Heave! Heave! Heave!" (and the other 19 understand what this means and agree), surges of total force can be generated which are greater than the maximum possible sustained force, and the boat moves in steps. Each person agrees to synchronize the pulls with the voice signal, thereby giving the signal the status of a command. To a bystander, it might appear that all 20 people have, suddenly turned into stimulus-response systems, with one of them mysteriously providing autostimulation.

Then I suppose you could have "managed control." The skinny captain of the lifesaving team watches the struggles on the beach for a few minutes, then claps his hands and shouts "Give 'er a yo-heave-ho!" The team, shamefacedly, agrees and starts singing "Yo, heave Ho!" and the boat starts to move as they pull in time with the song.

In all of these cases, the actual control lies inside individuals and is conditional on agreement and understanding.

I think just by remembering that control always lies in the individual, one can come to understand social phenomena without invoking some superordinate being or mystical force, much as Clark McPhail and Chuck Tucker avoid such things in their analysis of gatherings. When I say that there is no social system, I'm denying the widespread sense that there is an impersonal system run by some gigantic and implacable (and rather stupid) monolithic entity analogous to a single human being. Of course there is a social system: it is not, however, a unitary control system, but the outcome of all of the concerned individuals interacting, cooperating, conflicting, joining together in concerted effort, seeking each other, hiding from each other, looking for dependence and independence, enforcing laws and fighting or ignoring them, and so on. Small groupings of people in this system occasionally and for short periods get their goals and perceptions to run sufficiently in parallel to accomplish something together which they

could not accomplish alone. The rest of the time these same people interact differently with each other, often against each other. The net result, at any given time, can be any sort of system which is imaginable, including no-feedback and positive-feedback systems. The result can imitate a hierarchy, a heterarchy, a random network, or simply randomness. The only things determining what kind of system it is, aside from physical constraints, are the perception and goal structures of the individual persons which are in effect at the time.

I think that social laws can be deduced, but they will not be fixed universals. They will be contingent: *If* a group of people adopts such-and-such a mix of goals and has such-and-such skills, *then* the following phenomena of interaction will emerge. One example of this sort of law is the degrees-of-freedom concept. When there are enough people sharing a given environment that the number of independent goals possible exceeds the available degrees of freedom in the means of achieving those goals, conflict (and all its symptoms such as aggression and violence) will necessarily appear. The growth of social systems can probably be traced to the various feasible means which exist for resolving such conflicts: taking turns, specializing, developing the idea of concerted, coordinated, or managed control, and so on. Each person in a conflicted society has a personal motive for adopting methods which will reduce conflict: the restoration of personal control.

The real question is not whether there is a "social system." It is what kind of system it is at the moment and in a particular locale. I think that the answer varies with place, personnel, and circumstances. There is always a system, even in the inanimate world. The whole universe is a network of interacting variables, which is all you need to have a system.

Bruce Nevin: Kent, in the cases of interest to sociology and social psychology, it seems to me the shared reference values concern controlled perception of roles and relations and moves in a conventionalized, game-like sequence involving them. It's not "look, we can both keep this cursor on track," but rather °It's your turn to track that one now, and I know you know that, and I know you know I know it, etc., by prior agreement." Of course the notion "your turn to track" is "unrealistic," but only in the sense that any model can seem very simple and artificial by comparison with that modeled, and modeling the control of roles and relations required by convention to carry out tracking tasks might not be a bad next step.

I have no difficulty with (and argue for) "the notion that social conventions have a 'reality' external to individuals." I only argue against the supposition that this social reality reflects suprapersonal hierarchical control. Though control theory has enormous scope, it necessarily does not encompass all that is to be said about human and animal behavior. Necessarily? Relations among control systems, precisely because they are not matters of hierarchical control, are by definition not treated in it beyond the observation that our familiar presumptions about interpersonal control and power are wrong, and the beginnings of evidence that some patterns in social behavior are mere byproducts of individual control for values conceived as private, rather than social (arcs and rings in the crowd program). Will the latter suffice? Lots of muck shoveling, perhaps, before we get at claims of the social sciences which bear deeper scrutiny.

You ask: "Under what conditions can two (or more) independent control systems, working in parallel in the same environment, be modeled as a single system?" When are two (or more) autonomous control systems controlling for the modeling of themselves as members of a single system, according to mutually known roles and relations?

Bill Powers: Bruce says that "in the cases of interest to sociology and social psychology, it seems to me, the shared reference values concern controlled perception of roles and relations and moves in a conventionalized, game-like sequence involving them." The "game-like sequence," once adopted by one person, consists of rules like: "If he does or says A, I do or say B." As the other person, you can learn to perceive this rule experimentally. Of course, before that can happen, you have to perceive a principle: "Hey, that was a funny thing for this person to say (or do) in relation to what I said (or did)—ah, he's playing some sort of game. Let's see if I can figure out the rule."

Bruce also says, "I have no difficulty with (and argue for) 'the notion that social conventions have a "reality" external to individuals.' I only argue against the supposition that this social reality reflects suprapersonal hierarchical control." The relevant "reality external to individuals" is, of course, other people. Other people do things for their own reasons. They seem to march to inner drummers, and often a lot of them seem to be marching to the same inner drummer (as near as you can figure). They also build things and leave them around: chairs, houses, roads, television sets, dinners. Those things are just physical objects until you realize that someone had a purpose in building them, and figure out what that purpose might be, and try it out for yourself. Then you know what it feels like to march in cadence.

When you see enough people apparently reacting to you in accord with a rule of some game, and when you have deduced the rule well enough to predict how they will respond to your moves (or disturb you if you don't move), you might come to think, "OK, I guess that's the rule," and adopt it for yourself. This can leave you with the impression that this rule exists somewhere out there in space. It seems

to affect everyone, so it must be imposed from somewhere else. It isn't just that your mother likes to put the fork on the left with the napkin, then the plate, then the knife and spoon on the right. That's they way they are *supposed* to be placed. It's a rule of etiquette, and etiquette isn't something people decide to do: they do it because it's right. Now the rule has become reified; it no longer seems that you or anyone else has a choice.

Our language is full of words which have the specific function of making social rules seem to be something other than a personal choice adopted after considerable effort. You have duties, responsibilities. You must do what is right. People have something called "authority," and it must be "respected." Children must learn to "cooperate" (i.e., do what they are told). This is a government of laws, not persons. People can be upright or transgressors. People have "rights."

When you start thinking about all the facets of society (as it is or as it should be) to which you wholeheartedly subscribe, you come face-to-face with the real price of understanding control theory. The sense of being carried along and protected by some benign regulating system external to yourself disappears: you are faced with taking responsibility for fundamental aspects of your life which, long ago, you turned over to someone else. You see other people not as being in the grip of the system, but as the authors of their own choices and their own actions.

In truth this basic freedom was there all along, but in getting involved with figuring out all of the games which are going on, and in learning how to adopt the rules yourself and use them for your own ends, you, the adult, have forgotten what the point was. It's both liberating and frightening to realize just how much of your life is in your own hands.

A true model of behavior doesn't just describe the way people are. The way they are results from just one possible adjustment of the model, one possible set of parameters. A true model shows you other ways they might be, given changes in the parameters and in the alterable aspects of organization. One reason why control theory has taken so long to be recognized and adopted is that the older theory wasn't even recognized as a theory—it was simply the way things are. Something happens, and a person responds to it. That's just a fact. But when you realize that reference signals are adjustable, that stimuli are really disturbances of controlled variables, an apparent response to a stimulus suddenly becomes just one of the possible outcomes. If the reference signal changed or the perceptual function were reorganized, the same stimulus would be followed by a different response, or the opposite response, or no response at all.

Societies as they are now represent the outcome of one way human organisms can conceive of each other and interact with each other.

Control theory shows that there are other ways. The job of control theory is not just to describe social phenomena as they are, but to reveal those phenomena as a consequence of adopting just one mode of organization out of many which are possible. The same goes for language: language as we know it is just one way in which people can use conventions, rules, and principles to manage their interactions with each other. To understand language, we have to see how the system might be different from the way it is—merely fitting a control-theory interpretation to the situation as it exists is only a small first step.

Bruce Nevin: Bill, two angles on the "relevant 'reality external to individuals': the cellular consciousness angle and the furniture angle.

Cellular consciousness first. This is a point-of-view problem. The relevant reality (in the same sense) external to the cells in my body is the other cells. They are governed by and in part constitute a hierarchical control system, per theory and experiment so far. This is a thing of a radically different order from the cells and other structures in the body, and the cells, so far as we know, lack means of detecting or controlling for this higher-order reality. There is no convincing evidence that people together constitute social hierarchical control systems in an analogous way, and fundamental reasons (no way to implant reference signals, conflict instead of compulsory compliance) why, so far as we now can tell, they cannot.

We nonetheless seem to want to push this analogy and through reconstructable human history always have. The king is likened to the head in medieval society, the priest to the heart, the serf to the hands, and so on. Metaphors abound for finding one's place in social space. Are we just inventing to fill a need for top-level reference values?

Conversely, our ineptitude at this business of explaining a social level of organization to ourselves does not indicate that there is none. Assume that there is some higher level of organization of some sort perhaps inconceivable to us, as the organization of my body would be hypothetically inconceivable to my pancreas (were it capable of having conceptions of things). Just assume that there is, for the sake of the argument. For the next two lines of text. Statements like "The relevant 'reality external to individuals' is, of course, other people" are in that case clearly reductionist. (OK, you can stop assuming a higher level now. The pain will go away if you rub it.)

This is a point-of-view problem because it is not clear that anything can have a point of view in any usefully relevant sense if it is not a hierarchical control system. But what do we know?

Now the furniture angle. A person walks into a room and tries to sit down in a handy chair. "Don't sit there!" Shocked expressions. Only a person in a certain role can sit there. That person might arrive any

moment. The status of the chair, the role, the person holding the role, the visitor, etc.—these exist only by virtue of the human participants maintaining certain reference values internally. If they ceased to live, or ceased to maintain those reference values for whatever reason, those social realities would be no more. If the cop ceased to hold certain reference values, he would be a cop no more and might indeed kiss you on the cheek. But that does not demonstrate that things like roles, statuses, etc. are unreal. The existence of many things which are undeniably real is contingent upon hierarchical control, notably our own existence as living, conscious beings. Do you deny your existence as a person because that existence is contingent upon reference values held by elementary control systems at various levels of your control hierarchy?

The furniture of our lives is all social constructs. The fact that we do the constructing out of our perceptions of culture-free objects and events is no more relevant than the fact that the objects around us are "really" mostly empty space, or the facts of quantum mechanics.

And, indeed, the objects and events exist for us only as control-system constructs. It is not only social reality which is contingent. (Is that an orange flower? More energy in the UV range. And get a load of that gamma burst across the parking lot!)

The main concern on your part seems to be autonomy, rather than ontology. There certainly are a lot of rules, but is that all there is in the social realm? All if-then program steps and nothing else? No, clearly, that somewhat ill-defined range of levels between programs and configurations/transitions has culture-specific elements in it, all of the business of words and symbols and signs. These are not rules, nor are they likely to be constituted as they are only by virtue of rules stipulating how one is to interact with them.

"Societies as they are now represent the outcome of one way human organisms can conceive of each other and interact with each other. Control theory shows that there are other ways." On the one hand, anthropology shows that there are many ways. On the other hand, people do need to coordinate their goals without expending all their efforts on arranging to do so, and if control theory suggests a better way than learned social conventions, we should hear about it.

"When you start thinking about all the facets of society (as it is or as it should be) to which you wholeheartedly subscribe, you come face-to-face with the real price of understanding control theory. The sense of being carried along and protected by some benign regulating system external to yourself disappears: you are faced with taking responsibility for fundamental aspects of your life which, long ago, you turned over to someone else. You see other people not as being in the grip of the system, but as the authors of their own choices and their own actions."

Worth repeating, so I did. The same experience arises when one becomes multicultural, multilingual, multidialectical, able to shift adaptively to the prevailing norms. To a slight degree, we all do this. But the sense of those norms only appears when there is conflict with them, otherwise it is invisible, so to say it disappears with the epiphany of control theory seems to miss the mark. Rather, we offer ourselves different sorts of choices when conflicts about coordinated control do arise. A different way of saying the same, I think. Different means, different experiences, can lead to the same shift in how one experiences.

Chuck Tucker: It seems to me that several comments I made earlier this year on the Net are relevant to the issue of "external control." I stated the following:

Society, social structure, social class, culture, and group pressure do *not* make people do anything.

Personality, socialization, and social background do *not* make people do anything. (Rather, these provide resources for action, but determine none of it.)

Social life, by which I mean living and acting together, depends on arrangements people make.

People guide their actions by directions they give themselves.

Discovering the laws of social life is *not* possible, or even sensible.

Biological agents such as germs and viruses, or chemical agents such as alcohol, cocaine, and steroids do not make people do anything. (Rather, these can affect performance levels and the coordination and control of behavior.)

Technology does *not* make people do anything. (Rather, technology provides resources for action.)

Social norms, rules, values, beliefs, customs, traditions, laws, and social sanctions do *not* make people do anything. (Rather, these are devices people use to facilitate living and acting together.)

Genetic inheritance and other biological factors do *not* make people do anything. (Rather, these permit people to do what they do, and, undoubtedly, permit them to do much that so far they have not done.)

Without making arrangements, people are socially incompetent.

People *cannot* be made to do anything, unless they are literally and directly and physically forced to.

These comments speak quite directly to matters of "external control," but I don't believe that those of us who use the word "control" in the title of the model we use will ever be able to avoid the problem of interpreting that word as meaning "control by others," "force," "manipulation," "external influence," or "environmental cause" unless we clearly point out, as Bill has, that control is a technical term meaning stabilization of a variable against arbitrary disturbances. Most people

do not use this definition of control, and sometimes some of us forget and use it in a non-technical sense. All of the literature in sociology, psychology, and social psychology which I have examined uses the term to mean either control by outside forces or forces responding to the violating of norms, rules, or laws (this is also the case in my dictionary, where I find that "control" comes from the Latin *contra*, meaning "against"). So, to have others understand what we are talking about and are interested in, we will have to preface our remarks with the technical definition of control or make up other words for social control, like "influence" or "reciprocal influencing." Another alternative is always to use the phrase "perceptual control," and clearly distinguish it from "social control" and "reciprocal influencing," but never to use the word "control" alone. Of course, each of us can give that direction or instruction to him/herself, but following it is always a difficulty.

It seems to me that the recent posts of McClelland, Powers, and Nevin should be read as a set with the focus on how control theory deals with the "social." I see a wide area of agreement that language is crucial because it is *used* by people for their perceptions, to adjust reference conditions, and even to adjust loop gains, as well as being crucial in the reorganization process. Roles provide a handy illustration of how this is done, since a person will evoke a role not only to "control" his/her own action, but will ask another to "control" to do similarly, as in "I'm your Father," "Don't call your Mother "her"—she is your Mother," "This is Dr. Tucker speaking," "I said that when I was a member of the administration, but now I'm a Judge," "He's not Bush, he's President Bush," and "I'm transferring you a call from the President." Now, not everyone will act exactly the same when such statements are made, but my bet is that one would observe very similar actions from the receiver of such statements. We are not robots, but we can organize our conduct in ways which are quite predictable to ourselves and some others.

Rick Marken: Kent, it might be that a tracking task with two people can be modeled with one control system. But why do it if there are really two? As I recall, in one of Tom's demonstrations, one person controls one cursor and another person controls another, possibly relative to each other, but not necessarily. The social part comes from the fact that each person affects their own as well as the other person's input. I don't see how this particular task can be modeled with a single control system; there are two degrees of freedom to be controlled (the two cursor positions), implying two control systems.

Many of Tom's demos show that two control systems can act cooperatively even if that is not their goal. This is what happens in the case above. You could also have two people control the difference between two cursors—now you could run into conflict if there is a difference in the reference for what this difference should be. If you set it up so that both systems are affecting the input variable in the same way, then you probably could model this situation with a single control system, and the accuracy of the model's match with behavior would depend on the closeness of the two references, the relative gains of the systems and all of the other stuff you mentioned. I don't see why one would do this, especially when you know that there are two physically different systems working on the task, and you know how they are connected to the input variables. Are you thinking that social control has something to do with the degree to which two actually separate control systems act as though they were one? If so, then your notion of social control differs from mine. I think of social control as something which controls the interactions between two or more people. An interaction is behavior (actions) on the part of two or more people who influence variables which are controlled by one, both, or all of the people. A social controller would be something external to the people involved, which controls this interaction in some way.

I think what Tom is trying to show (rather beautifully, I think) is that "interactive" control requires no external social controller. The appearance of social control (as I use the term) emerges out of the non-conflicted interaction of multiple control systems. Another nice illustration of the "emergent" nature of social control is the crowd demo of Powers, McPhail, and Tucker. Here, complex, coordinated social behaviors emerge out of the mutual interaction of many control systems.

I think the best way to get at this "social control" issue is to define precisely what it is. Perhaps we could agree on one of Tom's demos as a prototype example of social control, and then see what's actually going on—and whether there is any evidence that there is more going on than interaction between two or more control systems each controlling their own inputs (and, in doing so, adjusting to the effects of other control systems).

If it turns out that there really is no such thing as "social control" as conceptualized by sociologists (and other social scientists), this does not mean the end of sociology—not by a long shot. Control systems *do* interact; they are social. So this is what sociologists will study—the phenomena which result from the interactions of multiple control systems.

So don't worry, Kent, there is still plenty (possibly even more) to do in a control-theory-based sociology.

The Control Systems Group is a membership organization which supports the understanding of cybernetic control systems in organisms and their environments: *living control systems.* Academicians, clinicians, and other professionals in several disciplines, including biology, psychology, social work, economics, education, engineering, and philosophy, are members of the Group. Annual meetings have been held since 1985. CSG publications include a newsletter and a series of books, as well as this journal. The CSG Business Office is located at 73 Ridge Rd., CR 510, Durango, CO 81301; the phone number is (303)247-7986

The CSG logo shows the generic structure of cybernetic control systems. A Comparator (C) computes the difference between a reference signal (represented by the arrow coming from above) and the output signal from Sensory (S) computation. The resulting difference signal is the input to the Gain generator (G). Disturbances (represented by the black box) alter the Gain generator output on the way to Sensory computation, where the negative-feedback km/ is closed.

Now available:

LIVING CONTROL SYSTEMS II

Selected Papers of William T. Powers

ISBN 0-9624154-2-1, 1989, 277 pages, illustrated, softcover, 5.5” x 8”, $22.00 postpaid (KY residents please add sales tax).

There are 22 previously unpublished papers in this book, including two chapters from the preliminary version of Behavior: The Control of Perception (“Emotion” and “An Experiment with Levels”) and the complete version of “CT Psychology and Social Organizations.” W. Thomas Bourbon wrote the Foreword, and there is also a list of the published works by William T. Powers on living control systems, 1989-1991, which updates the list in the first volume of Living Control Systems.

Order from: CSG Book Publishing, 460 Black Lick Rd., Gravel Switch, KY 40328 U.S.A.

Back cover

Closed Loop

Threads from CSGnet

Spring 1992

Volume 2

Number 2

Front cover

Closed Loop

Threads from CSGNet

Spring 1992 Volume 2 Number 2

Edited by Greg Williams, 460 Black Lick Rd., Gravel Switch, KY 40328

CONTENTS

Members of the Control Systems Group receive *Closed Loop* quarterly. For membership information, contact Ed Ford, 10209 N. 56th St., Scottsdale, AZ 85253; phone (602)991-4860.

CSGNet, the electronic mail network for individuals interested in control theory as applied to living systems, is a lively forum for sharing ideas, asking questions, and learning more about the theory, its implications, and its problems. The "threads" in each issue of *Closed Loop*, stitched together from some of the Net's many ongoing conversations, exemplify the rich interchanges among Netters.

There are no sign-up or connect-time charges for participation on CSGNet. The Bitnet address is "CSG-L@UIUCVMD" (use no quotes in this and the following addresses); "CSG-L@VMD.CSO.UIUC.EDU" is the Internet address. Messages sent to CSGNet via these addresses are forwarded automatically to all participants. Via CompuServe, use the address "> INTERNET :CSG-L@VMD .CSO.UIUC.VMD" to reach the Net. Initially, you should send a note to the network manager, Gary Cziko, at "G-CZIKO@UIUC.EDU" (Internet) or at "c CZIKO@UIUCVMD" (Bitnet); Gary's voice phone number is (217)333-4382.

Epistemology

Bill Powers: The aim of scientific objectivity, it is said, is to remove as far as possible all subjective bias on observations of the real world. In the physical sciences, this is done through the use of instruments, reducing observations to simple judgments of coincidences. But I want to put off that part of the subject and look more closely at the concept of observing without bias. The model of perception that is assumed makes a great deal of difference in the meaning of "observing" and of "bias."

The model I assume is this: the world we experience consists of signals in the brain created by the interaction of the nervous system with the world outside it. This means that neural signals are not *about* the world of experience; they *are* the world of experience. What they are about is another matter that calls for considerable investigation.

If the world we experience exists in the brain, we must then ask what objectivity could possibly mean. I think it means a certain attitude toward experiences, toward perceptions.

If you see a man carrying a briefcase hurrying along under an umbrella through the rain, you can interpret what you see in different ways. You might see a man trying to get to work on time, or someone late for an appointment, or a thief who has just stolen a briefcase and an umbrella. To see these things, you clearly have to add imagined information to what you are actually observing. The same would apply if you saw a man who seemed anxious, or angry, or oblivious to the world. The most objective way of reporting what you see would eliminate all imagined information, all that is not actually in the scene before you.

To be even more objective, you would have to examine the details of what you are seeing. The man seems to be hurrying, but all you are really observing is that he moves more rapidly than others. "Hurrying" is a characterization added to what you see. He seems to be carrying a briefcase, but it could be some other object. "Briefcase" is an interpretation of the shape you see. He seems to be carrying it, but perhaps it is shackled to his wrist. "Carrying" is an interpretation of the relationship between his hand and the object. He seems to be under an umbrella that shields him from the rain, but perhaps the umbrella is a signal to someone he is to meet and isn't being used to keep him dry.

To be most objective of all, you have to ignore all these characterizations, because no matter how you characterize what you see, the characterization always goes beyond the perception; a different characterization is always possible. To be completely objective, one must simply observe without the accompaniment of an internal explanation or characterization of the observation.

This is almost impossible to do. It is possible, however, to broaden the scope of what one thinks of as observation to include not only the scene being observed, but the internal explanations-interpretations-characterizations that come along with it. If one observes both, then it is clear which set of experiences is the interpretation added to the observation. Or at least it becomes more clear.

So the most completely objective observation is that which is totally subjective and silent. It is simply attending to appearances as they actually present themselves, without any attempt to add to them or manipulate them rationally, without saying that they are real or unreal, without theorizing, associating, or explaining. Doing this to the extent that is possible takes practice and discipline leading to a state of mind much like what Zen practitioners seek through meditation.

Now we can reintroduce the subject of instruments, asking on the way why it is that such instruments are used.

The object of scientific explanation is to explain experience. More exactly, it is to explain why some parts of the experienced world are related to other parts as they are. Why is it that when there is a flash of lightning, there is quite often, after a delay, a roll of thunder? All we experience is the sequence of events; any other relationship between them is hidden. Science is an attempt to guess the connections between the flash of light and the sound, to explain the sound as a natural or necessary result of the process that created the flash of light.

Past experience tells us that the world does not appear exactly the same to everyone; furthermore, observations are inevitably tinged by explanation and interpretation, which creep in under the cover of innocuous words like "hurrying." So to eliminate these subjective differences, science employs instruments.

To measure the flash of light, a scientist would use a photoelectric cell, which responds to light by generating a small current that can be indicated on a meter or recorded on magnetic tape. But what does it mean to say that the photoelectric cell "responds to light"? It means that when the photoelectric cell shows a response, a human observer sees light. When we examine the two things being compared here, we see that they are very similar: both the perception and the meter reading are outcomes of receptor processes, one organic and one inorganic. Both outcomes depend on something else, but the human outcome is a brain signal measured in impulses per second, and the light-meter outcome is the angle of deflection of a needle. Neither outcome is in units of "light."

Instruments, therefore, provide us with consistent indications of

something going on on the other side of the instrument, but they don't identify to us what it is that is being measured. Instrument readings are more objective than eyeball observations only because they are more repeatable and are not influenced by interpretation prior to the reading. They are not more objective in the sense of bringing us closer to a pure description of reality itself. The basic correlate of an instrument reading is not some real physical variable, but either another instrument reading or a human perception.

Now what about the claim that instrument readings are more objective because they reduce observations to a simple discrimination of coincidence? The claim would be that the photocell measurement of light intensity is more objective than the visual estimate because the photocell always responds the same way to the same light intensity. All that the human observer has to do is read the meter face carefully, or, for a digital instrument, write down the number on the display.

But what do we have then? Suppose the reading is 12.5678241. Just writing down that number is reminiscent of the joke that goes, "We interrupt this program to bring you a late score: six to nothing." The number by itself is meaningless. At the very least, you have to know that it is the reading from a photocell, not from a thermometer. To use it in relation to any other meter reading, you must know how the meter is calibrated: what are the units of this number? Foot-candles? Lumens? Ergs per second? What is the spectral range being measured? And to use this reading in the context of science, you must also explain what it is that is being measured: the absorbed part of a flux of photons, a flow of energy, a squared amplitude of magnetic and electrical vibrations at a certain frequency or with a certain wavelength. You must, in short, reveal the complete model of what the meter supposedly measures.

There is no way, in fact, to reduce an observation to a coincidence of a meter needle with a mark. If it is reduced that far, it ceases to mean anything.

What gives meaning to the meter reading is exactly the same thing that gives meaning to an uninterpreted human perceptual experience. It is the structure of interpretations and theories that depicts a world on the other side of the receptors. That world does not exist in unadorned, uncommented observation. It exists only in the adornments and comments added by human intellectual processes.

When we try to understand human perception scientifically, we automatically introduce something other than direct experience. We introduce a term "perception," indicating that there is a perceiver, a consequence of perceiving, and something to be perceived. This is like introducing a photocell, a photocell reading, and light-energy. There is an "inside" component and an "outside" component linked by a physical device. All three of these components are theoretical entities, part of a model of processes that underlie direct experience.

The popular conception of a light-meter is that the light is what exists, while the meter reading is only an indication of it. But considering how scientific modeling actually works, the priorities must be reversed. It is the meter reading that is given; what it indicates and how the indication is derived are matters for theory and conjecture. We must reason backward from the meter reading, taking into account the theoretical properties of the meter and the photocell (and doing the same for many other kinds of meters), to deduce what lies at the origin of the reading: reality, the world.

In a model of perception itself, we must do the same thing. The physical device we place between reality and perception is the nervous system. The properties of the nervous system are a matter of theory and interpretation of observations. The given is the perception, the experienced part of the process: the way the world appears. What remains to be deduced, by reasoning backward through the assumed properties of the perceiving device, is the external world.

It is possible to observe objectively only the outcome of this theoretical perceptual process. We can see whether the model consisting of a nervous system with its properties and a physical world with its properties (derived through studies with instruments) can be made to produce an outcome that matches what is in fact experienced. Objectivity then consists in observing what the model actually does, with as little interpretation as possible, and how the world actually appears, also with as little interpretation as possible. But objectivity has nothing to do with reporting on the world that is represented by our models.

Wayne Hershberger: Bill says that "neural signals are not *about* the world of experience; they *are* the world of experience." Almost, but not quite. Any claim that neural signals are *about* the world of experience is undoubtedly false, as you say, precisely because the reverse is the case: my world of experience is *about* my neural signals—and about the optic array in the ambient light, and about the stuff with resting mass which is said to give structure to the ambient light. It is not true that my experiences comprise neural signals. Rather, to quote an authority you might recognize, "this world presents itself... in three dimensions, stereo sound, and living color, chock full from edge to edge of continuously-present smoothly changing noise-free colors, shades, objects, motions, relationships, and operations in progress."

According to a coherence theory of truth, the perceptual objects comprising our experience (phenomenal world) may be said to depend upon (be about) the conceptual objects we "construct" (neurons, photons, electrons, input functions) to the degree and only to the degree that these conceptions (models) account parsimoniously for the

perceptions in question.

Perception is not imagination.

Bill Powers: Wayne, I was not describing what the world *is,* I was describing how it *appears.* This appearance is the world we directly apprehend. The control-theory model, based on the appearances of neurology and physics, leads me to conclude that this world is, physically, a collection of neural signals, although it does not look that way (that is, it does not look the way neural signals do when we visualize them on an oscillogram using electronic means of sampling limited aspects of neural activity).

I know that you insist that there is no other world than the world of appearances (or that we don't need to consider one in explaining perception). When we aren't concerned with explaining, but only with experiencing and living (no theoreticians required), I agree with you. Reality is precisely what we experience. But to add "and nothing more" is to assert what we can't know directly, and to deny, for no good reason, the implications of physics, chemistry, and neurology, all of which claim to represent a world of immense detail that is inaccessible to our senses.

The crucial difference between our views, as nearly as I can comprehend it, is in my assumption of another reality that is not part of the world of direct experience. If there is no such separate reality, then of course all that exists in nature is experience as it appears to us. Neural signals, physics, chemistry, and all such conceptions are just that: conceptions and nothing more. They are simply ways of ordering our experiences and have no significance beyond that. They do not refer to anything unexperienced.

I can understand that conclusion as a conclusion, but I can't accept it as a fact. It is simply another conjecture. Its truth, then, comes down to the evidence we have for and against it, and to how we reason about that evidence.

I think that control phenomena provide us with evidence that there is a universe beyond the limits of human perception; that this universe imposes its properties between our actions and their perceived results; that there are independent agencies in this universe that are capable of disturbing our control actions without our being able to detect the causes of the disturbances. I think we learn about these properties only indirectly, and as conditioned by the kinds of perceptual systems we have and do not have. I think we have to infer the nature of the disturbing agencies and the properties of the world, by building models that would, if they were true representations of the unseen world, explain how our experiences are related to each other. I do not think it is likely that we have arrived at models that just happen to capture every significant entity outside us, every significant functional relationship among those entities. And I do not think it is likely that the world of direct experience exhausts the degrees of freedom that really exist in the universe around us.

So that is my basis for accepting, as the most reasonable hypothesis, the existence of a real universe apart from our perception of it, and for denying, on the basis of the same reasoning, that our perceptions are likely to be veridical renditions of that universe.

I agree that perception is not imagination. Imagination is, however, a subset of perception. Some of our experiences are generated inside the brain and do not depend on the current external state of affairs, even though they might sometimes give a convincing imitation. But the rest do depend on something outside. In neither case, however, do perceptions without the aid of reason give us a picture of what is really causing them—however inadequate the picture.

Wayne Hershberger: Bill, it would be a comfort were you to agree with me, but I do not agree with the statement, "Reality is precisely what we experience." Of course, I agree even less with the obverse idea.

That is, as an empiricist, I do not endorse the idea that reality is precisely what *cannot* be accessed empirically. Hence, I cannot imagine ever saying that physical, chemical, or neurological phenomena are epistemically inaccessible. In fact, I have been championing the antithesis.

I guess I am not making myself clear.

Let's go to square one. I claim that the epistemological challenge is not to explain how the truly inaccessible can be accessed (a logical impossibility), but rather to explain how the truly accessible can have appeared to be inaccessible (a logical possibility). Do you agree?

Bill Powers: Square one it is. I agree that the truly inaccessible can't be accessed. That's a definition, not a proposition. The truly inaccessible is that to which we have no access at all, either direct or indirect.

The second part of your claim is not a definition: it asserts that the accessible sometimes appears inaccessible. I think this is an attempt to create a two-valued situation (either something is accessible or it is not) out of one that has more than two possible values.

I see accessibility of the workings of nature to be a matter of degree, with the maximum degree falling short of 100%. To explain my view, I will resort to a thought-experiment.

Suppose we have before us a black box which we have no means of opening. Let's call this box, to humor me, the Reality Box. On its surface are numerous buttons and lights. The buttons and the lights are undoubtedly real, because we experience them directly and unequivocally: they are totally accessible and cannot be mistaken as being inaccessible to our observation. So the box, the buttons, and the lights are not an epistemological problem.

When we press various buttons, we find that certain lights and combinations of lights turn on or off. With sufficient experimentation and record-keeping, we can discover consistent effects of the buttons on the lights. As our experience grows, we can discover that some buttons alter the effects of other buttons on lights, or make certain lights come to depend for their state on new combinations of buttons. We can discover that only certain sequences of button-pressing will have predictable effects on one or more lights. We can find that certain lights have mutually exclusive states; if a member of one set is on, another set is always off, and vice versa. We could uncover logical relationships, relationships corresponding to arithmetic operations, and so on. We could even develop heuristics: some ways of turning lights on work best under one set of circumstances (combinations of lit lights), usually, than other ways.

Thus, we arrive eventually at a sophisticated empirical understanding of the Reality Box. At no point have we asked what is inside the box. We have simply observed, recorded, and tried to recognize consistencies.

Perhaps I should motivate this investigation by saying that for reasons we only vaguely understand, certain of the lights on the Reality Box have extraordinary value to us; indeed, their states of illumination seem to us to be a matter of life and death, or at least make the difference between enjoyment and disappointment. So we have an interest in pushing the buttons to keep the most important lights in the states that seem the most desirable, especially as they will not stay in those states without the button-pushing. To be crass and less mysterious, we could say that each time we succeed in maintaining the critical lights in the critical state for one minute, we receive $5—that's $300 per hour for this job if we can learn to do it perfectly.

I think I have now described the state of human understanding of nature in the pre-Galilean era.

In fact, we find that we are a long way from making $300 per hour—the actual payoff isn't nil, but it's just barely a minimum wage. It would be greatly to our advantage if we were allowed to cheat: to open the box and trace out the circuits (or talk to the little men, or analyze the chemicals, or take whatever action is appropriate to what we find in the box). If we knew *why* the buttons affect the lights as they do, we could abandon the trial-and-error empirical approach and simply deduce the actions that would have the effect we want.

I am now describing the advent of the physical sciences.

We are not, however, allowed to cheat. Nobody knows how to open the box.

Nevertheless, once we get the idea of explaining the dependence of lights on buttons rather than just observing it, we might well decide that even a good guess about what is in the box might be more valuable than random experimentation. So we begin to construct a model of the internal workings of the Reality Box, trying to outguess its designer.

This project turns out to be extraordinarily successful. By imagining circuits and functional devices inside the box, we succeed immediately in explaining why some buttons cause some lights to change their brightnesses. Numerous revisions of the model are required, however, because just when we think we have the right connections, an anomaly turns up and we have to modify the design of the hidden devices or the connectivity between them. But by demanding that the model *always* work, no matter what combinations of buttons we press, we eventually get this model to the point where it never fails in any way we can notice.

We now begin to believe in the reality of the model. What appears to be happening is that buttons activate lights, but what is *really* happening, we say, is that the buttons are feeding their effects into a hidden complex device that in turn operates the lights. Gradually, the status of the insides of the Reality Box changes. Those insides no longer seem hidden to us. In fact, even though they are complex, they are far simpler than our records of empirical findings are. They also permit us to predict the effects of button-pushings, even combinations never tried before, with exceeding accuracy, whereas our empirical predictions, based only on unexplained frequencies of occurrence, are wrong nearly as often as they are right, and are essentially useless in unfamiliar circumstances.

I have now described the rise and maturation of the physical sciences, and their essential difference from the purely empirical sciences.

The penultimate stage in this development arises when someone notices a fact that by now is considered a very strange fact. Those who are engaged in the exploration of the Reality Box by now feel that its devices and connections are *in the box.* The lights tell them what is happening inside the box; the buttons let them influence what is happening inside the box. What this someone says that is thought so strange is merely a reminder that, in the beginning, nobody knew what was in the box, because only the lights could actually be observed, and no effect of pressing the buttons could be seen except in the lights. This is, in fact, still the case. So the model of what is in the box must exist in the minds of those who are observing the box. It is not in the box. In fact, it is perfectly possible that what is in the box is entirely different from what is in the mental model, but is equivalent to what is in the mental model under all of the button-pushing operations so far tried. Even

what seems to be a simple direct connection through a hidden wire might actually involve a hidden modulator that converts the button-press to a radio frequency and broadcasts to a receiver whose output lights the light. That would not be a parsimonious design, of course, but it might be the one that exists.

What this upsetting stranger is doing is reminding everyone that all they can actually observe are the lights, and the only effect they can know they are having is to press the buttons. All the rest is imaginary. Therefore, we should throw away all those figments of the imagination and admit that all we know is how the buttons affect the lights, and to remain pure of heart, we should talk about nothing else.

I have now described the advent of stimulus-response theory, behaviorism, biology, empirical psychology, and so on.

The final stage entails the epistemology of the Reality Box that I propose.

In fact, the model works much better than it should. Moreover, there is evidence in the relationship between buttons and lights that tells us something consistent is going on independently in the Reality Box. The lights that we can affect with our buttons sometimes turn on and off when we aren't pushing anything. Very often, we have to change which buttons we push in order to reproduce the same state of the lights, and there seems to be no way to predict when, by how much, or in what direction we will have to make these changes. Something else is interfering with the effects that the buttons have. This something else can be inferred, to some extent, because it might occur regularly, or in some regular pattern, as we can tell by watching what different buttons we have to press to reproduce the same effect, and when we have to do this.

So we are led, in the end, to recognize three major facts. First, our mental models of what is in the Reality Box have an unknown relationship to what is actually there. Second, the regularities implied by the model actually do occur, even through we can't know that they occur for the reasons we propose. And third, there is something in the Reality Box that can act independently of us. So we can say that in some regards, what is in the box is accessible to us, but we must also admit that our interpretation of its inner workings is not necessarily the only one that would be as good at explaining what happens.

The lights, of course, are our perceptions, and the buttons are our actions. The Reality Box itself is invisible; we experience only the input and output devices mounted on its surface. We conjecture that the buttons do something that we don't observe. We conjecture that the lights indicate something that is also not observed—if only the presence of a wire from the button to the light, and an invisible power supply.

So what does "accessible" mean? Does it mean that we observe

Reality exactly as it is, or that there is neither agency nor order other than what is evident to us in our sensory experiences? I feel that such questions are not matters to be deduced logically, so that we can know once and for all the truth about experience. I think that they are matters to be settled as we settle all factual questions: by the examination of evidence, and by settling for the inference from the evidence that seems most supportable by all the rest of our experiences and knowledge. Pure philosophy can't provide that sort of conclusion: it demands an end-point, a certainty. That, I think, is definitely inaccessible.

Wayne Hershberger: Bill, while agreeing that the truly inaccessible is inaccessible, you assert that much of that which is accessible is not all that accessible. Here you are changing the subject (I'm assuming that you are not simply contradicting yourself). That is, you are using the word accessible to refer not to the possibility of epistemic access, but to the amount of X that is accessible, or to the difficulty of achieving access to X, which could vary, of course, with the directness or complexity of the epistemic process. This is as misleading as referring to the length of a pregnancy as a degree of pregnancy (e.g., the unwed mother who claims she is "just a little pregnant"). More specifically, if 90% of X is truly accessible and 10% is truly inaccessible, it does not follow that X (i.e., every bit of it) is 90% accessible.

You say: "... the box, the buttons, and the lights are not an epistemological problem." On the contrary, the epistemological problem is usually stated in just such terms; for instance, it might be said that although the buttons "appear" to be solid and stationary, they "really" comprise a swarm of whirling dervishes known as atoms. Or alternatively, it might be said, as you are wont to do, that they really comprise a collection of neural signals. In fact, the buttons, lights, etc. are the phenomenal objects that the empirical process we call perception provides us, and the epistemological question concerns whether or not these phenomenal objects are as objective as the label "object" implies. The box, the buttons, and the lights pose the epistemological problem! Your elaborate example begs the question.

To suppose that the phenomenal objects are accessible appearances comprising indirect representations of an inaccessible reality is to embrace a radical skepticism, because there is no way to assess the fidelity of the representation without accessing the inaccessible (i.e., in order to test the correspondence between the reality and the appearance, one needs access to both, and that, by definition, is not possible). This question of correspondence between what is accessible and what is inaccessible (i.e., between what is internal to and what is external to the limits of experience) is readily confused with the correspondence between what is internal to and what is external to the nervous system. But

whereas the former type of correspondence is impossible (by definition), the latter type of correspondence is easily determined—neurophysiologists do it all the time.

You say, "What this upsetting stranger is doing is reminding everyone that all they can actually observe are the lights, and the only effect they can know they are having is to press the buttons. All the rest is imaginary. Therefore we should throw away all those figments of the imagination, and admit that all we know is how the buttons affect the lights, and to remain pure of heart, we should talk about nothing else." If I get your meaning, that stranger is no stranger. His name is George Berkeley. Later, Johannes Muller echoed Berkeley's Subjective Idealism in his doctrine of specific nerve energies; as Muller put it (in his article 5), the sensorium is aware not of the external object, but of the state of the nerves only. Having said that, Muller then seemed to recognize belatedly that Berkeley's thesis makes no sense expressed in physiological terms, because in article 8 he said that the sensorium is aware not merely of the state of the nerves, but of the external causes as well. (There's nothing like having your cake and eating it too.) Muller was confusing the two correspondence questions described above. Berkeley's philosophy concerns only the former type, as Kant observed; that is, the only objects we experience are phenomenal, not noumenal, things.

I share your concern with the essential nature of nature. That is, when an experiment asks a question of nature, "someone" answers. But I see no reason to exclude this final arbiter of empirical truth from the phenomenal domain. Banishing this arbiter to an inaccessible realm from which it creates accessible appearances (like the Wizard of Oz) makes about as much sense to me as claiming that today is but a representation of yesterday's *real tomorrow.* What does it buy one, but a big headache? Who needs it? What's wrong with immanent truth, as reflected in phenomenal coherence?

I am interested in your staking a claim to what I see to be the epistemological high ground (in my view, solipsism is *not* the high ground). My motivation is selfish. Because you are the principal champion of psychomodular control theory, I have a vested interest in your being in the best position to defend both your psychomodular theory and your epistemology. And since I do not see your psychomodular control theory as implying a solipsistic epistemology, I see no reason for you to defend that indefensible epistemological position. My inability to persuade you to give solipsism a wider berth than you do leaves me ambivalent about my efforts to that effect, because I am not interested in being a mere disturbance.

Bill Powers: A model might be epistemically correct, but we will never be able to prove that. I was not saying that we have complete epistemic success part of the time, and incomplete success the rest of the time. I was saying that we do not *ever* know whether our models of reality are successful or not, because the only way to check for that success is to repeat the process that led to the models in the first place: there is no independent check.

Perhaps we could get to the nub of this matter sooner if you would give me one example—any example—of a case in which we have complete epistemic success in verifying that there is a real counterpart of any perception. We apparently agree that there are some cases in which uncertainty remains, so there is no point in dealing with them. My claim is that there is *no* case in which we have reached certainty, so you should be able to demolish my claim with a single counterexample.

Wayne Hershberger: I seem not to be making myself clear; you are looking for me in the wrong direction. My argument (actually, Hume's) is not that any particular case affords certainty, but rather that *every case is entirely uncertain* (i.e., "verifying that there is a *real* counterpart of a perception" is absolutely impossible, even as a matter of degree). Hume's arguments to this effect are called Radical Skepticism for good reason. Consequently, modern science uses a coherence, rather than a correspondence, theory of truth—where reality has no ontological status.

When asked how he discovered the laws of chemical compounding, Linus Pauling replied, "I made them up." Pauling avoided any claim of having gained epistemic access to Reality—because such a claim would be gratuitous (God's Reality is a matter of faith, and serves no scientific purpose). Rather, Pauling made up a parsimonious model which provides a very *coherent* account of the chemical phenomena in question. Any claim that such human-made models correspond, in varying degrees, to some divine original is epistemically empty.

You have made up a parsimonious model of living control systems which control the value of inputs from their environments. Sometimes you have used the word "virtual" to refer to these controlled variables, because they are defined by the input functions which process the input. But it would overstate the case to claim that the environment contributes nothing to the values comprising these virtual variables. That is, only by overstating the case is one misled to suppose that your model implies solipsism. Your model addresses questions of correspondence between what is inside the brain and what is outside the brain, but that is physiology, not philosophy.

Bill Powers: Comes the dawn. I feel like a wrestler who has converted an advantage to a position flat on his back. Your previous arguments have given me the impression that you believe there is a reality outside

of perception *and* that perception somehow manages to represent it veridically. Now it seems that you're saying that human perception bears no verifiable relationship to any universe "behind" or "beyond" perception, which is, of course, the position I have also been taking.

Unfortunately, your language still leaves me wondering what precisely is your position on the constraints we can detect between actions and perceptions, and on the significance of models. You say, "Any claim that such human-made models correspond, in varying degrees, to some divine original is epistemically empty."

If it weren't for that word "divine," I would be more sure of how to take what you mean—I trust you're not accusing me of religious fervor. Would you still allow for a correspondence to a "non-divine original"—i.e., a natural universe that exists independently of our perceptions of it? In other words, are you opposing a religious view of reality, or any view that there is (or could be) a reality more inclusive than what is perceived, whether or not we can be certain about its nature? More on this at the end.

The main conundrum comes up when you say, "Your model addresses questions of correspondence between what is inside the brain and what is outside the brain, but that is physiology, not philosophy." A problem is created by talking about what is inside the brain and what is outside the brain. The problem arises when we assume that we, as conscious beings, are conscious because of activities in a brain. Allow me to elaborate.

In order for my model to be consistent with the physical model of reality (both, I quite agree, being "made up," so that physiology, too, is "made up"), there are certain relationships between physical-model variables and neural-model variables that must be assumed. The physiological neural model allows for no way of getting information from physical stimulation other than through interaction of physical variables with neural sensors. For example, in the physical model there is a made-up entity called the photon. The signals in the neural model's retina supposedly arise from absorption of photons. However, the neural signals carry no information about the origins of those photons; furthermore, there is an infinite number of different photon energies and fluxes arriving from an infinite number of directions that will yield exactly the same neural signal in any given receptor.

Given a model of physical optics and observations of reflection or emission sources, we can construct a model of the origins of the photons and show that this model is consistent with an array of neural signals that amounts to a map of the scene toward which the eye's lens is directed. So far, so good. But if we then look at the basis for accepting the physical model—which includes things like "lenses" and "objects" and "light rays" going through something called "three dimensional space"—we find that there can be no basis but observations made by the same means by which an "observation" of a photon is made. We identify objects by looking at them with our eyes; we verify that there is a photon flux by interposing a light-meter (which we see) and reading—with our eyes—its indicator. So, from the standpoint of the neural model, the physical observations we are using to assign an external source to the visual neural signals arrive in the brain by exactly the same means as the signals we are trying to explain.

This is a not a problem if we adopt a point of view from which we can see both models, the model constituting an exterior physical world and the model constituting an interior world of signals in a brain. It is not a problem if we add to the other two models a model of a non-neural conscious observer which is not confined to a brain. It becomes a problem only when we decide that the model of the brain must be a model of ourselves, the observers and thinkers.

When we adopt that view, as I do, we can no longer take the third-party omniscient view. The hypothesis is that we, who are thinking about perception, are brains like those in the model. Therefore, we must be dealing with the external world (represented by the physics model) by the same means just proposed: through neural signals. If this is true, then the physical model is *not* outside the brain. It must be located inside the brain-model, as part of that model. It is a construction existing as patterns of neural signals related not by physical constraints outside us, but by abstract rules and computational processes taking place in our heads.

When we apply this reasoning to purely physical models, there isn't much difficulty except with people who insist on reifying photons, electrons, quanta, phlogiston, and so on. The real difficulty arises when the external world we are thinking about is the world of subjective reality: the world we experience directly. This is clearly not the world of the physical models. Between the physical models and this world of direct experience, there are few points of contact. For the most part, physical models consist of entities and relationships that are not evident in direct experience. Here and there are points where, usually through the use of instruments, but not always, a physical variable corresponds to an experiencable variable. With the unaided eye, we can perceive an approximation to what a physicist calls "distance," although by using instrumentation like radar or optical range-finders, we can arrive at meter-readings much more consistent with physical theory than is the direct apprehension of distance. But when it comes to functions of distance such as gravitational acceleration or potential energy, direct experience remains blind.

So where do we put this world of direct experience, with all its objects and sounds and smells and relationships and people? It is not

represented in either the physical model or the (physiological) neural model. I have elected to put it into the neural model, but not in the form of neurons. It exists in the brain as a weightless, massless organization of neural signals, the appearance and behavior of which is precisely the appearance and behavior of the world we experience. Certainly this assumption creates a mystery; more than one. The main one is who or what is it that apprehends this collection of neural signals in such a way that it takes on the appearance that we experience?

The most obvious error to be made at this point is to say that this mysterious observer is the agent who imposes interpretations on the neural signals so they become objects, relationships, processes, concepts, and so on. But as everyone knows by now, that simply requires expanding the model to explain how these interpretations are made. My way of avoiding this error has been to propose the levels of perception in my model. By looking for classes of perception in the apparently real world around me, I attempted to show how neural processes can themselves create signals which contain the interpretations that are needed. While this initially seems to rob experience of some vital qualities, a close examination of any particular example of this problem shows that it does not exist. These vital qualities can't be pinned down by direct inspection, either. When one attempts to isolate them for a close look, they lose any special quality and become just an amount of something that can be more or less present. Just like a neural signal. The only specialness that there is exists in the entire collection of neural signals, each behaving in the context of all the others.

The other function of the levels is to enumerate and classify types of perceptions ranging between what have been considered "concrete" and "abstract" perceptions. By showing how successive levels of interpretation can form a link between the concrete and the abstract, the model removes the necessity for explaining these interpretations by assigning them to a homunculus. As each new level is considered, the subject-matter with which it deals is stripped out of the homunculus and returned to the physical brain. In the end, the homunculus contains only those functions of observation that are not accomplished by the brain model. And all that is left is awareness.

We now seem to have a model, itself a neural model, that contains a physical model in a nervous system. But the nervous system is basically a physical conception: it is a subset of the physical model. Logically, if one model is contained by and contains another model, there can be only one model. But there is another answer: it is that both models are contained in direct experience, and they interact with each other..

The ultimate reality, therefore, is direct experience. That is the superclass within which models exist. This leads us, finally, to the ultimate mystery.

We can divide direct experience into things we do and things that happen. Many things that happen proceed without any need for our action. Among such things, we can detect consistencies and dependencies. This leads us to formulate expectancies which, when formalized, we call laws of nature.

We can also take actions, which are the set of all those experiences that we can influence by an act of will. We find that these actions, themselves capable of being experienced, affect other experiences. We can learn to create some experiences that are not directly subject to acts of will by varying those experiences that are directly willable: the whole is an act of control. Through long experience with this kind of act of control, we have found regularities that show how we must act in order to control many kinds of experiences. The reasons for these regularities are not evident in experience—there is no a priori basis for expecting any particular act to have any particular effect on something else.

This is where we get the idea of a natural world of regularity that lies outside the boundaries of experience. And this is why we build models, both physical and neural. With models, we hope to probe into that mystery that is hinted at by these unexplained regularities. We hope to reduce the complexities of these piecemeal regularities by finding underlying simplicity; I think this is what we mean by "mechanism." In physics, simplicity is attained by imagining a hidden world of fields and particles, energy and momentum and entropy. The few kinds of variables in this world lead to the vast multiplicity of different-seeming phenomena in the world of direct experience.

The question, Wayne, that you and I have not brought out into the open and resolved between us, is whether these models constitute increasingly good approximations to something beyond experience, or whether they are simply "summaries of observations." The complex picture I have tried to lay out here should indicate my view. Clearly, I don't think that either physics or neurology is as good an approximation as is usually assumed. There is too much of the human observer entwined, unanalyzed, in all our models. The very name "particle" in physics shows this. But I think that there is evidence of agency outside us (other people, for example), and evidence of relationships imposed by unseen means (e.g., other people's intentions). I think that there is structure inside the Reality Box, and that while we can never arrive at a unique representation of it, we can arrive at an equivalent representation, equivalent in the sense that our models show one way it could be constructed inside, functionally equivalent to the way it is constructed. I see no contradiction in saying that all we will ever know for certain is what our own brains present to awareness, while maintaining that uncertain knowledge is not empty.

Is the remaining problem, perhaps, what is meant by "epistemic"? If

"epistemic" knowledge is certain knowledge, then the argument resolves itself: there is no such thing outside direct experience. But to say that a proposition is epistemically empty does not then mean that it is incorrect or empty of significance, because that would say that all of experience is, with complete certainty, incorrect or empty of significance beyond itself—an epistemic fact which, of course, we can never verify.

Martin Taylor: Bill, in answering Wayne, says: "But if we then look at the basis for accepting the physical model—which includes things like 'lenses' and 'objects' and 'light rays' going through something called 'three-dimensional space'—we find that there can be no basis but observations made by the same means by which an "observation" of a photon is made. We identify objects by looking at them with our eyes; we verify that there is a photon flux by interposing a light-meter (which we see) and reading—with our eyes—its indicator. So, from the standpoint of the neural model, the physical observations we are using to assign an external source to the visual neural signals arrive in the brain by exactly the same means as the signals we are trying to explain." But one actually can test the "existence" of the things detected through photons by using other senses—the acoustic effects and so forth. These form a set of converging operations that help to reduce the set of possibilities for interpreting the perceptions obtained through one sensory system.

If that were all there was to it, the same argument could be made, but extending the notion from "photon" to "physical energy exchange phenomena" or some such. But that is not all there is to it. There is the volitional aspect of what and how we choose to observe.

Let us presume a deceitful Nature, and a passive (multi-sensory) observer. This Nature could present us with any of an infinite number of sources for PEEPs (Physical Energy Exchange Phenomenon, plural) that had the same effect on our sensory organs. But when *we* choose which aspects of the universe to test, and in what way, the deceit becomes much harder to sustain. That's the fundamental difference between an observational science like astronomy and an experimental one like physics. Psychology is somewhere in between.

The difference between active observation and passive observation was clear to Gibson, who distinguished "haptic" from "tactile" perception. You can try it yourself. Have a bunch of objects available, and a friend. Close your eyes; have your friend take one of the objects and touch it to your open hand in various orientations and ways. What you perceive is a set of touches, some soft, some warm, some sharp, and so forth. Now have the friend place the object in your hand for you to manipulate. What you perceive is not a set of touches, but an object.

In either situation, you might be able to determine which object from the set was touching your hand, but when you yourself choose where and how the object contacts (or fails to contact) your hand, it has a completely different subjective quality.

Wayne Hershberger: Bill says, "... it seems that you're saying that human perception bears no verifiable relationship to any universe "behind" or "beyond" perception, which is, of course, the position I have also been taking."

Yes.

And that, I believe, is what Gibson was saying as well, or at least trying to say. And Kant, too. Your prose is more lucid than most of theirs, but as I read you (Hume, Kant, Gibson, and yourself), you all seem to be motivated by this same epistemological insight. However, the other three believed that this insight also implies that this hypothetical Reality has no empirical basis. So did Plato, who claimed that Reality's basis must, therefore, be rational.

You, on the other hand, seem to accept a Reality appearing to require neither basis. That is, the Reality to which you persistently refer, despite your above remark, appears to be neither an induction nor a deduction, but rather an abduction; you seem to pluck it out of thin air. For instance, while admitting that a perception cannot be proven to be a veridical representation of Reality, you are wont to claim that neither can it be proven that perceptions do not approximate Reality to some degree. This begs the question of the Reality itself!

I would encourage you to accept the harsh implications of your own epistemological insight and not backslide, admitting through a back door what you have banished from the front. The challenge, remaining to be addressed, is what exactly is meant by the term real or true, that a perception may be identified as veridical, as opposed to illusory, *and* at the same time *not* be regarded as a representation of Reality? That is, what is it that distinguishes a veridical from an illusory perception if *not* the perception's degree of correspondence to some transcendent reality—something you claim to eschew as an arbiter of truth?

In practice, we seem to use a coherence theory of truth. Laymen and scientists alike regard a perception that cannot be replicated as illusory. A perception that does not survive the layman's double take is an illusion. An empirical observation that science cannot replicate is no fact.

You say, "So, from the standpoint of the neural model, the physical observations we are using to assign an external source to the visual neural signals arrive in the brain by exactly the same means as the signals we are trying to explain." Although this is often the case, it is not necessarily the case. Let me address the flip side of this question—to which you also allude: is there "a natural universe that exists independently of our perceptions of it?"

I submit that the ability to register luminous flux with virtually any retina or photomultiplier tube provides the very sort of independence referred to here. Only those perceptions which are demonstrably replicable across observers are objective perceptions, or "objects," as we are wont to say, for short. This is the type of independence required of the objects comprising our *natural universe.* When it does not matter who or what makes the observation (i.e., the results are independent of the particular observer), the perception is said to be objective, or to be an object. The natural order is immanent in experience and not to be confused with some hypothetical Reality that transcends all experience.

You ask, "So where do we put this world of direct experience, with all its objects and sounds and smells and relationships and people?" In the phenomenal world of time and space—which Kant recognized as intuitions (meta-models?). "Silicon Babies," an article in the December 1991 issue of *Scientific American* said something relevant here. Speaking of robots as Rodney A. Brooks conceives them (which is similar to the way you conceive them), the author of the piece said: "Subsumption architecture relies largely on the nature of the outside world rather than sophisticated reasoning to structure the robot's actions. For example, if the robot encounters an obstacle, the important thing is to go around it... The robot may not need even to remember that the object is there—after all, it will detect the obstacle perfectly well the next time it approaches it. (p. 128) The expression "outside world" in this passage obviously refers to the robot's environment. This world outside the robot is not outside the robot's realm of experience. Neither is the robot's world a re-presentation (copy or memory) in the robot of a world actually transcending its experience such as our model of its environment. The robot merely *registers* its environment in its "inimitable way." Call it modeling the environment. Of course, the way a robot registers/models its environment is not actually inimitable; identical robots would register/model in the same way.

Similarly, if your expression "a natural universe that exists independently of our perceptions of it" refers to something other than the natural order immanent in the psychophysical flux we call experience, it is surely a reference to the perceptual/conceptual models that are constructed out of that flux registering that immanent order.

"So where do we put this world of direct experience, with all its objects and sounds and smells and relationships and people? It is not represented in either the physical model or the (physiological) neural model." Right you are. The objects of direct experience are not part of either of these scientific models (physics or physiology), because these objects of direct experience are themselves models—empirical, if not scientific. They are the layman's perceptual models, analogous to science's conceptual models; both types seem to involve a lot of neural processing, just as your own theoretical/scientific model says. Both types of models are modeling the same natural order. They are twin born of experience. One type of model is not modeling the other type of model. (Only psychological theories such as your model are reflective, modeling the process of modeling itself.) Neither is the basis of the other. The basis of both is the natural order which, as you say, "exists in the entire collection of neural signals, each behaving in the context of all the others." Strike the word "neural," and I think I could buy it. That is, if this psychophysical flux has any essential characteristic, surely it involves informing and being informed—in a word, signaling. Matter which can neither influence nor be influenced by other matter, doesn't matter. But the signaling does not begin and end in the nervous system.

Bill Powers: Wayne, you are telling me that Hume, Kant, Gibson, and I all seem to be "motivated by this same epistemological insight"—yet "... the other three believed that this insight also implies that this hypothetical Reality has no empirical basis." They might not have had the same attitude toward empiricism that I have. For instance, they might have been of the opinion that only empirical facts can be true and real. This would put them in good company, but it would be the company of those who customarily elevate statistical preponderances to universal certainties. I have met very few empirical facts that did not contain easily discernible uncertainty; certainty is achieved by ignoring the actual data and plunking a dot down in the middle of the scatter: that's the *real* value. Most of the time, there is a background of approximations, arbitrary assumptions, and interpretations without which empirical data would have no meaning. These assumptions, and the fact that someone is interpreting, are not mentioned in polite company.

So to say that a hypothetical reality has no empirical basis is not the indictment it might be if there were such a thing as pure empiricism untainted by human imagination and interpretation.

Even when we confine our observations to the omnipresent psychophysical flux, we see things that are contradictory. Our judgments of width and height do not agree with readings from calipers. Our judgments of straightness do not agree with straight-edges. Our judgments of relative temperature do not agree with thermometers. Our judgments of relative brightness do not agree with photometers, and our judgments of relative color and, especially, color composition do not agree with spectrographs. In realms of more complex observations, we do not agree with each other about palatability, difficulty, comprehensibility, spelling, grammar, or miracles. We observe nonexistent phantom arrays created by what we know to be a single stationary flashing light. We

don't even agree on a color like "green." When we do agree, the spectrograph can tell us we are looking at different colors; when we disagree, the spectrograph can tell us we are looking at the same color.

You say, "An empirical observation that science cannot replicate is no fact." Is there *any* empirical observation that science can, literally, "replicate"? Replication never in fact occurs: perhaps that is the *only* replicable fact. What happens is that we make a series of meter readings that disagree with each other, and then we say that there is a Real value that lies somewhere within the range of the readings. We can't actually "replicate" a reading (in fact, if we get exactly the same reading twice in a row, we tap the meter from then on). We can't even replicate the scatter in a series of readings. We replace the scattered, variable, inexplicable, individual observations with an idealization that we conceive of as the real observation. In doing this, we create precisely the reality I am talking about: a reality that we do *not* observe, but accept because it makes sense of experience. Empiricism itself leads to acceptance of a reality that underlies observation, but is not the same as what we observe.

Much of what seems to be replication is a product of the human capacity for categorization. We can make observations that vary widely, yet make them appear to be the same by classing them together. If we ask 100 people, "Are you in favor of abortion?" and 60 percent of them say "no," we lump the 60 "no" answers together and say that they indicate the same opinion about abortion. In fact, we don't know what question the respondents were actually answering; all we know is the question we heard ourselves asking, and what it implies to ourselves. Some were thinking, "No, not even to save the life of the mother," while others were thinking, "If it's necessary to save the life of the mother, but in general, no." Those are both "no" answers, aren't they? "No" is "No," isn't it?.

'Empiricism fails as soon as you go beyond a description of a snapshot of the psychophysical flux. Take as simple a thing as a lever with the fulcrum in the middle. Pushing down on one end, you observe that the other end goes up, while the lever itself tilts. You can easily satisfy yourself that if you do push one end down, the other end will rise, and if you don't push that end down, the other end will stay where it is. There is little more to be determined, empirically, about the behavior of this lever. Now, is there any connection between your pushing on one end of the lever and the subsequent tilt of the lever and the rise of the other end? Have you given an adequate account of the lever by reporting just the facts of what happens in the psychophysical flux? Obviously not. You have reported three facts: pushing down, tilting, and rising. Is the first fact directly influencing the other two facts? Of course not: at least one physical property of the lever, its rigidity, is needed in order for the one fact to lead to the others. Facts do not influence other facts just by existing. Given two identical-looking levers, one may behave as you expect, while the other simply bends. The difference that explains the difference in behavior is not among the empirically observed facts. It is an imagined property of the lever, deduced from its observed behavior. That property is part of the unseen reality of which I speak. Does this "rigidity" actually exist? It might. But it might also be a consequence of unseen factors such as intermolecular forces, none of which itself is "rigidity," but which together have the consequence of imparting rigidity to the lever.

The "physical" part of the "psychophysical flow" is imagined or deduced (it is deduced, then imagined). It is not observed. There is no meter that measures rigidity, or whatever factors give rise to it. And if there were a meter measuring the straightness of the lever, it would not be measuring the causes of rigidity.

Or consider another case that tells us even more about our relationship to reality: a wall switch that operates a light in the ceiling. We can easily determine that when the switch is up, the light is on, and when it is down, the light is off. We can manipulate the switch and reliably observe—for a while—that the state of the light obediently changes. If all we care about is operating the light, we are finished.

But some of us assume that there is some connection between the switch and the light. We can't observe this connection without destroying the wall and ceiling, yet we have little doubt that there is one. This imagined connection has no basis in our empirical observation of *this* switch and *this* light, yet as a matter of faith we accept the existence of the connection. We do not accept effects at a distance, in most cases.

Does it now shake our faith if we flip the switch up and the light fails to come on? Not at all; we deduce that the switch has failed, the bulb has burned out, or there has been a power failure. We don't observe those explanations empirically, either—although we would like to check them out by some indirect means, like turning on a floor-lamp to eliminate the general power-failure explanation. We don't actually need to visit the power plant.

We begin to suspect the switch when we wiggle it up and down and observe that the light comes on again. But now we observe an odd thing: the light now comes on when the switch is *down* rather than up. We can still toggle the state of the light by moving the switch, but the relationship has reversed. Have the innards of the switch suddenly turned upside-down? That seems ridiculous.

Then we remember that high-school puzzle, the two-way switch. There is, we realize, *another switch* somewhere else that also controls this light. There is someone fiddling with that other switch!

Is this an empirical observation? No, it is a memory-based guess

about a hidden reality. Can we be certain that there is another person fiddling with the other switch? No. How could we be? We haven't seen the other person or the other switch yet, and if we have to catch a plane, we might never do so. It's possible that the spring in the other switch broke and let the switch flip down, with nobody operating it. It's possible that there's a relay in the circuit that short-circuited. Yet there is value even in our wrong guesses, because they are possible explanations and in other circumstances might be the correct ones. These possibilities relate not to the empirical world, the psychophysical flux, but to a world beyond what we are sensing: inside the wall, in another room we haven't visited, in a power plant we have never seen.

Most of the world within which an individual human being makes empirical observations is outside the scope of that person's perceptions, yet its imagined state forms the context within which what is observed is interpreted.

There are some aspects of the hidden reality that we strongly suspect to exist, but which we will never be able to verify. Is there really an electromagnetic flux propagating through empty space, a flux that we call "light"? There is no way to check this. We can only say that when some sort of receptor *stops* the imagined propagation, we get some sort of meter reading. There is simply no way to detect light in flight. Human reason screams at us that *of course* light has to propagate through space in order to reach our detectors—but that is not and never will be an empirical fact.

Nearly all of our meter readings interfere to some extent with what is being measured; meter readings lie. An ordinary volt-ohm meter draws current when it measures voltage. That current causes the actual voltage to drop a little—sometimes a lot. In electronics, we learn to measure the meter's resistance and calculate that of the circuit, and correct the reading to the "true" value. So the meter reading we see has to be corrected to indicate the voltage that really exists.

In order to estimate how hard a suitcase is pressing down on the rug, you have to lift it. Data from polls has to be corrected to show what the true opinions would have been if everyone had been telling the truth. When we bargain with another, we try to estimate from the offers the other is making how much that person is really willing to pay. When we see a car in the convex right-hand outside rear-view mirror, we see a• label saying, "Objects appear farther than they are." When a PhD candidate fumbles a question, we make allowances for her nervousness. When an agent says, "I'll give you a call if anything comes up," the actor ceases to expect a call. We are always making adjustments to observations, denying the validity of empirical data, to bring our actions and expectations more in line with a world that underlies appearances. We are better off doing so than not doing so, even though we are sometimes mistaken in not taking literally what is before our eyes. Sometimes the phone rings and it is the agent telling the actor to report for rehearsals.

There are really no justifications for denying the existence of a reality that is different from the one we experience, even the one we experience through the use of scientific instruments. The scientific instruments themselves shout at us that there is something going on that is invisible to us. If we were strict empiricists, we would report analog meter readings in radians, not volts or pH or counts per second or RPM or pounds per square inch or quarks per cubic meter. But we do not: doing so would leave us with a world that made no sense. We have constructed an elaborate network of imagined entities and relationships that purport to live in the world on the other side of the meter readings. While this conceptual world might miss the mark and might describe only a projection of a much larger space onto the dimensions to which the meters are sensitive, it might be correct in some respects, particularly respects having to do with derivative notions like conservation of energy or control. We will never know, of course; our meters and our sensors stand between the observer and the reality. If information is coming in to us through these channels, we still don't know what it is about. The incoming information carries no identifying labels.

Still, it pays to guess, as long as we are alert for the evidence that says we should change the guess. But today's empiricism is tomorrow's illusion, and often today's.

Wayne Hershberger: Bill, you use the word reality very differently in the following two passages. The first usage refers to a created reality; the other refers to a hidden reality (my italics below), which you say is unverifiable. I find the former concept very useful, but not the latter. The latter is essentially a contradiction of terms, a paradox. This paradox is readily resolved, however, using your own very cogent arguments. That is, if the words "created reality" used in the first passage below are substituted for the words "hidden reality" in the second passage below, what you then say makes sense, don't you think?

Passage 1: "We replace the scattered, variable, inexplicable, individual observations with an idealization that we conceive of as the real observation. In doing this, we create precisely the reality I am talking about: a reality that we do *not* observe, but accept because it makes sense of experience. Empiricism itself leads to acceptance of a reality that underlies observation, but is not the same as what we observe."

Passage 2: "Most of the world within which an individual human being makes empirical observations is outside the scope of that person's perceptions, yet its imagined state forms the context within which what is observed is interpreted. There are some aspects of the *hidden*

reality that we strongly suspect to exist, but which we will never be able to verify. Is there really an electromagnetic flux propagating through empty space, a flux that we call 'light'?"

There is not just one created reality.

There are many created realities, some comprising perceptual "objects" (e.g., a laser beam), some comprising conceptual stuff (electromagnetic flux), both dealing with the same phenomena. When we attempt to order these created realities by their truth value, we find ourselves using parsimony and replicability as our criteria.

However, I am not sure that the conceptual stuff is any more or less empirical than the perceptual stuff. The mean of a set of data is a datum too. Further, although the arithmetic mean is not the only measure of central tendency which could be defined upon the raw data, it is very precisely constrained by its raw data.

It seems to me that our created realities transcend the raw psychophysical flux in essentially the same way that a mean transcends its data; that is, they are precisely constrained by the psychophysical flux—barring miscalculation. The input functions in your psychomodular control-system model create perceptual realities in this constrained way. That is why your theory is not solipsistic.

On another matter: "Empiricism fails as soon as you go beyond a description of a snapshot of the psychophysical flux." On the contrary, empiricism fails when you artificially restrict it to snapshots of the psychophysical flux.

Bill Powers: Wayne says: "There are many created realities, some comprising perceptual 'objects' (e.g., a laser beam), some comprising conceptual stuff (electromagnetic flux), both dealing with the same phenomena."

Is this a correct diagram of what the above says?

perceptual stuff→perceptual created reality
"The same phenomena"→
conceptual stuff→conceptual created reality

In other words,

"The same phenomena" are contained in
Perceptual stuff, which is contained in
perceptual created reality, and
Conceptual stuff, which is contained in
conceptual created reality.

Or, is it more like

Phenomena→
  Perceptual stuff→
    Conceptual stuff?

Wayne Hershberger: Neither alternative is quite correct, but the first is the closer of the two.

I would say that the natural order immanent in the psychophysical flux is realized both in the form of perceptual stuff and in the form of conceptual stuff. And when these two types of realization deal with the same phenomena, they often seem to depict a contradictory nature. As perceptual stuff, the desk at which I am sitting (an objective perception) is a static and solid object, but I can simultaneously regard it conceptually as a collection of whirling dervishes hurtling about within a confined space virtually as empty as an inflated balloon. The two types of stuff (i.e., the two types of reality) are, as I've said before, twin-born of the psychophysical flux (there is nothing like reason to make sense of something). This is not to say that objective perceptions (e.g., meter readings) are not useful in testing hypotheses derived from a theory representing a putative conceptual reality, but only that the perceptions involved in such tests should *never* directly involve the phenomenon being conceptualized. That is, you do not settle the geocentric vs. heliocentric world-view issue by watching a sunrise.

Further, I do not mean to say that there are only two levels or types of reality. For example, within the sphere I am calling perceptual stuff are to be found such things as real illusions (e.g., Ames' window, colored shadows, etc.), which are as different from the other perceptual stuff called real objects as the real perceptual objects differ from the conceptual stuff.

Furthermore, speaking of realities, as if there were a limited number of types, is misleading. Reality is but a dimension upon which we order the truth value of our countless epistemic creations—using replicability and parsimony as our criteria.

Bruce Nevin: An observation about the ongoing discussion of epistemology: I believe that in your language, Bill, you assume for yourself the perspective described by your (our) theory. This is an excellent way to test its adequacy, but perhaps deserves calling out for notice.

The theory or model requires there to be a "boss reality" in the environment to complete the feedback loop. Without it, perceptual control is impossible. So, from its perspective, there has to be a reality there.

However, the theory does not describe this reality or prescribe any attributes, other than that it be present and in at least some respects

stable so it is capable of being modeled within the perceptual control hierarchy.

In an important sense, this environmental reality is hidden from the perceptual hierarchy. Its only access to it is proximal stimulation of intensity sensors.

In an important sense, this environmental reality is not hidden from the perceptual hierarchy. Its model of it is presumed reasonably veridical because it in fact accomplishes perceptual control requiring feedback through the environment.

There are two senses of "model" possible here. The perceptual hierarchy may create a model of reality at higher levels of control. It is likely to use language to do this. Like any model, it is imperfect and requires periodic amendment. Because it is maintained at a high level of the control hierarchy, response to conflict is slow. The capacity for amendment is in the control hierarchy that holds the theory, not in the theory itself.

The second sense is that there is a model immanent in the perceptual control hierarchy as a whole. In its capacity to control perceptions, whatever the feedback through the environment might be, the perceptual control hierarchy is a kind of reflection of the environment. Like any model, it is imperfect and requires periodic amendment. Unlike many models, it includes this capacity for amendment in itself. Speed of response to conflict varies with the level of the conflict. Call this model 1 and the other sense of "theory" model 2.

I think it was the more primitive sense of the implicit, immanent model 1 that applied when I said "capable of being modeled within the perceptual control hierarchy." Must "boss reality" be capable of being modeled in both senses? Does it have to be able to sit still for its picture to be taken, so to speak?

Bill, you are assuming the perceptual-control-theory model 2 (theory) as your perspective in talking about knowability. Anything not countenanced in that model you suspect is illusory. And you are using your model 1 to test the model 2, as indeed are we all.

We have a primitive sense that what appears to be there in the environment is real (naive realism). We can talk ourselves out of this if our model 2 (theory) calls for it. We can also ignore perceptual signals if our expectations say they are not there—if associated error at higher levels is not significant. The two cases seem to me entirely alike.

Suppose there were a physical, mechanistic basis for our primitive sense of the reality of our perceptions. This doesn't entail that this sense be articulate enough for us reliably to distinguish illusion and hallucination, something that we appear not to do. I would base this sense in the continuum of the environment outside the skin with the biochemical and biomechanical environment inside it. Awareness of this would probably have an intuitive and emotional quality rather different from the attention to perceptual signals that usually concerns us. At its peak, an awareness of being part of a larger unity, perhaps.

But anything not countenanced in our model 2 or theory (the one we have adopted as our perspective) we suspect is illusory. For good reason: naive realism runs into well-known difficulties. And these intuitive apprehensions are global rather than particular. I think, too, that this sort of apprehension of the reality of reality is undemonstrable within the perceptual hierarchy. It is only apprehensible to it. Does that mean it is part of the environmental feedback for the control hierarchy? I don't know.

As we all know, "boss reality" doesn't really sit still for its picture to be taken. A model 2 or theory is possible only by categorization, subsumption with neglect, conventionalization, language. Only a model 1 with its continuous, live tracking can be veridical, and that only in a limited and local sense.

Theories are models 2 of perceptions in our model 1, which is *a* model or reflection of reality. All but the specific environmental feedback being tracked and controlled for is hidden from the model 1. Potentially, nothing is hidden from models 2, but their precision and accuracy are suspect. Partly this is because they are constructed using conventionalized social products such as verbalized categories. Partly it is because their responsiveness in the face of *aniccha,* impermanence, is too slow.

The obvious generalization is to speak of a level -1 model as most local and most accurate, a level -2 model, and so on, up to the models 2 of the system-concept level (some idealization here about the sequential separation of levels, as we know). Assuming of course that our model 2 of perceptual control is veridical.

As students, we take your verbalizations about error signals from your comparators as indications for setting reference values in our own. Which we might do, or we might verbalize error signals in turn. A reciprocal process called communication, of course.

This is a test. This is only a test.

Martin Taylor: Bruce, I would like to raise a flag to signal my objection to your claim that model 1 can be veridical "only in a limited and local sense." It is true that mathematical theories of physical "reality" take us a lot further than intuitive physics in predicting the behavior of the world, but it is not so clear that this is true for the less simple sciences. Physics is, after all, the only science so simple that the most intelligent humans have a reasonable hope of understanding some of it.

But we do behave reasonably successfully in the much more complex world of nutritious and poisonous foods, friends and enemies,

and so on, for which linguistically (e.g., mathematically) based models do a lousy job.

Bill Powers: Wayne: "I would say that the natural order immanent in the psychophysical flux is realized both in the form of perceptual stuff and in the form of conceptual stuff."

You have picked the first of my representations.

So does your model looks like this?

A:

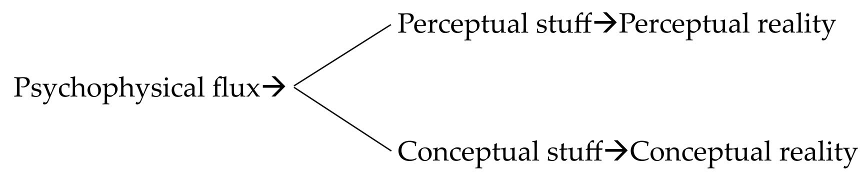

Or would you draw it this way?

B:

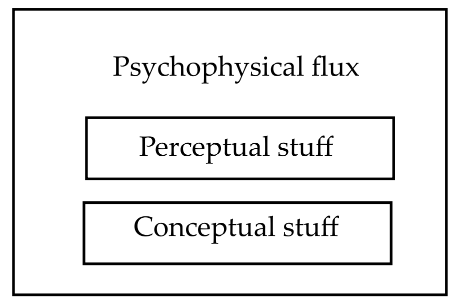

(Psychophysical flux includes more than perception and conception together.)

Or this way?

C:

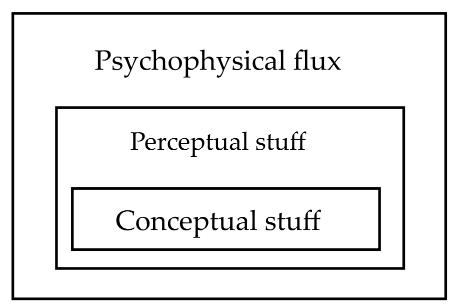

(Or with perceptual and conceptual interchanged.)

Or this way?

D:

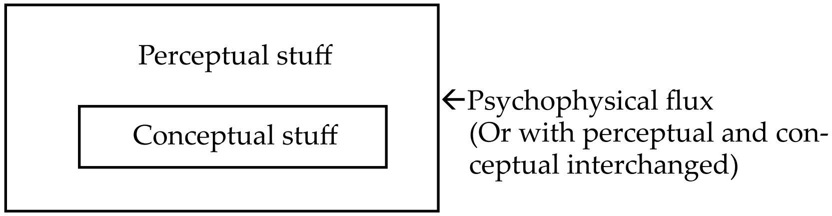

Or this way?

E:

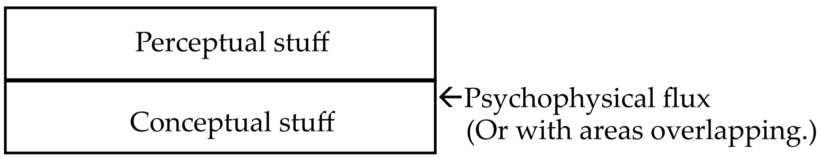

Or some other way?

Bruce: Isn't the concept of models 1 part of models 2?

What we know is the world of direct experience. If this is what Wayne means by the "psychophysical flux," then I would agree with him, but I would not use that term. To say either "psycho-" or "-physical" is to introduce description, characterization, and classification, which are added to what we observe and go beyond what we observe. The term "psychophysical" introduces a model and a theory right at the beginning of the discussion, which for me is too early. I want to explain the psychophysical flux, not assume it as a premise.

I do take the viewpoint of the hierarchical-control-theory model in explaining the structure of subjective experience. I take the viewpoint of the physics model in explaining the structure of the environment and brain. The purpose behind doing either one is to explain the way the world seems to me, as it is directly experienced. The ultimate criterion of truth for me is not any principle of philosophy, physics, neurology, or logic. It is simply whether these viewpoints, considered together, explain in an honest, testable, and self-consistent way what is puzzling to me about the world that I experience directly. I am not puzzled about its existence. I am puzzled about why it works as it does.

Why do I prefer honesty, testability, and self-consistency? Because

I like such things. Everything seems to work better when such principles are accepted as constraints. Highly recommended. Best buy. But not proven.

Model 1 is an assertion about the brain's built-in abilities to generate a perceptual world based on an external physical world. So it already contains an assertion about an independent physical world and a world of brain function. Model 2 extends model 1 to higher levels of brain function, by explicitly introducing modeling in terms of symbolic processes. But aren't these two models simply ways of classifying subsets in the general hierarchical-control-theory model?

Is either one more "real" than direct experience?

All of science, in my view, revolves around direct experience. We don't require models to explain each other: they are all required to explain what a human observer can experience, and how human actions and spontaneous changes in the observed world affect human experience. A theory or a model must bring order into the relationship between actions and perceptions, where "perceptions" includes both unaided human observation and observations of the readings generated by instruments (there's no fundamental difference).

To me, it is simply a fact that I don't experience anything but the surface of the world. I don't understand how anyone can claim that this is all there is. All you have to do is dig a hole, and you will see that the surface of the ground is held up by something else. That something else is hidden from the senses until you dig the hole. What holds *it* up is hidden until you deepen the hole. I see no hope of ever seeing what holds it all up, at the center of the Earth: long before we could get there, our shovels would melt. I'm willing to entertain the possibility that there is really a nickel-iron core in the center of the Earth. I do not, however, confuse that possibility with an actual experience of the Earth's core, the only incontrovertible verification of the possibility. Another theory might claim that there is a black hole at the center. The universe is not expanding: we are shrinking into the event horizon of the black hole. How would we verify that?

Hierarchical control theory is verifiable to the extent that it predicts classes of perceptions that we can actually experience and control. Control theory is verifiable to the extent that it predicts relationships among actions and perceptions that we can actually experience. In neither case, however, can we verify intermediate processes required to make the model work but which don't themselves correspond directly to aspects of direct experience. None of us, for example, can verify that these processes take place in a brain. That is conjecture. We will only know that these processes take place in a brain when we can link each process to a perception or measurement of activities and relationships in a brain and show beyond doubt that affecting each process as measured affects direct experience exactly as predicted. On the way to doing this, our conceptions of the intermediate processes in the model will undoubtedly change, and radically. The only things that must not change are the correspondences between variables in the model and aspects of direct experience. They provide the anchor points in reality.

In perceptual control theory and hierarchical control theory, certain identifiable aspects of direct experience are labeled "perceptions," and they correspond one-to-one with specific signals in the model (or they would if the model were complete). This does not change direct experience. It does change what we think about direct experience. We are led to think of all descriminable aspects of the experienced world as "perceptions," not just as givens. The perceptual signals in the brain model are linked theoretically through physical properties of neurons to other signals, and eventually to variables in the physical model of the world. The physical model deals primarily with variables and relationships that do not correspond to perceptual signals: a world beyond the senses. As predicted, we do *not* experience electrons, light waves or quanta, force fields or energy. The physical world becomes directly experiencable only at contact points established by meter readings of various sorts. What we experience is a meter reading, not the physical process that gives rise to it. Processes intermediate to those contact points and the physical variables on the other side of the meter remain conjectural and unverified. Therefore, the two models together imply that what we perceive is not necessarily in direct correspondence to the entities and relationships in the world proposed in the physics model. If we choose to use both models, the viewpoint we must take is that the world of experience is derived from or dependent on another world that is not experienced, just as the surface of the ground that we can see and touch is held up by deeper layers of unknown composition that remain invisible and intangible.

This is the only viewpoint I can see that is consistent with physical models, neurological or biological models, functional models of the brain, and direct experience. What we experience is not a model. Everything we say about experience is.

Wayne Hershberger: Bruce says, "As we all know, 'boss reality' doesn't really sit still for its picture to be taken."

What!? You boggle my mind.

Your sentence implies what it denies: that is, although we cannot picture it, "we all know... 'boss reality.'"

Perhaps you meant to say that, although we can picture it, we cannot know boss reality. But, of course, such a transcendental reality as that smacks more of heaven than earth.

The relationship between a hierarchical control mechanism and its

environment is a much more mundane affair than picturing a transcendent reality. You imply in the following two paragraphs two different avenues of access:

"In an important sense, this environmental reality is hidden from the perceptual hierarchy. Its only access to it is proximal stimulation of intensity sensors."

"In an important sense, this environmental reality is not hidden from the perceptual hierarchy. Its model of it is presumed reasonably veridical because it in fact accomplishes perceptual control requiring feedback through the environment."

I am not sympathetic with the first point. To say that the environment is hidden by all the proximal stimuli is to paraphrase the fellow who claimed not to be able to see the forest for all the trees. (E.g., "Gee officer, I couldn't see the fireplug; my eyeballs got in the way.") Also, don't forget that the "intensity sensors" are spatially arrayed and sensitive to various forms of energy—over time. Further, transducers such as radar scopes vastly expand the range of our biological transducers. More trees to obscure our view?

However, I am favorably impressed with your second point, which is very similar to one I addressed last year—before you logged on to CSGnet. At that time, I observed that sensed efference affords a significant window to the world; that is, when an environmental variable is being controlled, sensed efference reflects the environmental disturbance (e.g., the weight of an object is proportional to the effort required to heft it). This principle provides a basis for the ideas of J. G. Taylor. Since that time, I have come across a delightfully lucid example from physics. Some physicists (Gerd K. Binnig & Heinrich Rohrer) won a share of a 1986 Nobel Prize by capitalizing on this principle in their design of the scanning tunneling microscope, STM.

> The STM operates by passing an ultrafine tungsten needle over the surface of a sample to be studied. A low voltage is applied to the needle, creating a tiny electric potential between the tip of the needle and the atoms on the surface. Although the needle and the sample never touch in the classic sense, quantum fluctuations enable electrons to "tunnel" through the intervening distance, hence the microscope's name.
>
> The current passing between surface and tip depends on the distance between them. A feedback mechanism continuously repositions the needle as it scans over the surface to maintain a constant voltage: the undulations of the needle are studied to re-

construct the sample's contours. (Scientific American, June 1990, p. 26)

Bill Powers: Wayne, you comment to Bruce: 'Perhaps you meant to say that, although we can picture it, we can not know boss reality. But, of course, such a transcendental reality as that smacks more of heaven than earth." Is that "of course" an argument against the proposition, or a bit of innuendo associating Bruce with a proposition that you *do* know how to refute? You must have a better reason than that for rejecting the possibility of a boss reality. Are you arguing against the uses of imagination?

Bruce Nevin: Wayne—sorry to butt in where angels fear, etc, I was supposing (out loud) why Bill might speak of aspects of reality being hidden. If Bill assumes the point of view of a perceptual-hierarchy model, and we assume a perspective supposedly outside of both that model and that which it is modeling, then we see that the only contact that a perceptual-hierarchy model has with "boss reality" is proximal stimulation of intensity sensors.

What might lie beyond that, accessible or potentially accessible (directly or in a further mediated, i.e., inferred, way) by way of proximal stimulation of intensity sensors might be reflected or imaged or modeled in the connections, input devices, and neural signals on up the hierarchy from those initial input devices and effectors.

Is the fidelity of that reflection or image or model verifiable? We postulate that coherent, successful behavior (however we define that) as an outcome of ongoing perceptual control constitutes a demonstration of fidelity. But the existence of conflict and reorganization must then be a demonstration of less than full fidelity. Since everything is connected to everything else, I suppose it might be argued that the "representation" immanent in the control hierarchy is complete—the universe in a grain of sand. But completeness in the same sense must be accorded the control hierarchy of a turkey.

All of which is only to say. there are grounds for assurance that the world of forces and impacts is there, but not for assurance that one knows everything going on in it. This is different from saying that some knowledge of it is in principle inaccessible. I know of no basis for either affirming or denying that.

As we all know, our pictures of "boss reality' are imperfect. (Our pictures: our snapshots, portraits, models, theories.) We know this by internal inconsistencies (conflict), and the very provision of means for revision (reorganization) in the model itself indicates that coevolutionary mutual adaptation is an aspect of that which we are modeling. A moving target indeed. I think this formulation is not ambiguous so as

to allow the pernicious interpretation entailing that we "know 'boss reality,'" an interpretation that I did not intend in the original formulation. It relies only on our own perceptions, and on the assumption that these reflect reality, etc., as above.

I wonder if it would be useful to consider Bateson's distinction between pleroma and creatura, the old Gnostic terminology by way of Jung, in place of the mind/body dichotomy that is the usual starting place for epistemology. From this perspective, the perceptual control hierarchy in a living control system is part of a continuum of cybernetic feedback loops extending throughout "boss reality." The fact that a control hierarchy is more strictly organized than other parts of this cybernetic soup is an important distinction as regards the control (behavioral) aspects of perception, but does not bear so strongly on the receptive (observational) aspects of perception.

A sense-intensity receptor is a difference detector, as I understand it. A perceptual signal is then news of a difference as it enters the control hierarchy by way of a receptor from some other part of pleroma, and as it passes up the hierarchy being transformed into other differences that make other differences in turn. The combining of signals to make a signal of a different type is unique to control hierarchies, I suppose.

Wayne Hershberger: Bruce, a cybernetic perspective is certainly very appropriate. In fact, that is exactly my point. The environment is an Integral component of cybernetic systems, and Bill's model is no exception. To speak of the environment as being outside Bill's model makes no cybernetic sense to me. It is OK sometimes to linguistically "zero" the environmental part of Bill's model (he certainly has greater proprietary claims on the internal hierarchical part) just so long as we don't forget that the loops are closed through an environment. Cybernetically, the environment is part of the epistemic system.

Bruce, you ask, "Is the fidelity of that reflection or image or model verifiable?" If I understand what you are saying, "verification" cannot possibly entail a demonstration of any correspondence between the "model" and what you are calling "boss reality" (there is no one to bring the boss). So, it seems to me that boss reality is really a gold brick: a charming fellow who is nowhere to be found just when you need him.

Bill, if I were to draw Venn diagrams, I think I would want to label them as shown below, meaning that the natural order immanent in the psychophysical flux is realized both perceptually and conceptually. Further, there is more natural order in the psychophysical flux than is currently dreamt of in our philosophies, meaning only that the subset boundaries are not fixed.

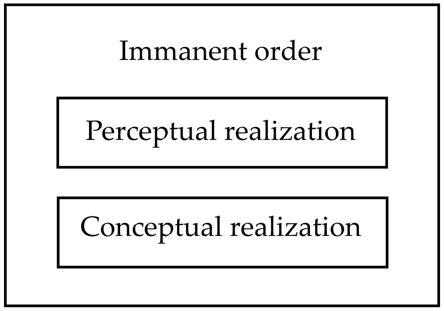

Beyond this, I am reluctant to go, because it seems that I would then be doing what I claim we should not be doing: confusing control theory with cosmology.

However, I admit that the expression "psychophysical flux" does reflect my control-theory perspective. When I think of "the psyche," I tend to think "reference values," and when I think of "the physical," I tend to think "disturbances." Each of these is an input to the canonical control loop, giving the loop psychical and physical poles. These poles are as inseparable as the poles of a magnet, making the canonical loop (incorporating the two inputs) a psychophysical whole.

The canonical control loop may be partitioned into separate arcs by a mechanism-environment interface, but the location of this interface is an accident of nature and does not separate matter from mind. The loop itself is *not* psychophysical, in the sense of comprising a mental arc plus a material arc separated by receptors and effectors.

That is what I think—I think.

Bill Powers: Wayne, Bruce has reiterated the basis in the control-system model for entertaining the concept of a boss reality. Your response basically says that if it's impossible to find the correspondence between the boss reality and perception, why bother with the concept?

You open that comment with: 'The environment is an integral component of cybernetic systems, and Bill's model is no exception. To speak of the environment as being outside Bill's model makes no cybernetic sense to me."

Your comment and Bruce's finally, maybe, perhaps, have joggled me into the right point of view for explaining my recalcitrance and possibly bringing our mysterious controversy to an end.

Yes, in my model there is always an environment and a behaving system. Neither makes sense without the other. I have always taken both into account. So follow me as I outline a chain of reasoning, and see if there is any point where you detect a weak link.

We're being modelers now. Imagine a sheet of paper on which we

draw two boxes, an Environment on the left and an Organism on the right. We don't need to model the environment; that has already been done better than we could do by physics, chemistry, and if you want to include raw meat, anatomy and neuroanatomy. We can put physical variables into that Environment together with all the laws that express relationships among them.

What we're trying to model is the organism part. So we draw two arrows: one from the environment to the organism, representing effects the environment has on the sensors of the organism, and one representing effects the output devices of the organism have on the environment. We are sitting up here with a good view of the paper, so we can see what is in the environment and what we're putting into the organism.

The challenge is to build a model of the organism so it will interact with the environment exactly as the real organism does. This means that basically we can give the model no help other than to provide it with the functions and interconnections that will, by their operation, generate some sort of behavior. When we guess wrong, we find that the functions and interconnections do *something,* but it bears no resemblance to real behavior. We just keep fiddling with the model until it behaves correctly. This leads us to a hierarchy of control systems, and so on.

If this model is to be complete, however, it has to reproduce not just behavior, but experience. In other words, the physical environment over on the left has to appear to this model just as it does to us. If we see intensities, the model has to see intensities. Simple receptors excited by various forms of physical energy will do for that. If we distinguish sensations in which different intensities are interchangeable, the model must do so. No problem: weighted sums seem to make sensation perceptions depend on physical variables as they should.

As we go higher, the problems become tougher, but we know what we're working toward. We want the model to contain signals representing configurations, transitions, events, and the rest, because we can see the world in such terms. We can't just tell the model about such things, of course; it has to contain the equipment that will, all by itself, derive such perceptions from its inputs. At the moment, we're pretty far from being able to do that, but we can at least draw boxes into the model showing where we will put the machinery for deriving the signals once we know what it is. As we know what the signals have to correspond to in our own experience, we can label them: "event perception," "relationship perception," "category," etc., corresponding to our subjective analyses of private experience. The model has to have those same private experiences. It has to have *all* the private experiences that we can discriminate into "natural kinds." That includes thought and reasoning.

If we now want to go far beyond where we are in the process of building this model, we may want to ask about epistemology. From our perch above this sheet of paper, we can see both the physical variables in the environment and the perceptual signals inside the organism model, the model of the person. It's perfectly clear that the perceptual signals are derived in systematic ways from energy fluxes connecting the physical variables to the sensors. As we fill in the boxes, we come to understand the details of that correspondence: just how an object in the environment, through the properties of light and optical devices, and through the photoneural receptors, comes to give rise to signals indicating its size, its distance, its shape, its orientation, and so on.

But now we come to the crux of the problem. We want to let the model figure out what there is external to it that corresponds to its perceptual signals. For example, the object it is looking at is actually a hologram, and all that actually exists in the environment is a set of wavefronts of light that don't actually originate at the surface of an object. How does the model go about checking into the reality of the object? We have no problem; we can see exactly what is going on. But how can the model figure it out without us to whisper in its ear? The model doesn't necessarily understand holograms (this has to be a model of any person).

One way is for the model to extend a limb to bring its visual image into the same region of visual space as the apparent object. If no contact is felt, the object could be considered intangible (that being what intangible means). But is it an intangible object in that position, or is there no object at all? Is this some kind of plasma object, or a less familiar trick of nature?

Solving this problem would clearly require a lot of sophistication and experience on the part of the model. It would have to compare what one set of sensors reports with what another set reports. It would have to form hypotheses and test them by performing appropriate acts. In the end, it would probably narrow the possibilities down to a small set, and on the basis of preference or niceness or some general principle, pick one of them as the answer.

Would it pick the same answer we would give from our omniscient point of view? Possibly, possibly not. In truth the model would have to know everything we know about the environment, and interpret its information exactly as we interpret it, and know what operations take place inside its own perceptual functions (which are not represented in the signals) to arrive at exactly the correct conclusion about what corresponds to any of its perceptual signals.

There is one thing we can be certain that this model can't do. It can't rise out of the plane of the paper and peer across at the environment

model to see what is going on there. We have given it no abilities that would allow it to see the environment except through the raw sensitivity to energy at its input sensors. The line separating it from the environment is a barrier that can be crossed only at the most primitive level, by physical energy.

So, for this model, as we have constructed it, *we* can know for certain how its perceptual signals correspond to what is happening in the environment model, but it can't know for certain. All it can do is entertain possibilities. One of those possibilities might be absolutely correct. But it can't know which one, if any.

So that is the epistemology of the model. Now what about our own?

If this is indeed a model of a human being, if we've got everything right, then it is a model of the observer, of ourselves. It is a model of us sitting up here and looking down at a sheet of paper on which there are diagrams of an environment and of a nervous system. The model has eyes and limbs; they are models of our eyes and limbs. The model has sensors and neural signals which are supposed to represent our own sensors and neural signals. The model, if it were looking at a sheet of paper with diagrams on it, would know of those things only in the form of neural signals inside itself. As the model can't rise out of the plane of the paper to see what is really in the other diagram, the diagram of the environment, so we can't rise in a fourth dimension out of our brains, to peer at whatever it is that is causing our neural signals. As the model can't sense the internal workings of its perceptual functions and use that information to deduce what is causing any given perception, so we can't deduce the transformations that lie between the environment and our perceptions.

The model might conclude correctly that it doesn't have access to an authoritative picture of the environment model; it could reach this conclusion simply by noticing that several plausible alternative interpretations exist. On that basis, it might decide that there is no point in guessing about a boss diagram that it now realizes it can never experience directly. It might decide that all it can do is compare one perception with another, and take that as the beginning and end of reality. The boss diagram is an unnecessary frill, a religious superstition; it is to laugh.

Of course we, sitting up here, would laugh at that, knowing what a mistake it is. There really is a diagram of the environment there, and it really does have a particular state, and the model hasn't been so far off the track as to be completely hopeless. At least it could survive in its interactions with the environment on the basis of what it has deduced. What it thinks it is controlling is at least equivalent, in the necessary ways, to what it is actually controlling. It might have omitted a conformal transformation or two here or there, but because it omits the same transformations from perception of its own actions, the two mistakes cancel for all practical purposes. And if it gives up now, assuming that all there is to be known exists already in the perceptual world it has constructed for itself, it's going to miss most of the fun.

And what of us? We sit up here, experiencing our own perceptions, and debating whether or not they are connected to a physical world, and if so, what kind of physical world. If we believe what the model of the person seems to imply, then we are in the same fix it is in: we experience our perceptual signals, but there is nobody sitting in a higher place still who can tell us what the environment diagram really looks like. We have to figure it out on our own, each in an individual private world.

So that's where my epistemology comes from. It comes from trying to think of a model that behaves and experiences like a person, and is built the way a person is built with sensors and a nervous system and effectors. The final step, to my personal epistemology, is simply an application of the model to myself. The model contains my best understanding of how the nervous system on the right, and the environment on the left, work and interact with each other. If I now don this model and imagine that I am experiencing the world from inside it, I transform my understanding of the physical world that seems to surround me. I realize that a very plausible thing to say about it would be: it's all perception.

But it is not implausible to add "... of something else."

Wayne Hershberger: Bill, a thought has occurred to me that you might agree would prove helpful. It seems to me that the little man model might be said to know where the target is, while knowing nothing either of computers or of yourself, his creator. That is, it seems to me that if you were to rewrite your essay, putting the model in a computer where it can function as a simulation, the epistemological implications might appear more clear-cut.

Bill Powers: Wayne says, "... if you were to rewrite your essay, putting the model in a computer where it can function as a simulation, the epistemological implications might appear more clear-cut." Actually, that is the route I took to my present position, only it came from building real systems more than from simulations. In the late 1950s, for example, Bob Clark and I built an "isodose tracer" that used an analogue computer as a control system to make a tiny radiation probe move along curves of constant radiation intensity in the beam of a Cobalt-60 treatment machine (VA Research Hospital in Chicago). In the early stages, we got some strange curves, because the long stem that held the probe turned out to be radiation-sensitive. The control system was

keeping what it assumed to be sensed radiation at the probe tip constant, but it couldn't know where that radiation was being detected. The variable under control wasn't quite the one that was supposed to be under control.

I've mentioned the voltmeter effect before: the reading on the voltmeter is not the "true" voltage because the meter draws current. In my electronics ventures with radiation probes and photosensitive equipment, often incorporated into control systems, it was almost always necessary to correct the meter readings when measuring low-current high-voltage sources. Automatic control of such voltages required compensation so that the "real" voltage, not the measured voltage, was controlled.

The whole world of electronics is fraught with examples. A simple circuit board is, to the electroniker, largely imaginary. The surface appearance of the board has almost nothing to do with what is "really" going on. Every component carries in it mysterious properties like resistance, capacitance, inductance, and amplification that are never experienced directly. (Voltage is one example, but it doesn't feel like voltage. It feels like hell.) Usually, such things are known only after calculations based on the few contact points with direct experience. Yet when you assume that such things exist in some boss reality, as you must in order to make any sense of "correcting a meter reading," the result is the power to make things happen in highly predictable ways. You adjust a tuned circuit a little below resonance, so it will be exactly at resonance when you remove the capacitance of the probe you're using to measure the response. The true operation of a circuit is what you deduce would take place if you weren't measuring anything!

You've picked this diagram from the possibilities I suggested:

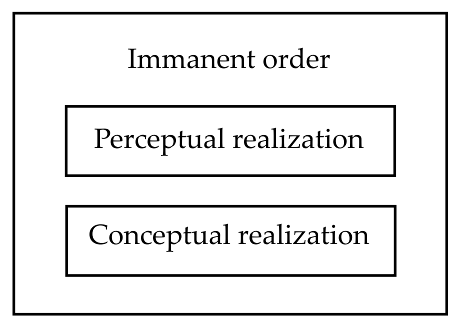

"Immanent order" wouldn't be a bad term for "boss reality." From my viewpoint, it has the nice implication that there can be order without our knowing what it is. But I have some more questions.

By choosing the diagram in which the immanent order extends beyond the boundaries of the realizations, you have agreed with me that there is more to know than meets the eye. By making the two realizations independent and non-overlapping, you have said that each has its own relationship to the immanent order independently of the other.

In this diagram, there's no connection shown between perception and conception, nor any indication of how these "realizations" might relate differently to the immanent order. You describe the figure as a Venn diagram. This implies that within the outer boundary there is some immanent order, and that it's simply marked off into regions, with the elements of the largest field being no different inside and outside the two "realizations."

As you didn't specify the difference inside and outside the realizations, there are two possibilities:

A realization is simply a noticing of something that was always there, the noticing in no way altering what was always there but merely bringing it into the field of attention.

A realization is some transformation or projection of the immanent order, so that the realization is an invention or at least an expression of the nature of the system becoming acquainted with the immanent order.

In both cases, there is an implicit relationship between a realization and the immanent order. In the first case, the realization is completely passive; it is merely recognition. In the second case, there is a difference between the realized and unrealized states of portions of the immanent order. Does either of these choices fit your conception?

I take it that the rationale for the term "immanent order" is that neither perception nor conception is random; that both reflect some orderliness that constrains them. Does this not imply some effect of the immanent order on the realizations?

Wayne Hershberger: Bill, you said: "Yes, in my model there is always an environment and a behaving system. Neither makes sense without the other. I have always taken both into account. So follow me as I outline a chain of reasoning, and see if there is any point where you detect a weak link."

That is like waving a red flag in front of a bull. How can I refuse? Here goes.

Although you said, as quoted above, that the environment is in your model, you soon spoke as if it were external to it: "But now we come to the crux of the problem. We want to let the model figure out what there is external to it that corresponds to its perceptual signals. For example, the object it is looking at is actually a hologram, and all that actually exists in the environment is a set of wavefronts of light that don't actually originate at the surface of an object."

The basic (perhaps only) problem I see is your confusing separate issues. In your elaborate analogy, there is one dichotomy (creator-created) and one dyad (organism-environment). You confuse the two, shifting from one to the other as though they were one and the same. As I have said before, this confuses physiology with metaphysics.

For example, you say that "... we can't rise in a fourth dimension out of our brains, to peer at whatever it is that is causing our neural signals." The first part of this remark, mentioning a fourth dimension, is alluding to the inability of the created dyad to assume the epistemic perspective of the creator (epistemology), whereas the last part is concerned with the relationship between the two parts of the created dyad (sensory physiology). Apples and Oranges.

Then you reverse course and switch from sensory physiology back to epistemology with the following remark, effectively by substituting the word "perceptions" for the words "neural signals" (moral: perceptions are not to be equated with perceptual signals).

"As the model can't sense the internal workings of its perceptual functions and use that information to deduce what is causing any given perception, so we can't deduce the transformations that lie between the environment and our perceptions." If a neural model could monitor its role in the perceptual process, could it deduce the nature of the transformations that lie between the neural model's signals and the neural model's environment?

You also say: "'Immanent order° wouldn't be a bad term for 'boss reality.' From my viewpoint, it has the nice implication that there can be order without our knowing what it is." Yes, exactly. It seems to me that the expression "immanent order" (or natural order, or what have you) would be a much better term for your purposes than "boss reality," for precisely the reason you mention. The word reality connotes a verifiability you are denying to "boss reality," making the expression an oxymoron.

And: "By choosing the diagram in which the immanent order extends beyond the boundaries of the realizations, you have agreed with me that there is more to know than meets the eye." No. I think there is a point of agreement here, but not for the reasons you say. The aspect of the diagram that implies that there are some things which can be known but which do not directly meet the eye or the ear or the other sense organs (e.g., your example of voltage) are the elements which are both *in* the set labeled conceptual realizations and *not in* the set labeled perceptual realizations. In contrast, the elements that are in neither subset (neither type of realization) simply imply an immanent order which is not realized—either perceptually or conceptually. Whether this unrealized order is potentially realizable is something a static Venn diagram doesn't capture. But if one takes the view that at least some of the immanent order unrealized at present might be realized in the future, it is presumptuous to suppose that this realization *cannot* be perceptual. Further, any immanent order which cannot possibly be realized at any time in either way is simply not to be known; it does not mean that there is more to know than can be known. I readily admit that there can be more than what-can-be-known, but I cannot agree that there is more to be known than what can be known, without contradicting myself. Nor can you. We are talking here about the limits of the epistemic process, not the limits of a man—obviously, there is more to be known than any one man will ever know.

"By making the two realizations independent and non-overlapping, you have said that each has its own relationship to the immanent order independently of the other." I would say that one is not a subset of the other, but their intersection is not nil, meaning that the two realizations are independent of each other. Your drawings did not seem to include this alternative, so I selected the one which I thought would "suggest" independence (actually non-overlapping subsets depict mutual exclusion, a form of dependence).

"As you didn't specify the difference inside and outside the realizations, there are two possibilities:

1. A realization is simply a noticing of something that was always there, the noticing in no way altering what was always there but merely bringing it into the field of attention.

2. A realization is some transformation or projection of the immanent order, so that the realization is an invention or at least an expression of the nature of the system becoming acquainted with the immanent order."

Your two alternatives are not a matched set. They are not necessarily mutually exclusive.

First the latter: The expression "system becoming acquainted with the immanent order" seems to me to suggest that the "system" transcends (stands apart from) the immanent order. That would insinuate a gratuitous wild card. For me, the system becoming acquainted with the immanent order must be part and parcel of the immanent order. Perhaps they are even coextensive. The system responsible for the two types of realizations is best characterized as an ecological system (i.e., an organism-environment dipole). If we attribute the "becoming acquainted" merely with the organism pole and the immanent order merely with the environment pole, we are being arbitrarily inconsistent. Therefore, if one is to be consistent, it seems to me that the realizations would inevitably be "an expression of the nature of the system becoming acquainted with the immanent order," because the system (the ecological dipole) becoming acquainted with the immanent order is part and parcel of the immanent order.

Now the former: Does "self-acquaintance" rule out the possibility that acquaintance is simply a registration of what is "there"? Fortunately, the question appears to be academic. If some aspects of the immanent order are hidden by the recursiveness of self-acquaintance, or whatever, I would say, so what? Call it Noumenon, and let the faithful worry about it, because, by definition, it is not to be known.

Finally: "I take it that the rationale for the term 'immanent order' is that neither perception nor conception is random; that both reflect some orderliness that constrains them. Does this not imply some effect of the immanent order on the realizations?" The relationship is not cause-effect, but *yes*, realizations of both types reflect some orderliness that constrains them.

Bill Powers: Wayne, we keep going around and around on the same points without getting anywhere. You keep saying that I am missing the distinction between modeling and metaphysics, and I keep saying that metaphysics is just one of the things a brain can do. Let's take it from the top.

You say, "It seems to me that the expression 'immanent order' (or natural order, or what have you) would be a much better term for your purposes than 'boss reality,' for precisely the reason you mention. The word reality connotes a verifiability you are denying to 'boss reality,' making the expression an oxymoron." So in your book, "reality" is identical with "verifiable reality." It's not, in mine. I don't need to understand electricity to comprehend that touching certain objects is highly unpleasant. I can generate acts like touching objects, but I can't decide what their consequences will be. That is decided for me by something I don't sense and only partially conceptualize. I can choose whether to repeat a consequence or to avoid it, but I can't make an act have a different consequence. In that department, something else is boss.

Are you saying that I must realize in perception or conception the connection between an act and its consequence, and *verify* the nature of that connection, before I can accept that there really is a connection? Or are you saying that it is sufficient to verify only that the consequence reliably follows the act, and never mind why? I would argue against the latter as being simply pre-Galilean empiricism, and reject it because it works so poorly in comparison to the method of modeling. The method of modeling posits an unseen reality mediating between act and consequence, and has most profitably interpreted nature in those terms. The assumption has repeatedly been vindicated. How could the purely empirical approach ever predict a new perception, and experimentally reveal the link explaining the surface appearance of a causal sequence?

Later in your post, you say, "... if one is to be consistent, it seems to me that the realizations would inevitably be 'an expression of the nature of the system becoming acquainted with the immanent order,' because the system (the ecological dipole) becoming acquainted with the immanent order *is* part and parcel of the immanent order." This would be consistent. It would also be an empty generalization, a true statement of which one can legitimately ask, "so what?" To say that all of knowledge is an expression of the immanent order (whatever that is) is meaningless: any statement that is true of everything is trivial. Even that statement and my response to it are part of the immanent order. I repeat: so what? Knowing that does not contribute to our understanding of any specific phenomenon—in fact, it seems to discourage asking questions and conjecturing. All of our useful understanding comes from discriminating one part of the immanent order from other parts, and from realizing that different parts of it have characteristics of their own unlike the characteristics of other parts. It is out of these differentiations that all knowledge comes. From these differentiations, we come to realize that organisms and environments are *not* alike. We realize that some parts of organisms function differently from other parts. We realize that brains exist.

And ultimately we are faced with a paradox, the one you and I have been arguing about. We find by experimentation that the presence of certain signals in a brain is the sine qua non of perception. Remove those signals and you destroy, as far as the victim is concerned, a chunk of the immanent order. Yet you don't destroy it for anyone else. What other conclusion can we reach but that perception is absolutely contingent on those signals? That puts us, as perceiving entities, inside the brain. To deny that would be to destroy the whole structure of perceptual and conceptual organization we have so painfully built up. That structure is at least as well worked out as any metaphysical argument in words, and a lot better tested experimentally.

I don't see that any philosophical conception, any combination of words, any exercise of pure reason, can be more persuasive than these simple observations. By simple and straightforward reasoning based on close attention to experiment and observation, we are led to conclude that the object of perception and thought is a world existing inside, not outside, a brain. We can see how this world of experience is related to what we conjecture to exist in a physical environment outside of us, but we can also see that the relationship is not a simple or direct one, nor is it wholly verifiable because of our peculiar circumstance of being inside the very system we model and by necessity having to perceive and think using its equipment.

Until you can come up with an equally persuasive set of observations and deductions that lead to a different conclusion, I will continue

to be satisfied with my view of the relationship between consciousness and reality. Simply reiterating your point of view without revealing and justifying each step of the way that leads to it will not win me over. I understand that if I believed as you do, all would be explained. But I do not.

Rick Marken: I have gotten behind on this epistemology debate (or maybe I just don't understand it). Could Bill or Wayne give me a short (like two-sentence) description of what is being debated? I am wondering if Wayne is arguing that there is no physical environment, or that the physical environment is an unwarranted assumption, or what?

Wayne Hershberger: Bill, we are back where we began. You say: "And ultimately we are faced with a paradox, the one you and I have been arguing about. We find by experimentation that the presence of certain signals in a brain is the sine qua non of perception. Remove those signals and you destroy, as far as the victim is concerned, a chunk of the immanent order. Yet you don't destroy it for anyone else. What other conclusion can we reach but that perception is absolutely contingent on those signals? That puts us, as perceiving entities, inside the brain. To deny that would be to destroy the whole structure of perceptual and conceptual organization we have so painfully built up. That structure is at least as well worked out as any metaphysical argument in words, and a lot better tested experimentally." And I say that ablation (to which you refer above: "Remove those signals..."), the technique pioneered in the 17th century to localize mental functions, identifies certain *necessary* components of the various functions we call mental (e.g., vision). It does *not* identify the *necessary and sufficient* components. Without the photon, there is no vision—for anyone. Ablate photons and we are all blind. This means that the *proprietary* aspect of our respective experience (my perceptions versus your perceptions) are contingent upon our respective brains. That is the argument you are making, right? But that does *not* put us in our respective brains! Our feet are too big.

So, you are right, this is where we came in. Perhaps it is time to take a different tack. Let me reciprocate by asking you what, if anything, is wrong with the following remark (using your terminology): Bill Powers' hierarchical-control-theory model models an aspect of boss reality (the organism aspect), the other aspect (the environment aspect) already having been well-modeled by contemporary physics.

Bill Powers: Wayne, what are these mythical "photons" of which you speak? I don't know anyone who has the ability to "ablate photons." We can perform various acts, like shutting our eyes or pulling the chain on a light, that result in loss of vision, but to attribute that loss of vision to a loss of "photons" goes far beyond anything that is observable.

By accepting "photons" as necessary precursors to vision, you are leaping ahead to the conclusion you want to reach; namely, that photons actually exist just as we imagine them to exist. But I can't accept that mode of argument: I want to know the operational basis for every critical entity you use in your proofs. I will accept that turning off the lights results in loss of vision. Those are both observables, perceptions. I do not accept that you have shown photons either to exist or to have anything to do with this phenomenon—not until you tell me your basis for knowing that

"This means that the *proprietary* aspect of our respective experience (my perceptions versus your perceptions) are contingent upon our respective brains. That is the argument you are making, right? But that does *not* put us *in* our respective brains! Our feet are too big." Cute comment, but irrelevant. Our perceptions that we call "feet" are certainly not too big to fit into a brain: they are precisely small enough to pass through a neural fiber. All aspects of our perceptions are proprietary, including our convictions that some are not. If that were not true you would have convinced me by now. But you have nothing objective to show me to help make your case.

You ask, "... what, if anything, is wrong with the following remark (using your terminology): Bill Powers' hierarchical-control-theory model models an aspect of boss reality (the organism aspect), the other aspect (the environment aspect) already having been well-modeled by contemporary physics." Sensing a bear-trap, I answer cautiously. Both the hierarchical-control-theory model and the physics model purport to represent aspects of a boss reality. Both are tested (by a person, using a brain and body) by assuming the model to be correct, and predicting the effects of actions on this boss reality that have consequences we can perceive. It is the boss reality that determines whether our predictions work out as we expect, or whether different consequences occur. If the consequences are different, we modify our imagined pictures of the boss reality in a direction that promises to lessen the difference. This process converges to some minimum-error condition where we declare ourselves satisfied with the models. According to both the physics model and the hierarchical-control-theory model, this process of acting and testing takes place inside a brain. It is not necessary to assume that the model in the brain has any particular correspondence to the boss reality. It is necessary only to assume that whatever that correspondence might be, it is stable over time.

I ask a similar question of you: is it fair to say that you believe (a) that there are non-proprietary aspects of our respective experiences and

(b) that we can say unequivocally what they are?

Avery Andrews: The first wisdom of linguistics is that speakers are always wrong when they try to explain why they say what when (the stories are pathetic, and tend to fail within 30 seconds). Getting behind the descriptions to the explanations will be reverse engineering all the way (I regard current linguistics as being essentially descriptive, in spite of the presence of a lot of talk about explanation).

Bill Powers: Avery says that "... getting behind the descriptions to the explanations will be reverse engineering all the way..." Beautiful. Precisely.

Gary Cziko: I would appreciate Avery and/or Bill giving me a description and example of "reverse engineering." I have a hunch that *all* science and *all* nontrivial engineering (i.e., finding engineering solutions to new problems) is in fact reverse engineering, but I want to know more about what this term means before making this claim.

Avery Andrews: Gary: All I meant by "reverse engineering" is that there is no quick substitute for figuring out how it works on the basis of analyzing what it does.

Bill Powers: Gary, reverse engineering is a term from (I believe) the semiconductor industry. It refers to duplicating the function of someone else's integrated circuit. What with copyrights and patent laws, modern reverse engineering gets pretty complex. One team analyzes the function of the competitor's chip and prepares a specification stating the relationships between inputs and outputs (and other aspects of visible behavior) that the "unknown" chip creates. This specification is then passed on to a design team which is never given access to the chip itself, only to the specification. The design team is never allowed to communicate directly with the analysis team. From the specification alone, the design team generates a completely new chip design, from scratch, that will accomplish exactly the specified functions. I'm sure there has to be some cheating—the design team has to know that the specs describe a computer, for example, and not a sewing machine.

At any rate, the result is a new chip that can be plugged into the same socket that the original chip occupies and works exactly the same way, down to the last detail of functioning. This is the ultimate in the method of modeling.

In fact, the final chip might not accomplish the functions in exactly the same way the original did. Sometimes the new chip proves to perform some functions more efficiently than the original—in fewer steps, or faster. Presumably, if those aspects of functioning had been part of the spec, the design team could have deliberately slowed some circuit operations and matched the slowness of the original too! But the design team, prior to releasing its product, never can know whether it has accomplished the functions in the same detailed way that the original does. In the final comparison, it is often found that some functions were reinvented exactly as in the original, while others do the same things in a different way. That is what is hoped for—what avoids a suit for patent infringement.

This is basically what I am arguing with Wayne Hershberger about. We are trying to reverse-engineer evolution (or whomever you want to blame). In doing so, we come up with a model of underlying design features constituting a system that interacts with its environment just as real organisms do. Of course, in doing this, we try to reproduce only those functions we understand, and we ignore many others, such as skin color, weight, exact lengths of appendages, and so on through a long list of "unimportant" parameters. As initial models succeed, we bring in more detailed parameters to match, even to the level of neural functions in a few cases.

But we can never know that we have accomplished something in the same way that an organism accomplishes it, in every detail. For that matter, we have no reason to think that every organism of a given species accomplishes its functions in the same way as other organisms of the same species. Judging from the very large differences in brain anatomy that exist from one person to another, in fact, its unlikely that all people are internally organized in the same way even if they behave in roughly the same way. The brain is plastic and its organization is influenced by the experiences of a single lifetime. Our reverse engineering is fundamentally limited by this fact: no one model can ever reproduce to the last detail the inner functioning of all examples of any kind of higher organism, because the originals are not all designed in exactly the same way. We will always be limited to modeling the "general idea" behind an organism, because that is the limit of consistency in the originals. The method of modeling is primarily a method of understanding individuals, and only secondarily a way of saying general things about all individuals. Models must always contain parameters that can be adjusted to fit the "general idea" to a specific organism.

This, naturally, has some serious implications concerning the nature of scientific research into human nature. It's usually assumed that one is dealing with a standard instance of Homo sapiens—the very idea of assigning such a term to the whole human race is to assert that fundamentally we are all the same. In the psychology lab, great attention has been paid to using a standard animal model—the Sprague-Dawley rat, during my formative years. If you have a standard rat or a standard

person, you should get standard responses to standard stimuli. If any human being is as good an example of *Homo sapiens* as any other, you can study groups of people as interchangeable units, drawing generalizations from the data which you assume to be measures of common underlying properties fuzzed out by uncontrolled stimuli.

But what if, below some level of observation, there are no common underlying properties? Then the whole rationale of statistical studies of populations collapses. The specification team can't come up with a spec that fits all instances of the chip that is to be reverse-engineered. All they can describe, for each parameter, is the average spec. As Russell Ackoff said in a lecture that Dag Forssell has transcribed, there's no way to design the optimum human being by combining the optimum spec for each function making up the person. This would be like trying to build a perfect car by using the engine of a Rolls-Royce, the suspension of a Ferrari, the body of a Chevette, the carburetor of a Chevrolet, and so on. The functions all have to work together in a single person; the final workable form of each function depends on the final forms of all the other functions. Each part of a person is adapted to all the other parts of the same person, not to the same parts as they are manifested in other individuals. And the process of mutual interadaptation never ceases.

I use the term "generative model" as Humberto Maturana defined it (perhaps following someone else). A generative model is one that will reproduce the phenomenon of interest by operating strictly from the interplay of its own properties. A generative model of control behavior is a control system with an input function, a comparator, and an output function, in an environment that links output to input in a specific way. There is no component in a control-system model that "controls." Control is the result of operation of a system with these functions in it, connected as specified by the control-system model, and operating as dictated by the input-output properties of each component.

So, given inputs, constraints, and parameters, a generative model must always produce some kind of behavior. We can't necessarily anticipate what such a model will do, but whatever it does is rigidly set by the properties we have given it, and by the surroundings with which it interacts. We hope that the behavior of the model will resemble the phenomenon we're trying to explain. If it doesn't (and few models do, the first time they are set in motion), we have to modify the model. That's how models grow and improve.

Wayne Hershberger: Bill, you said, "Our perceptions that we call 'feet' are certainly not too big to fit into a brain: they are precisely small enough to pass through a neural fiber." Neural signals in the brain might be said to be relatively small, but the replicable perceptions (phenomenal objects) those signals help mediate are not necessarily small. Smallness is an aspect of phenomena, and it is a mistake to suppose that the size of a phenomenal object is in any way related to the size of the neural signals which help mediate it. You have yourself been championing this sort of argument in many of your recent posts.

You also said, "All aspects of our perceptions are proprietary, including our convictions that some are not." No. A proprietary aspect is immanent in all experience, or so it seems. But this does not imply that there are no other aspects.

"This is basically what I am arguing with Wayne Hershberger about. We are trying to reverse-engineer evolution (or whomever you want to blame). In doing so, we come up with a model of underlying design features constituting a system that interacts with its environment just as real organisms do." Yes. As I see it, we are trying to reverse-engineer the phenomenal domain, and the "spec" that I think is of the first importance in this venture (also, as I think Kant was saying) is that phenomena are bipolar: in a word, psychophysical. Control theory appears to be uniquely compatible with this psychophysical specification, providing one continually recognizes both ends of the dipole—a control system *and* its environment. Perhaps we should change our language habits and speak of control subsystem, since the control system is only one part (or pole) of the system being captured by our reverse engineering.

Bill Powers: Wayne says, "Neural signals in the brain might be said to be relatively small, but the replicable perceptions (phenomenal objects) those signals help mediate are not necessarily small. Smallness is an aspect of phenomena, and it is a mistake to suppose that the size of a phenomenal object is in any way related to the size of the neural signals which help mediate it." But you assume, in order to say this, that phenomenal objects and attributes of objects are something other than neural signals. I assume they are the same thing. How do we get past that?

"As I see it, we are trying to reverse-engineer the phenomenal domain, and the 'spec' that I think is of the first importance in this venture (also, as I think Kant was saying) is that phenomena are bipolar: in a word, psychophysical." Why do you assume the "-physical" part of psychophysical? There is nothing in the physical domain that is not derived from perception and thoughts about perceptions. It seems to me that you slip your conclusion into your premises. I do not see the "psychological" aspect of experience as being on an equal footing with the "physical" part. The physical part is a set of ideas, and so is a subset of the psychological part.

I find the topology of your point of view baffling. It seems to involve

some magical way of knowing things without perceiving them, and some way of checking on the meanings of perceptions other than comparing them with other perceptions. I can't grasp the role that you give to perceptual signals, or for that matter, to the brain. I can't understand what position you're assigning to the Observer—if the observer isn't in the brain, where is it? And where, then, are the objects of observation?

Joel Judd: I'm starting to lose track of what is being claimed as individual responsibility and what's being foisted on the environment. Don't we all agree that the "world" is constituted in our perceptions?

Martin Taylor: Joel, I can't speak for Bill, but in his discussions with Wayne, I think I agree with him. I assume that there exists something outside ourselves, but it can be known only through our perceptions. Our perceptions can be constructed only through the feedback of our actions to our sensors, but we can develop internal things (which I call structures to avoid words like "simulated worlds" or "world models" or "imagined worlds") that enable us to perform as if there were certain objects and relationships in the (unknowable) world and not get into too much trouble by doing so.

Bill Powers: Mental representations, in hierarchical control theory, are identically neural signals arising from sensory receptors. Each level of signals enters a higher level of perceptual functions (neural computers), many functions acting in parallel, which re-represent subsets of the incoming signals as a new level of mental representations. There are 11 such levels in my model, covering (as far as I could) all phenomena of perception, all aspects of the experienced world, inner and outer, concrete and abstract. I refer to the mental-representation signals at all levels as "perceptions," rather than using different terms for low-level and high-level representations.

Comparison implies two things to be compared. In the hierarchical control-theory model, one of them is a mental representation, a perceptual signal, indicating the current actual state of the perceived world, or one aspect of it. The second is also a mental representation, a signal, but it represents the state of the same aspect of the perceived world as it is intended to be perceived. This is the reference signal. A comparator is simply a device that receives these two signals and emits an "error" signal indicating the difference between the two inputs to the comparator. A less pejorative term is "deviation." An error signal does not indicate a mistake. It simply indicates by how much and in what direction the current perception deviates from the current setting of the reference signal. That indication drives the corrective actions of the control system.

All that the organism can know about the environment exists in the form of mental representations, perceptual signals. The organism can't know the actual states of its physical inputs (although an intelligent enough organism can certainly make models of the external environment, and thus provide itself with a highly plausible story about what they are). When I say "an organism," I mean every human being, as well as our coevals of other species. The environment that is directly experienced by a human organism is confined to the set of all perceptual signals (although they are not all consciously experienced at once). Wayne Hershberger disagrees with me. But I agree with me.

Evaluation of behavioral-path consequences can be done through the imagination connection. A system of higher level normally acts by sending reference signals to lower-level systems. Those reference signals specify the states to which individual lower-level systems are to bring the kind of perception that each controls. Copies of the resulting perceptual signals become inputs to the perceptual function in the controlling higher-level system. When lower-level control succeeds, as it usually does, the result is that each lower-level system sends upward a perceptual signal that matches the reference signal it is receiving from the output of the higher system.

Exactly the same effect can be achieved if the higher system sends its output not to the comparator of the lower system, but back into its own perceptual function. It is just as though the lower system had succeeded perfectly and instantly. This is what I call the imagination connection. With this connection in effect, the higher system can quickly go through possible outputs (I assume a level where complex logical processes are occurring) and judge their effects on the controlled variable. Thus selection of lower-level actions (and their perceptual consequences) can be done without actually producing any actions.

This process of mental planning is undoubtedly more complex than I make it here. Modeling must be involved, in the imagination path, because the properties of the outside world (which includes all lower-level control systems) must be taken into account. But the basic picture of how imagination works seems to explain the broad outlines of planning of all kinds—not just behavior-path planning.

Behavior always follows some path. The question is whether the paths are in fact always planned, or whether they are simply the result of the way a control system gets from a state of error to a state of no error. Planning of behavior paths is not necessary in all cases—in fact, it is necessary in very few cases. To see whether a path is planned, one can introduce disturbances and see if their effects on the path are corrected, or if the organism simply accepts the deviated path and reaches the goal anyway. The latter is probably the more likely outcome. Paths would be planned in advance only when they make a difference to the

organism. Control systems do not have to precalulate behavior paths.

You can say that an apple is redder than an orange, or cheaper, or better-tasting. But the control process is separate from that comparison, which is really a judgment of relationship. Given the perceptual comparison, you must still specify what the goal is: are you going to paint the orange to make it as red as the apple, or is the difference in redness OK with you? Are you going to raise the price on apples, or inject something in the orange to make it taste better? The goal has to be stated if control is to be involved. And then the comparator—an element of the model—must take the perceived relationship between apple and orange, compare it against the desired relationship, and judge it as being not sufficient, just right, or overdone—relative to the preferred state.

We determine what the goal will be; the environment doesn't. The environment may provide us with a selection of experiences from which to pick feasible goals, but it doesn't do the picking. The environment determines what we must do in order to have the desired effect on experience. If we can't do what it requires, or if the desired result is impossible, then we fail to control.

Neurologists tell us that human beings are basically a set of neural connections. Biochemists tell us that behavior is controlled by interactions among molecules. Sociobiologists tell us that it is genetic fitness to reproduce that determines how we shall act. Physicists tell us that thermodynamics and quantum uncertainty are the keys. Radical behaviorists tell us that schedules of reinforcement are what do the trick. Personality psychologists tell us that traits and attitudes and feelings and aspirations account for behavior. Sociologists tell us that the individual is simply an expression of the society. Existentialists tell us that individual being is at the core of it all.

Doesn't this strike you as a bit suspicious? All these answers, and they all show that the particular interests of the explainer just happen to contain the correct solution to it all. But when you ask any of these explainers how their explanations work, you run into a blank stare. The explanations *are* how it works. They don't ask what lies beneath the explanation. They don't try to link their own field of study to the fields of study of others. It's all extremely provincial and, aside from the specialized expertise involved, superficial.

Control theory crosses all these boundaries because it is concerned with the how of behavior more than the what. It has nothing specifically to do with society, or even with any particular individual behavior. All examples of behavior, all aspects of behavior in any discipline, are grist for its mill. The world it addresses is larger than that of any existing discipline.

Wayne Hershberger I have a nagging itch demanding to be scratched. Bill thinks I wants physics to be part of the immanent order. No. I would say that physics is a science, involving conceptual modeling, as I imagine Bill might say. I would say that there is order immanent in the phenomenal domain that is modeled by physics. I use Bill's word *model* to refer to the intellectual achievements of physicists. That is, I use the word model to denote something human-made. Unfortunately, the word model has another, unintended, connotation: a replica of an original. Like Linus Pauling, I do not regard scientific models as being replicas of divine (Noumenal) originals. Theoretical physics does not involve "reading God's mind." I view Einstein's saying that it did as a metaphor.

Bill says, "... you assume... that phenomenal objects and attributes of objects are something other than neural signals. I assume they are the same thing. How do we get past that?" As I see it, the issue is a difference between what your theory assumes, and what you say your theory assumes (or implies). I seriously doubt that your hierarchical control theory necessarily implies (or assumes) that phenomenal objects are neural signals. In claiming that your theory is not solipsistic, I find myself in the paradoxical position of arguing that your theory is better than you say it is. That is a sort of disagreement, but one that I think belies a fundamental agreement.

Let me say some things about phenomenal objects, because such descriptions comprise the specifications which we are attempting to reverse-engineer. Please understand that what I say is not presented as an alternative to your theoretical model. What I am trying to do is describe some of the specs that all our psychological models must be able to realize.

Phenomenal objects are simply the particulars of experience. They are the constituents of the empirical world that we are wont to call things. The layman calls them objects or physical objects, and supposes that their substance is essentially material. In contrast, philosophers such as Bishop Berkeley called them perceptions and supposed that their substance is essentially mental.

It seems to me that arguing whether phenomena are substantially mental or material is much the same as arguing whether a magnet is essentially a north or a south pole. The argument makes no sense to me, because phenomena, like magnets, appear to be bipolar, with each instance involving an observer-observed (knower-known) dipole. A dipole, not a dichotomy. For instance, the visible surface of every phenomenal object in my study is the one facing that ubiquitous phenomenal object I have learned to call myself. Inasmuch as this personal "perspective" inheres in every phenomenal object, there is more of me to be found in the phenomenal world than is to be found in the phenomenal object I call myself.

This widely distributed aspect of myself which permeates the phenomenal world lends a proprietary aspect to the phenomenal world, making it mine, as it were. That is, the phenomenal world presents itself as a personal "perspective" with that unique point of view being tied to the phenomenal object I call myself. (I put the term "perspective" in quotation marks to signify the observer-observed relationship noted above: a dipole, not a dichotomy.)

Locating oneself is an empirical matter, and does not involve merely locating one's brain, as Dennett, for one, has nicely illustrated in his delightfully humorous essay, "Where am I?" Locating oneself involves a determination of the spatial relationship obtaining between what might be called the sentient self and the sensed self, or what William James would have called the relationship between I and Me. In my own case, Me is the human male residing (i.e., located) at 436 Gayle Avenue in DeKalb, Illinois. I, on the other hand, am distributed throughout my phenomenal world. If *I* am to be assigned a single spatial location, it must be in terms of an interpolated personal station point, or personal point of regard, defined by the personal perspective immanent in the phenomenal world called mine. Normally, my personal point of regard (i.e., I), appears to coincide with Me, particularly Me's head.

When persons are asked to point directly at themselves, they tend to point at the bridge of their nose (i.e., at Hering's virtual cyclopian eye). The fact that they are then pointing at their brain is accidental. Imagine a set of Siamese twins in which the brain in head X is connected to the nerves of the body attached to head Y, and vice versa. If a flash card bearing the request, "please point at yourself" is presented only to the eyes in head X, at which head would the pointing arm likely point? At X, surely. And if the request were "please point at your brain," at which head would the arm likely point? Might there not be a different reply to the two questions? And if the hand points at heads X and Y, respectively, in response to these two requests, who would have the authority to question those answers? (By the way, I see none of this as being inconsistent with hierarchical control theory.)

Bill, the same can be argued about the relationship between the We and the Us. The two of Us, You and Me, are in Durango and DeKalb, respectively, but We, You and I, have come together in a dialogue, searching for a common perspective, point of view, or parsing of the world. That is, the proprietary aspect of the phenomenal world includes Our as well as Mine. For one thing, I can imagine (project) my phenomenal world as if from various points in phenomenal space, including those that are currently occupied by other individuals. More importantly, I escape an exclusively personal perspective simply to the degree that I demonstrably share a common perspective with others.

That is, I escape epistemic isolation (solipsism) not by dint of effort, but simply by default. I cannot imagine how colors look to a dichromat (they sort pigmented chips differently than I), but I've got an excellent idea about the trichromat's phenomenal world of colors, without even trying—because we judge (see/sort) pigmented chips alike. Claiming that people who sort all possible pigmented chips perfectly alike do not necessarily see colors alike, as some mischievous philosophers are wont to say, presupposes a fictitious absolute standard of comparison (Noumenal color), because the claim of a difference without a superordinate frame of reference is totally meaningless; further, if such a fictional frame of reference is assumed, for sake of argument, in order to allow the claim to acquire a certain syntactical sense (as does the statement "all invisible things are red"), it still is devoid of empirical meaning. I submit that a putative difference that makes no difference in phenomenal fact, is in fact no difference.

Whereas it is easy to escape epistemic isolation from others, it appears to be impossible to transcend the phenomenal world itself except metaphorically, that is, by a leap of intellect. We might *imagine* a noumenal world of "things in themselves" that transcends all experience, but that is not what science does or should be doing, according to the likes of Pauling and Bridgeman. The theoretical models that scientists conceive must be able to generate precise predictions in the phenomenal domain, because that is where the truth value of the models must be tested.

Science models the order that is immanent in the phenomenal domain. Physics is the branch of science that models the aspect of the phenomenal domain that we call the environment. Physiology models the aspect of the phenomenal domain that we call organisms. That is, physiology and physics conceptually model those aspects of the phenomenal world laymen perceptually model as Me/Us and The Environment, respectively. In contrast, Psychology is a science concerned with the conceptual modeling of the I and the We. The psychology of perception is that branch of the science concerned with the problem of modeling the observer-observed dipole as such. That is, when one models the putative process said to underlie the perceptual aspects of phenomena, one may be said to be modeling a modeling process. In other words, you and I are here involved with conceiving perceiving, or of conceptually modeling perceptual modeling.

When I try to imagine phenomena's substance from a psychological perspective (i.e., the essential substance of the epistemological dipole) I find myself coming up with words like immanent order or detectable structure or information—all of which are compatible with physics and physiology. It does not appear inappropriate to call such information "signals," but it does appear inappropriate to call them "neural

signals," thereby excluding all other signal types, because that is to forget the bipolar nature of the phenomena. *The epistemic unit is the dipole.* For example, I comprise a dipole characterized as me *and* my environment. In your model, this epistemic unit takes the form of an ecological control loop having two poles, characterized as a unique organism and its environment, including all other organisms. Because there are as many dipoles as there are organisms, with each organism being part of many dipoles, your control theory model is not necessarily solipsistic.

Because a single organism plays a unique role in each of these dipoles, it is tempting to suppose that the dipole is within that unique organism. That is, it is tempting to suppose that I am in my head, but that notion is not only illogical, it is also contraindicated by the fact that my phenomenal head is in my phenomenal world—along with a bunch of other phenomenal heads. Therefore, whenever I use the word "perception" to denote this personal aspect of phenomena, I try to remember that I am referring to a personal perspective or point of view rather than to a personal replica.

The *bipolar* nature of *objective* phenomena is what our reverse engineering must explain. Your hierarchical-control-theory model accounts for both of these in terms of interacting control loops. As far as I can see, your model poses no epistemological problems, and it disturbs me to hear you imply, sometimes, that it does. If anything, your model promises to resolve epistemological problems, not create them. That's the way I see it.

Bill Powers: Wayne, I'm going to avoid the temptation to get back into the epistemological argument; I'll let you have the last word. I like your exposition considerably, but there are still problems to work out—like what we should say neural signals are for, in our models.

Levels of Perception

Mark Olson: (To Bill Powers:) I'm having difficulty understanding the quality of the input signals for higher-level systems. I know that an input signal for one system is an integration of ( ) from lower systems—what's in the ( )? Are the signals that become integrated simply the same as the signals the first level receives? Or are these signals somehow adapted? All the diagrams I've seen make it seem as if they are not adapted, that they simply go all the way up. But if they are integrated, then something is different. I'm close but I don't quite have it.

Bill Powers: Mark, you ask about the functions relating one level of perception to another. This is indeed the question that hierarchical perceptual control theory (HPCT) poses—but doesn't answer. What lies behind HPCT is not any proposal as to how each level of perception is derived from the one below it, but a proposal as to what the levels of perception are and how they are related. This is the phenomenon that any model must in the end explain.

The "H" part of HPCT can be taken in two ways: first, as a general sketch of a hierarchy of control in the abstract, with the communication between levels consisting of a series of perceptual re-representations of reality and a corresponding set of reference signals used to control lower levels; second, as a series of proposed levels of perception (and control) based directly on an analysis of experience with the hierarchical-control concept as a guide. This is a beginning model; there might well be other modes of communication between levels, but the basic one is probably valid.

The definitions of levels define the modeling problem. We can see that the sensation level is probably derived by weighted summations of intensity signals, the weights defining a vector in a perceptual space having fewer dimensions than there are different sources of intensity signals. But that answer to the modeling problem comes after noticing that sensations seem to depend on intensities in a particular way, a way that could be modeled as weighted summation. The phenomenon to be modeled comes before the model.

And that's as far as I can go. I don't know how configurations are derived from sensations—how it is that we can get the sense of, say, a particular person's face over a range of distances and orientations and expressions. If signals standing for the dimensions of a face existed, then it's possible to make a rough guess that transitions of the face from one state to another would be sensed using time functions and

partial derivatives; that's a feeble start toward a functional model that you could run on a computer. As to the rest of the levels, the kinds of computations involved are mostly a mystery to me. The few guesses we have come up with are strictly stabs in the dark. You can use words like "integration" to describe how some kinds of perceptions are put together to create others, but the word is just a noise. It doesn't tell us anything about the processes involved.

Behind this exploration of perception lies a fundamental postulate; if you don't internalize it, I don't think you can even get started on the problem of modeling the brain's perceptual systems, or, for that matter, in understanding HPCT. The postulate, simply put, is this: it's all perception. By that, I mean that no matter what you attend to in the world of experience, whether you refer to inner or outer experiences, concrete or abstract, verbal or nonverbal, the object of your attention is a perception. You are looking at or otherwise experiencing the brain's perceptual activities, not the objective world itself.

Vision is the most important sense to understand this way if you're sighted; understand vision and the rest (touch, taste, sound, etc.) will follow. The world you see begins as pixels (individual picture elements). The pixels are so close together that you see no spaces between them, although the sensory nerves do not overlap and in fact do not completely fill the retina. There's a world between the pixels, but we don't see it unless the view shifts slightly—and then what we had been seeing disappears into the cracks between the pixels. This is invisible to direct experience; the world seems continuous over the whole visual field. We get a sense of seeing the world at infinite resolution, and can't imagine what the whole field would look like if we had, say, ten times as many retinal receptors and the optical acuity and brain power to take advantage of them. This would be like seeing the world through a magnifying lens, except that the whole world would look that way, not just one little part of it (which we still see at human resolution). The only way to imagine this is to go the other way: view the world at a lower resolution, as in a halftone photograph or a television screen seen close up, and imagine that the result is the only world you can ever see. That's how our picture of the world would look to a different organism with higher visual resolution. But we experience it as having continuous detail right down to the level where it appears smooth. I suppose the fly sees the world in the same way. But its world is smoother than ours.

Building up definitions of the rest of the levels in the hierarchy is then a matter of noticing persistent types of structure in this world of picture elements. The first level above the pixels themselves is sensation, a type of perception that can't be analyzed in any way except into variations of intensity. Color is a sensation, as is shading.

Perhaps things like edges are sensations, derived in one step from the pixel distributions. When analyzing perceptions, however, don't use any data but your own experience. Theory and neural data will tell you that in the visual field, in the retina itself, all edges are enhanced, so that there is a strong outlining effect. But look at the edge of a sheet of paper on a dark tabletop. There is no outline. The closer you look at the edge, the more nearly it seems to be an infinitely sharp line separating uniform white from uniform dark. The edge itself is there—but you can't see it as an object. It's just a sense of edgeness. Only under special conditions, as in looking at a smooth gradient of illumination going over a relatively short distance from white to black, do you see edge effects like the "Mach band," the only clear subjective evidence of edge enhancement. However those neural signals enhanced at edges are processed, the result is that step changes look like step changes, not outlines as in cartoons. Whatever model we come up with for how the nervous system processes pixel information, it must result in edges that look this way, without borders. If it doesn't, the model is wrong.

The next step is to notice that the edges and corners and broad white areas of the piece of paper add up to—a piece of paper. If you've made this transition properly, it will come as a surprise. Where did that piece of paper, or piece-of-paperness, come from? It wasn't there in the edge, or the corner, or the whiteness, or the darkness. It comes into being only when all those elements are seen grouped into a thing, a configuration with a familiar shape, orientation, distance, size, and so on. The Gestalt psychologists of old spent a lot of time looking at things like these. They should have kept going. Or perhaps they shouldn't have been cowed by the behaviorists.

You have to go slowly and by the smallest steps you can devise. If you go too fast, you'll miss the smallest steps; if you miss the smallest steps, you'll lose the sense of examining perceptions and start projecting the visual field into an external world again. You'll jump to the more abstract levels and lose the connection from one level to the next. This is, if you like, a form of meditation on experience in which you distance yourself from experience and look at it merely as a display. You're not trying to see anything about the world, but only something about the display. You're trying to see what features the person who constructed it thought of putting into it, just as when you read a program, you think to yourself, "Now he's setting up an array to hold the results," instead of just reading the code, or when you read a novel as a literary critic, you think "Now he's introducing tension," instead of just getting tense. Who the "he" is is immaterial—the point is to see what is before you as a construction that has inner organization, and to try to see how it is put together.

The general principle is that when you have found a level, like sensation, the next level is going to depend on it; also, the current level depends on the one below it. If you analyze a perception to see what it is made of, at first you see just more perceptions of the same level—big configurations are made of little configurations. But when you analyze in just the right way, you suddenly realize that all configurations, of whatever size or kind, are made of sensations, which are not configurations of any kind. And you realize that if it weren't for the presence of those sensations, there couldn't be any configuration to see: a field consisting of a single sensation, such as white, can't lead to any sense of configuration. There's a relationship between these levels of perception. That gives us a hint about building models of perception, a hint about how the brain's perceptual system is constructed.

Sometimes you will identify what seems to be a higher level of perception, some characteristic common to all perceptions, unconnected to lower levels you have previously seen. Then you can use this kind of analysis to try to fill in the gap. What is this new perception made of, besides smaller perceptions of the same kind? When the gap is large, the missing steps are obvious. You can, for example, look at spatial relationships such as "on"—something being "on" something else. You can see the on-ness clearly, it's right in front of you. But what is it made of? If you said "sensations," you would clearly be making too large a jump, because on-ness involves objects, things, configurations. Some kind of object is "on" some other kind of object. If it weren't for the impressions of distinct objects, there couldn't be any sense of the relationship between them. But is that step small enough? I've had to put two levels between relationships and configurations: transitions (which can be zero) and events (which can be as simple as mere duration). Seeing something "on" something else involves more than a brief contact; there must be duration.

Perhaps someone else could find smaller steps still, or would characterize the intervening steps differently. There's still a lot of room for improving the definitions of the phenomena we're hoping ultimately to model.

I'm not talking here about the models themselves. I'm talking about the attitude you take toward your own experiences when you're trying to notice phenomena that need modeling. If you were a physicist, you wouldn't be taking this attitude. You'd treat the world of perception in the normal unanalytical way, as if it lay outside yourself where everyone could see it, and you'd search for laws relating changes of one kind of perception to other kinds of perceptions. You would call these "natural laws" or "behavioral laws" and assume you were discovering truths about an objective universe.

As a control-theory psychologist, however, you have a different objective: to grasp the natural world as a manifestation of human perception (your own), and to ferret out of it some regularities that tell us about perception rather than about the world perceived. If you stumbled onto this attitude accidentally, without understanding what you were doing, you might well find yourself in a state with a clinical name: dissociation. I don't recommend this attitude as one suitable for ordinary living. It's difficult and uncomfortable, and it tends to strip the meaning from experience (until you get past a certain point, after which you realize that meaning, too, is perception, and let it back in). If you're afraid that understanding your girl friend as a set of intensities, sensations, configurations, transitions, events, relationships, categories, sequences, programs, principles, and system concepts in your brain might strain your feeling toward her (and hers toward you), don't do this with your girl friend. Do it with somebody else's, or a laboratory rat. It doesn't matter who or what you do it to, because you're really talking about your own perceptions. This is a private experience valid only in one person's world. It can become public only to the extent that different people independently arrive at the same analysis. I've always hoped for that, but only a very few people, to my knowledge, have tried this for themselves. Most people just memorize my definitions, which unfortunately are in words. It's easier to push words around than to shut up and examine direct experience.

You'll hear objections to this process, alluding to introspectionism, which failed to get anywhere a long time ago. But introspectionism didn't fail because it looked at the kinds of things I'm talking about here. It failed because it confused the subjective with the objective (and so did its critics). The world that I'm speaking of examining here would be called, by most conventional scientists, the objective world, not the subjective one. I'm not recommending shifting attention off the objective world and plunging into the dim and uncertain world of inner phenomena—or what we imagine to be inner phenomena. I'm recommending a change of attitude toward the world we normally consider to be the objective one, which includes the world outside us and our bodies as we experience them. I'm saying that you will learn something if you look on this world as directly experienced evidence about the nature of your own perceptual system, and only in a conjectural way about the world that is actually outside you.

Instead of treating relationships like on, beside, after, with, and into as properties of the external world, look on them as perceptions constructed on a base of lower-level perceptions. Instead of seeing categories as made of things that are inherently alike, think of categories as ways of perceiving that *make* things appear to be alike—things that are actually, at lower levels of perception, different. Instead of seeing sequential ordering as *a* fact of nature, see it as a way of putting ordering into an otherwise continuous miscellaneous flow. In short, take nothing about experience for granted, as if some aspects of experience were really outside and others were inner interpretations. Put the whole thing inside, and see what you come up with when you understand that it's all perception. All of it.

In HPCT diagrams, we show signals coming out of perceptual functions and going into higher-level ones (as well as the local comparator, if the signal is under control). I think of these lines as representing single neural signals that vary in only one dimension: how much. This can be confusing, because we don't experience single signals under normal circumstances (when we do, they cease to be meaningful). Instead we experience all the signals within the scope of awareness, at every level in the state we call conscious. To understand what the single-signal concept means, you have to break this world of simultaneous perceptions into its components, the individual and independent dimensions in which the totality of perception can vary. You have truly identified one isolated perception when it can vary only in the degree to which it's present, which we experience as its state. If the perception varies without in the slightest changing its identity, you have probably noticed a single signal.

This can be important when you talk about control. We talk loosely about controlling "a dog," for example. But that way of talking is really lumping many independently variable aspects of the dog together. You don't control its species, or its eye color, or the length of its tail. You don't even control its behavior. If it's behavior you're controlling, you always control *some particular variable aspect of the dog's behavior.* You might control the radius within which it can move by putting it on a chain. You might control its speed of walking by saying "stay" or "follow," and its path by saying "heel." Whatever you control, it must come down to a single variable or small sets of variables independently controlled. If you're controlling in more than one dimension, you must sense more than one variable and have a control system operating independently for each one. That's because independent dimensions can be independently disturbed; you need independent control systems so that a disturbance in one dimension can be corrected without necessarily causing an error in another dimension.

None of this answers your question as to how perceptual signals in a diagram depend on perceptual signals lower in the diagram. The only general answer I can give is that some computation lies between them. The input data consists of lower-level perceptions; the output data, the higher-level perceptual signal, represents the value of the function being computed over and over or continuously. At each level, I presume (judging from the way the context changes every time you consider a higher level), a new type of computation is involved, not simply a repetition of the kind of computation at the lower level. The process of deriving categories from sets of relationships can't be carried out by the same kind of computation that derives relationships from sets of events or lower perceptions. There is no one kind of computation that could serve at all levels.

But as I say, I am—we all are—a very long way from grasping what these kinds of computations are. Every time people come up with a new computer program for recognizing objects, they try to establish this new computation as the blueprint for the whole perceptual system. This is a waste of time. The blueprint changes with every level. Weighted algebraic summation is simply not going to suffice to model our capacity to recognize and execute a program described in words: a rule. Even though such networks are purported to recognize categories, I think that the category-ness is read into the results by a human observer. I don't think that any category-recognizing back-propagation model will actually create what human beings experience as categories—for example, the category "wife." Of the 11 levels of perception in my model, I think we know how to model two of them: the first two. All the rest of our modeling presents to us what a human being might recognize as a higher-level perception, but which the circuit or program itself does not recognize or control.

In that I could be wrong, of course, because I speak the truth when I say I don't know how the higher levels of perception work. That means I don't know how they don't work, too. I'm just expressing a hunch.

On Modeling

Bill Powers: Of the hundred-odd people on this net, I don't suppose more than a handful understand what some CSGers mean when they talk about modeling behavior. So I thought I'd explain it a little, at least as the process appears to me. Talking about modeling is a little like talking about control—most people have some concept to go with the word, but not many outside the engineering professions (and not everyone in them) mean what I mean by it.

I'm working now on a model of pointing behavior. On the surface, it's not very impressive. The computer screen shows a little stick man with one arm who reaches out and touches, or continuously tracks, a floating triangle that the user can move around from the keyboard in a perspective drawing of a three dimensional space. It looks like a cartoon of a not very interesting behavior. While movements are a bit more realistic than you find in most cartoons, most people have seen more impressive TV cartoons in which more interesting action occurs. But behind this surface appearance is the model; what's interesting is not so much *what* happens on the screen, but *how* it happens. To explain how it happens, I have to distinguish the kind of modeling I use from other kinds.

The first distinction of importance is that this kind of model is not an animation. That is, the various movements of the arm (and head—the little man always looks at the target) are not simply drawn frame by frame as in the Disney Studios. It's not done the way interactive video games are done, by switching from one animated sequence to another depending on what the user does at the keyboard. Instead, the program is reacting directly to the location and movements of the floating triangle, which are totally unpredictable by the program. I can guarantee that the program makes no attempt to predict the target movements, because I wrote it.

The second distinction of importance is that in this kind of model, there is nothing in the program that computes the actual movements of the arm as we see them. If the arm's fingertip moves in a straight line, this is not because something in the program computes the detailed actions needed to produce a straight line. Likewise for curved movements, or movements that begin fast and slow down as the fingertip nears the target. None of these aspects of movement corresponds to any specific calculation of path or speed in the program.

In some approaches to modeling, such calculations are the heart of the method. One looks at the actions and figures out what commands would be needed to produce them. If the fingertip is to move along a path and intersect a moving target, such a model would use the target movement information as input and find a path and a speed profile that would bring the finger to the same place as the target some time in the future. Then it would drive the computed arm so as to achieve that path and speed profile, thus bringing about the predicted intersection. Basically, this concept of modeling attempts to reproduce the visible behavior by calculating its details, given all of the physical factors of the situation.

The approach I use is more properly called "simulation." Inside the computer are program modules. Each module computes what some simple element of the real system would do when presented with continually varying inputs. Some of the modules are perceptual modules: they compute what certain nerve signals would do as the aspect of the environment to which a sensor is sensitive changes its effects on the sensor. For example, one module represents a muscle spindle, which emits a signal that depends both on the length of the muscle and on another neural signal, the gamma efferent signal. Another represents the tendon receptors that are affected by the muscle tension.

One of the modules is an effector module: it represents the muscle's response to a motor signal from a spinal motorneuron (including the shortening of its contractile part and the consequent stretching of its spring-like component to produce a force). And there are many more modules that represent the way hypothetical sets of neurons respond to neural signals by producing more neural signals. There are sets of modules that are repeated, with the same interconnections, for each muscle in the model.

In this model, by the way, I don't use actual models of individual neurons, although I could. Such a level of detail would not add anything to the performance of the model and would increase the size of the program and slow its operation. What I do instead is use simple calculations similar to what a neural model would do: add signals, subtract one signal from another, amplify signals, and do time integrations and (rarely) differentiations. Nothing more complex.

Each module is meant to represent the way some small part of the real living system works, as nearly as I understand it. Many of the modules represent guesses based on hints from neurology or even from waving my own arms around and paying attention to the details, and they constitute the conjectural parts of the model.

The model is not just a collection of computing modules: it is also a pattern of connections joining one module to one or more others. For example, there are modules representing the static and dynamic parts of the stretch receptors in muscles. The outputs of these modules, conceptualized as neural signals, become inputs to the module representing the spinal motorneuron. This motorneuron module produces an output that is the sum of several positive inputs from other modules and a negative input from the tendon receptor module. The output of the spinal motorneuron module becomes the input to the module that computes the muscle force output. And so on. Each module is woven into the whole model through its input and output connections from and to other modules.

A more subtle aspect of this process is that the model contains adjustable parameters in the links between modules. The dynamic stretch receptor module, for instance, sends its signal to the spinal motor-neuron module, but there's a parameter that determines how much effect this signal is to have at the spinal motorneuron, and the sign of the effect. If the parameter is set to a high value, the simulated arm behaves sluggishly or, at the extreme, breaks into high-frequency oscillations. If it's set to a low value, the arm begins to wobble around, and even goes into ever-increasing low-frequency oscillations. If the parameter has the wrong sign, the arm will behave more and more wildly, until the whole program blows up.

So it's not enough to model the right kinds of components of the real system, or even to connect them into a network like the real neural network, with the right signals going to the right places. The quantitative parameters can be adjusted to make a model, with any given components and any given pattern of interconnections, do completely different-looking behaviors.

Finally, there's a real-time aspect of this "simulation" kind of modeling. All the computations in all the modules are carried out effectively in parallel. One such parallel computation covering all modules represents one increment of real time, dt. In the arm model, dt represents 0.01 second of physical time (regardless of how long it takes the computer to finish all the computations). The last computation is to recompute all the outputs of the modules, so they have all changed before the next cycle when they will be treated as inputs to other modules. This sometimes requires paying close attention to the way the program is written, so that things supposed to be happening at the same time don't accidentally happen in sequence—one dt too late. In an analog computer, this requirement would be easy to meet, because all of the computing components would be acting at the same time. But in a digital computer, where there is only one busy central processor that has to do everything, achieving the effect of simultaneity isn't always easy.

After each round of calculations, all of the modules have new outputs, which become inputs to other modules (or even the same module) at the start of the next time increment. With a dt of 0.01 second, the result is very close to continuous operation, with all signals (inputs and outputs of modules) varying smoothly and simultaneously. The test to see whether the incremental approach is sufficiently like a true continuous computation is to decrease the size of dt—let each complete computing cycle represent, say, 0.001 second. If the same behavior results, but in smaller steps of movement, then the larger time increment is short enough. It's nice to use longer intervals, so the movements of the model become fast enough to see between breakfast and lunch. The arm model in its present form runs at about one/fifth of real time (on a 10-MHz IBM-AT-compatible programmed in C).

One of the modules is a physical model of the arm. The inputs to this module are three torques being applied by the muscle modules to the three joints during one time increment. Using kinematic equations, calculating Coriolis forces and all that, these torques are transformed into angular accelerations around the three joints (taking the moments of inertia and masses of the arm segments into account). Those accelerations are integrated to produce angular velocity, which is integrated to produce angular position. The three angular positions are inputs to the behavioral model, determining the new joint angles and angular velocities, and the new muscle lengths and rates of change of muscle length for the start of the next dt.

There are two inputs to each muscle control system: an alpha efferent and a gamma efferent. When these signals are varied (for testing purposes), the arm will go through certain motions on the screen. I use a standard test signal which simply switches from a positive value to zero and back again, with a half-second interval between transitions.

What the arm segment being tested *should* do is move quickly from one angle to another, stay there for a half second, and move and dwell for another half second, over and over. What *does* happen, of course, is initially something very different. There are five parameters to adjust, representing five meaningful aspects of the control system: three sensor sensitivites, one sensitivity of muscle contraction to driving signals, and the spring constant of the muscle. Only the muscle spring constant can be estimated from observations and data in the literature. The other four have to be guessed at. Finding the right combinations of values can be done in part through computations, but there are so many interactions and nonlinearities in the model that exact predictions are impossible (certainly for me). So what one ends up doing is changing the parameters experimentally until the arm begins behaving properly, or as nearly properly as possible, without adjusting the parameters of the other control systems, too.

Estimates of parameter values and, especially, of the behavioral effects of varying parameters can be made, but only for small segments of the model such as a single control system for a single joint. Such estimates get you in the right ballpark for each control system's parameters. But it's impossible to write the equation for the whole model and

solve it for the best values of parameters. The equations are all nonlinear differential equations (made more nonlinear when the visual part of the model comes into play), and the interactions among parts of the model are large (extending the arm at the elbow joint affects both arm segments, for example, through inertial interactions). This brings us to the heart of simulations.

The reason we do simulations is precisely that we can't analyze or even understand the whole model at one time. The postulates of the model are in the definitions of the modules. These modules are each very simple and are closely related to simple properties of the nervous system and muscles. So we can easily understand what each module does or is postulated to do.

What we can't easily understand is what will happen when we conned the modules together in some specific way, with specific interconnection parameters. Our postulates about the modules completely determine the behavior that is implied; the only problem is that we can't deduce our way from the postulates to their actual implications.

A simulation shows us the implications directly. It says to us, "I don't know what you thought you were modeling, but here's what you *did* model." It's just like a computer program, which does what you told it to do, instead of what you wanted done. A simulation cuts through all the fuzz of verbal explanations and imprecise reasoning about what a particular model *ought* to do. A simulation is a way of finding out the implications of propositions that are linked together in such a complex way that human reasoning is inadequate to reach a conclusion.

Human reasoning becomes inadequate for most real systems with more than three or four components. Even mathematical analysis is usually impossible in the real world, which doesn't fit the idealized forms that we know how to handle analytically. One result of this fact is that people regularly try to fit the real world to those mathematical methods they *do* know how to handle. Every new discovery of some tractable mathematical phenomenon is followed by a hoard of people trying to make nature behave that way. Hence, chaos theory and its application to literally every unsolved problem, particularly in the nervous system. There are phenomena to which chaos theory applies; in fact, chaos was discovered through observing a working simulation of the weather. But in other contexts it's a solution looking for a problem.

An alternative to analysis is simulation. You hook up a model of the system in which the simple components are represented or plausibly conjectured, turn it on, and gape at what it does. The model then becomes an experimental object. You can play with it, altering its components, their interconnections, and the connection parameters, and learn the effects of each kind of change. Each variation leads you to understand something about the real system. You find out why a given connection is positive instead of negative. You find out why certain connections are present in the real system and others are not. The "why" in every case is simply that the model doesn't act like the real system in some relatively dramatic way. And you can *see* why it doesn't.

A simulation is like an X-ray into the real system, showing you aspects of its functioning that can't be observed directly. Like an X-ray, the simulation can be ambiguous; the observed behavior can be accomplished by more than one plausible model. As with X-ray interpretations, however, we don't have to rely on ambiguous indications; we can think up alternative diagnostic tests that will rule out some possible models, and with increases in technical skill, we can even open up the system and see some of the connections, even monitor some of the circuit activities. Every added piece of observational evidence narrows the field of models that would behave correctly *and* work by the right means.

There's another side to the subject of observational evidence. Often the observational evidence is available, but isn't understood. To say it isn't understood is to say that there's no model that needs that evidence. The combined stretch and tendon reflexes are a case in point. These reflexes have been known for close to a century, but nobody has understood what they are for. There have been vague qualitative conjectures, of course. But the arm model I'm working on shows quantitatively what these reflexes do. The tendon reflex controls applied force. The dynamic stretch reflex controls the integral of applied force, or angular velocity. The static stretch reflex controls the integral of velocity, or angular position at a joint. The model shows that with certain values of the parameters, this combination of control systems makes the arm extraordinarily stable, quick to respond to driving signals, and consistent in response over a wide range of external conditions and internal conditions of the muscles. While I haven't demonstrated this yet, it's clear now that this combination of reflexes easily compensates for the extreme nonlinearity of the muscle's tension-extension curve. In fact, when I realized finally how this system works, I was amazed at its cleverness and simplicity.

But those who traced the circuits and measured their details couldn't have seen that cleverness and simplicity, because not having modeled the system, they didn't see all the problems that it solves with such economy. These reflexes can be seen as a remarkable design only after you have looked into the problem of controlling a jointed arm in some detail. I couldn't have designed that system. I simply designed the model to be as much like what I knew about the stretch and tendon reflexes as possible, turned it on, played with the parameters, and discovered beauty.

The whole arm model is built up this way. It behaves as it does because of the interactions among its modules. It reaches out and touches the target, and follows the target around when it moves, and looks at the target, and resists gravity, and moves at various speeds and along various paths in the process, because there is nothing else it can do. We are seeing in this kind of behavior the necessary consequence of organizing a system the way the model is organized. Maybe another organization would also have to behave this way. But this one behaves like a human being, at least at these levels of organization, and to the extent possible, its modules are similar in function to known modules in human systems. The external physics and optics in the model conform to what is known about physics and optics, near enough. Some parts of the model are in one-to-one correspondence with direct observations. Some parts are conjectured. But the X-ray seems to be showing a convincing shadow of the real system, at least as it is seen from this angle.

The Control Systems Group is a membership organization which supports the understanding of cybernetic control systems in organisms and their environments: *living control systems.* Academicians, clinicians, and other professionals in several disciplines, including biology, psychology, social work, economics, education, engineering, and philosophy, are members of the Group. Annual meetings have been held since 1985. CSG publications include a newsletter and a series of books, as well as this journal. The CSG Business Office is located at 73 Ridge Rd., CR 510, Durango, CO 81301; the phone number is (303)247-7986

The CSG logo shows the generic structure of cybernetic control systems. A Comparator (C) computes the difference between a reference signal (represented by the arrow coming from above) and the output signal from Sensory (S) computation. The resulting difference signal is the input to the Gain generator (G). Disturbances (represented by the black box) alter the Gain generator output on the way to Sensory computation, where the negative-feedback km/ is closed.

Back cover

Closed Loop

Threads from CSGnet

This reproduction of *Closed Loop* was created by Dag Forssell in 2009. Addresses and phone numbers have not been updated. Most are obsolete.

Posted at www.pctresources.com

Spell checked, but not proofread. Please report any errors to webmaster. Email address at website.

Summer 1992

Volume 2

Number 3

Front cover

Closed Loop

Threads from CSGNet

Summer 1992 Volume 2 Number 3

Edited by Greg Williams, 460 Black Lick Rd., Gravel Switch, KY 40328

CONTENTS

Members of the Control Systems Group receive *Closed Loop* quarterly. For membership information, contact Ed Ford, 10209 N. 56th St., Scottsdale, AZ 85253; phone (602)991-4860.

CSGNet, the electronic mail network for individuals interested in control theory as applied to living systems, is a lively forum for sharing ideas, asking questions, and learning more about the theory, its implications, and its problems. The "threads" in each issue of *Closed Loop*, stitched together from some of the Net's many ongoing conversations, exemplify the rich interchanges among Netters.

There are no sign-up or connect-time charges for participation on CSGNet. The Bitnet address is "CSG-L@UIUCVMD" (use no quotes in this and the following addresses); "CSG-L@VMD.CSO.UIUC.EDU" is the Internet address. Messages sent to CSGNet via these addresses are forwarded automatically to all participants. Via CompuServe, use the address "> INTERNET :CSG-L@VMD .CSO.UIUC.VMD" to reach the Net. Initially, you should send a note to the network manager, Gary Cziko, at "G-CZIKO@UIUC.EDU" (Internet) or at "c CZIKO@UIUCVMD" (Bitnet); Gary's voice phone number is (217)333-4382.

Statistics vs. Generative Models

Bill Powers: Before I spend time trying to explain a phenomenon, I want to know if it's real or just statistical. I want to know things like how many people show the phenomenon, how you find out that there's a phenomenon, how many trials show the effect, and how many don't—all that stuff. Once I'm convinced that there's a real phenomenon, it's time to think up explanations.

I'm not interested in 80 per cent correlations. That's way too low to define a phenomenon.

"Superficial" knowledge is knowledge gained by observing apparent causal or coincident relationships without any generative model of underlying processes. Statistical studies yield superficial knowledge.

I think that all attempts to apply abstract physical principles and advanced mathematical trickery to human behavior are aimed at solving a nonexistent problem. They all seem to be founded on the old idea that behavior is unpredictable, disorderly, mysterious, statistical, and mostly random. That idea has been sold by behavioral scientists to the rest of the scientific community as an excuse for their failure to find an adequate model that explains even the simplest of behaviors. As a result of buying this excuse, other scientists have spent a lot of time looking for generalizations that don't depend on orderliness in behavior; hence information theory, various other stochastic approaches, applications of thermodynamic principles, and the recent search for chaos and quantum phenomena in the workings of the brain. The general idea is that it is very hard to find any regularity or order in the behavior of organisms, so we must look beyond the obvious and search for hidden patterns and subtle principles.

But behavior *is* orderly, and it is orderly in obvious ways. It is orderly, however, in a way that conventional behavioral scientists have barely noticed. It is not orderly in the sense that the output forces generated by an organism follow regularly from sensory inputs or past experience. It is orderly in the sense that the *consequences* of those output forces are shaped by the organism into highly regular and reliably repeatable states and patterns. The Skinnerians came the closest to seeing this kind of order in their concept of the "operant," but they failed to see how operant behavior works; they used the wrong model.

Because of a legacy of belief in the variability of behavior, scientists have ignored the obvious and have tried to look beneath the surface irregularities for hidden regularities. But we can't develop a science of life by ignoring the obvious. The regular phenomena of behavior aren't to be found in subtleties that can be uncovered only by statistical analysis or encompassed only by grand generalizations. The paydirt is right on the surface.

The simplest regularities are visible only if you know something about elementary physics—and apply it. Think of a person standing erect. This looks like "no behavior." But the erect position is an unstable equilibrium, because the whole skeleton is balancing on ball-and-socket joints piled up one above the other. There is a highly regular relationship between deviations from the vertical and the amount of muscle force being applied to the skeleton across each joint. There is nothing statistical, chaotic, or cyclical about the operation of the control systems that keep the body vertical. They simply keep it vertical.

The same is true of every other aspect of posture control and movement control, and all controlled consequences of those kinds of control. Just watch an ice skater going through the school figures in competition. Watch and listen to any instrumentalist or vocalist. Watch a ballet dancer. Watch a stock-car racer. Watch a diver coming off the 30-meter platform. Watch a programmer keying in a program.

It's true that when you see certain kinds of human activity, they seem disorganized. But that is only a matter of how much you know about the outcomes that are under control. The floor of a commodities exchange looks like complete disorder to a casual bystander, but each trader is sending and receiving signals according to well-understood patterns and has a clear objective in mind—buy low, sell high. The confusion is all in the eye of the beholder. The beholder is bewitched by the interactions and fails to see the order in the individual actions. When you understand what the apparently chaotic gestures and shouts *accomplish* for each participant, it all makes sense.

Of course, we don't understand everything we see every person doing. It's easy to understand that a person is standing erect, but *why* is the person standing erect? What does that accomplish other than the result itself? We have to understand higher levels of organization to make sense of when the person stands erect and when the person doesn't. We have to understand this particular person as operating under rules of military etiquette, for example, to know why this person is standing erect and another is sitting in a chair. But once we see that the erectness is being controlled as a means of preserving a higher-level form, also under control, we find order where we had seen something inexplicable. We see that an understanding of social ranking, as perceived by each person present, results in one person standing at attention while another sits at ease. Each person controls one contribution to the pattern that all perceive, in such a way as to preserve the higher-level pattern as each person desires to see it.

It seems reasonable that once we have understood the orderliness of simple acts and their immediate consequences, we should be able to go on and understand more general patterns that are preserved by the variations that remain unexplained. As we are exploring a very large and complex system, we can't expect to arrive at complete understanding just through grasping a few basic principles. We must make and test hypotheses. But if we are convinced that the right hypothesis will reveal a highly ordered system, we will not stop until we have found it. If, on the other hand, we are convinced that such a search is futile, that chaos reigns, we will give up the moment there is the slightest difficulty and turn to statistics.

I claim that human behavior is understandable as the operation of a highly systematic and orderly system—at least up to a point. I say that it is the duty of any life scientist to find that orderliness at all discoverable levels of organization, and to keep looking for it despite all difficulties. We must explore all levels, not just the highest and not just the lowest; what we find at each level makes sense only in the context of the others.

We have a very long way to go in understanding the obvious before it will be appropriate to look for subtleties. I have no doubt that we will come across mysteries eventually, but I'm convinced that unless we first exhaust the possibilities of finding order and predictability in ordinary human behavior, we won't even recognize those mysteries when they stare us in the face. I don't think that anyone is prepared, now, to assimilate the astonishments that are in store for us once we have understood how all of the levels of orderly control work in the human system.

We won't get anywhere by looking for shortcuts to the ultimate illuminations that await. Most of the esoteric phenomena of physics that are taught in school today were occurring in the 19th Century. But who, in that century, would have recognized tunneling, or coherent radiation, or shot noise? If we want to see a Second Foundation of the sciences of life, we have to begin where we are and build carefully for those who will follow us. If we succeed in trying to understand the obvious, the result will be to change what is obvious. As the nature of the obvious changes, so does science progress.

Chuck Tucker: I think that the majority of those who have difficulty accepting our approach simply hold to the assumptions about the world attacked by Dewey in *The Quest for Certainty,* and rejected by us: that the real world will be revealed to you if you just use the "proper" methods and work hard enough. If we tell these people that their approach won't reveal the "true forever world," then they seem to have much less interest in what we have to say. Another feature of many of those who reject our view is that they are not "problem-oriented"—that is, they do not tolerate ambiguity, uncertainty, and problem-solving activity for very long; they want the answer quickly and cheaply (or statistically). But our approach does not offer such a magic solution; just hard, dirty, difficult work, with no absolute assurances that a solution will be fashioned, let alone work. Think about it: would you give up such a pleasant life of certainty and bliss for the one we offer? Probably not.

Bill Powers: David Goldstein and I have been conducting an argument for several years. David tends to win many of the rounds because he is working with clients who have both real and severe problems, and I often have to admit that when you're faced with solving such problems, you have to do what's possible. If a person is so depressed as to be on the verge of suicide, you give the person a pill that takes the edge off, and you're glad that such a pill exists. Afterward, you can think about trying something else. Even control theory can't cure a dead client.

A lot of our arguments are conducted in the context of such practical limitations. But I don't have David's responsibilities, so I can argue against conventional methods even if I don't have an immediately applicable alternative to propose. One of these arguments has to do with the utility of testing, particularly testing that involves questionnaires and other means of self-description such as Q-sorts. Basically, I argue that verbal tests are too imprecise to do much good, and that they inevitably put us in the position of applying statistical methods to individuals. I argue that we should be trying to apply control theory directly, trying to find out what individuals can and can't control, and trying to find out why they are having trouble. This means abandoning old diagnostic categories and old attributions of traits and conditions in the attempt to explain what's wrong. I claim that we must make a conscious effort to break free of cultural assumptions, which always steer us back toward the conventional categories. David doesn't exactly disagree with me, but—well, he can speak for himself.

David has proposed "qualitative modeling," the sound of which I rather like. He says, "Suppose that we plotted the urge to perform action X against time. The lowest point of the curve can be taken to be the reference level for whatever perceptions are being controlled by action X. Suppose that on a scale of 0 to 10, the intensity of perception Y1 = 2 and the intensity of Y2 = 5 at the lowest point. As a person deviates from these values, control theory leads us to expect increasingly stronger urges to perform action X the further we move away from these reference level values. If we do not obtain a U-shaped function around these values, then the particular clinical hypothesis may be rejected.

What do you think?"

I think that the method as stated predetermines too many variables. The first objective should be to see what perceptions are under control. To do that, you have to allow the action-variable to be free. If the perception is "people like me," the action that will contribute to that perception will be different under different circumstances (meaning different disturbances of the sense that people like me).

Under the conventional approach, we would be most concerned with the action, because that is what other people experience. But to understand the acting person, we first have to understand what perceptions are under control. A given perception can be controlled through many different actions, so no one action is significant by itself. Furthermore, we might see both an action and the opposite action being taken as a means of controlling the same perception, depending on whether disturbances are pushing the perception above or below its reference level. The object of control theory can't be to explain one particular action.

So I would propose backing up a step or two, and starting by testing Yl, Y2, Yn to see if they are controlled variables. This is hard to do using a verbal test, first because while taking the test, the person isn't experiencing the perception, but only a description of the perception, and second because the only way to apply disturbances is hypothetically, by describing them and asking how the described disturbance would affect the described perception (and, presumably, what the person would do if the perception changed). I much prefer direct interaction in real situations, with perhaps a discussion afterward if you want to cast the interaction in verbal terms. Maybe role-playing would be a compromise that allows setting up hypothetical situations while still allowing real perceptions and direct interaction with disturbances (supplied by the experimenter).

Gary Cziko: I have read Philip Runkel's book, *Casting Nets and Testing Specimens* (Praeger, 1990), and I believe I understand his arguments about why multiple regression (MR) and other "relative-frequency-based" analyses based on group data cannot tell us much, if anything, about the functioning of organisms. Bill Powers has suggested that MR cannot even be profitably used for predictions about individuals. But everything I've learned about MR tells me that this indeed can be done.

Let's use a medical example. I can draw a random sample from some population of interest. I want to be able to predict blood pressure, so I obtain data on weight, per cent body fat, smoking, dietary habits, and perhaps even have each person fill out some questionnaire relating to stress. I can then do an MR which will provide me with a weighting of independent variables best predicting the dependent variable, blood pressure. If I get a high multiple correlation (r-square), I can then use this regression equation to predict the blood pressure for someone whose blood pressure I have not yet measured, but for whom I know the values of the independent variables. Of course, this person must be a member of the original population. I know that I will not be able to predict his or her blood pressure exactly, but if I do the statistics right, I should be able to attach probabilities to ranges of values, i.e., establish confidence limits for his or her predicted blood pressure.

I realize control theory says that such a study does not necessarily tell me anything about what causes blood pressure to rise or fall in people in general or in any individual (Runkel's book makes this point well). And I realize that it would probably be easier just to measure the blood pressure instead of predicting it (it's a poor example from that viewpoint). But why can't I use this technique for predicting for individuals?

Bill Powers: Gary, I wish Phil Runkel were on the net, but I'll try to defend my statement without an expert's help (with the usual risk of getting it all wrong).

My basic argument is that you could use the MR method to predict the average relationship of various factors to blood pressure in *another group of the same size from the same population,* but you have only a tiny chance of guessing right about any individual from either the old group or the new group. I won't even get into the problem of how you know you're drawing from the same population, a subject on which Phil Runkel has some cutting remarks.

The reason for my opinion is that the "independent variables" (or the factors you get from them) are not known to be physically causative of high or low blood pressure: they are simply associated by experience with blood pressure. When you use multiple tests, the intuitive thought would be that getting at the relationship from many independent angles ought to improve your ability to predict for a single person. I'm quite sure that it doesn't, but let's see if I can work up a coherent justification for saying that.

If you looked at the raw data from the tests, you would find that some people high in each factor had high blood pressure, while others did not. Let's be generous and suppose that 80 per cent of the people in the original group who scored high on each factor actually had high blood pressure.

If that is true, and if 1000 people participated in the study, 800 of them who scored high on the first test had high blood pressure, while 200 of them didn't. We now have 800 people left whose scores on the first test truly indicated high blood pressure, or seemed to. Now we give the second test. After this test, we have 80 percent of 800 or 640 people who indicated high on both measures and did indeed have high blood pressure. After the third test we have 512 people left, after the

fourth test, 410 people left, and after the fifth test, 328 left. Therefore, out of the original 1000 people, only 328 who scored high on all five tests proved to have high blood pressure. So if you give all five tests to an individual, and the individual scores high on all five measures, the chances of high blood pressure are about one in three. In other words, you'd be safest in betting that a person who scores high on all five "indicators" does *not* have high blood pressure.

Why this counterintuitive result? I think the reason is that we confuse association with causation. If it were true that, for example, a high load of body fat *physically caused* high blood pressure, then there would be no way for an otherwise normal person to have high body fat and not have high blood pressure. The only room for error would be in measuring body fat or in finding the right curve relating body fat to blood pressure. A deviation would basically be a measurement error, not a matter of chance membership in a population. Body fat would amount then to a measure of blood pressure.

In the same way, each other measure, if it were truly a physically causative factor, would also amount to a way of measuring blood pressure, and you would expect using these multiple measures to reduce the error of measurement. But these measures are *not* measures of blood pressure. They're not "measures" at all. They are simply factors that common sense tells us might have something to do with the matter. That being the case, we are not perturbed by finding that a person who has high body fat happens to have low blood pressure. If there were a physical chain of causation involved, we would be very perturbed indeed to find our measuring instrument suddenly indicating the wrong way. This is the difference between physical or model-based measurements of relationships and statistical inference of relationships. There are no physical principles operative in a statistical inference, and of course the only model is pretty elementary.

This misuse of statistical "facts" is encouraged by the habit into which most empirical scientists fall, which is to say not that "80 per cent of people with high body fat have high blood pressure and 20 per cent don't," but that "high body fat predicts high blood pressure." The customary wording implies that this is *always* true; this makes the factor look like a physical cause. Just look at any summary of findings in a statistical study. Does it tell you the chances that a given person does not show the effect or shows the opposite effect? It does not. It says "A is associated with B." In *everybody.* That is why you expect the result to apply to *anybody.*

In truth, nobody knows why, in some people, the reference level for blood pressure is set to a high value. Nobody knows, because all the big research money goes into statistical studies instead of into developing a competent model of how the human system works. I wouldn't recommend that we just do studies of physical causation, because I don't think that's how you come to understand a system, but I do recommend that we study the ongoing networks of relationships that constitute a functioning body and brain. Until we do that, none of this statistical crap is going to do much good for an individual who has to make decisions based on an N of I and gets only one chance to bet right.

I smoke, eat eggs and bacon, weigh about 30 pounds too much, don't get a lot of exercise, and have, at last measurement, a blood pressure of about 125/80. Just a statistical fluctuation, that's me.

One last consideration. I think that studies involving very large numbers of people, like the cholesterol studies, are probably worse indicators of an individual's characteristics than studies involving only a few subjects. My reasoning is that large studies are necessary only when the effect is very small—when the number of people showing the effect is only slightly larger than the number not showing it. If 80 or 90 per cent of subjects in a pilot study showed the effect, why on earth would anyone then expand the study to huge numbers of people? In a large study we are justified in suspecting that the split is not 80/20, but more like 51/49. The numbers are needed to get statistical significance out of an effect that's just barely there.

In medicine, the practices are even worse than that. I recently saw a glowing report on a drug which statistics proved to help 16 per cent of the people who took it. In other words, 16 per cent got better and 84 per cent didn't. I think that result leaves room for a lot of questions about just why those people actually got better, and what effect the drug had on those who didn't. This sort of mindless application of statistics goes on all the time. Remember that the next time someone tries to get you to pop a wonder pill (unless you have as many chances to try to get well as necessary). Ask for a warranty.

One more last thought: Suppose it happened that all five tests together were a very good predictor of high blood pressure. Is that any reason to think that reducing all five factors would reduce the blood pressure? This is another elementary logical error: thinking that an implication works both ways. Suppose that the blood pressure is high for the same reason that leads to high values of these other factors. Statistics says *nothing* about causation.

See my paper in the *American Behavioral Scientist* issue edited by Rick Marken (September/October 1990) for a demonstration of how a statistical analysis can yield an apparent relationship that actually goes the wrong way.

Gary Cziko: OK, Bill, here's some thought data: 0 indicates low on a factor, 1 indicates high; A through D are independent variables, Y is

dependent (blood pressure):

Subject	A	B	C	D	Y
1	0	1	0	0	0
2	0	0	1	0	0
3	0	0	0	0	0
4	0	0	0	1	0
5	1	0	0	0	0
6	0	1	1	1	1
7	1	0	1	1	1
8	1	1	1	0	1
9	1	1	0	1	1
10	1	1	1	1	1

Note that only 80 per cent (4/5) of those scoring high on A have high blood pressure; the same holds for B, C, and D. The one person who is high on all four independent variables has high blood pressure, the one low on all four independent variables does not. In addition, *everyone* scoring high on at least four out of five independent variables has high blood pressure, and no one who scores low on four out of five has high blood pressure. And so perfect prediction is possible with these data. Of course, things might not be so pretty when I get another sample, since this sample is very small. But if with a larger sample I still don't get individuals deviating from this pattern, I would feel pretty confident in predicting an individual's blood pressure based on his or her characteristics as defined by the independent variables.

Looks pretty good to me.

Bill Powers: For those finding my statistics hard to swallow: If you propose that each of five conditions is associated with high blood pressure, but have no model and no knowledge of the physical means by which each condition has its effect, you can only assume that each association is independent of the four others. There is no a priori reason to assume that testing high on one measure predicts testing high on another.

The upshot is that you must assume that, on each test, the distribution of people measuring high on that parameter is independent of the distribution for any other parameter. When you isolate the 80 per cent who scored high on a given measure *and* had high blood pressure, you have not thereby isolated those who will score high on any other test (as Gary's example assumed). You have only eliminated those who tested high on one test but showed low blood pressure. Among those who are left, however, only 80 per cent, again, will score high on another test *and* have high blood pressure. Having high blood pressure is not sufficient to predict how a person will score on a test that seems to predict high blood pressure. It is a common error to suppose that this is true, but it's not. Implications don't work backward, as I said. Getting on a train at the next-to-last station implies—very reliably predicts—getting off at the last station. But if you see a person getting off at the last station, this does not imply that the person got on at the next-to-last station.

Finding, through factor analysis, a factor related to blood pressure *reduces* the credibility of an individual measure having a causal role. The hidden factor correlates better with the dependent variable than do the individual measures, which indicates that the hidden factor might be having a direct effect on the dependent variable and a lesser effect on the initially proposed independent variables. Of course, the hidden factor could itself be a side-effect of an even more important cause that also affects the dependent variable. It's simply a mistake to assume that an association implies a dependent and an independent variable. The fact that it's commonly assumed doesn't make it right.

Suppose that a person were in conflict. This can mean being physiologically prepared to act but not being able to carry out the actions that would normally "use up" the prepared state. One consequence of this state might be an elevation of the reference level for blood pressure. Among other consequences would be the tendency to measure high on stress, to seek comfort in good food or to gobble fast food, to be unable to act vigorously (a direct effect of conflict that equates to "little exercise" and thus being overweight), and so on. So it is not at all farfetched to propose a common reason for the high blood pressure and for the high scores. When that is the case, lowering the test scores will have no effect at all on the blood pressure.

Phil Runkel has laid out the circumstances in which statistical studies are appropriate and meaningful. These do not include the prediction of individual behavior or the exploration of natural laws. You learn through statistics what masses of people actually do, but you learn *nothing* about the underlying processes that lead to individual behavior. Statistics, when applied to individuals, is not science. It is organized superstition and systematized prejudice. It gives the illusion of knowledge, which is probably worse than ignorance.

Gary Cziko: Bill, please note that I have read (several times!) Runkel's book and find his arguments quite convincing that group statistics do not necessarily tell you anything about how individuals function. I do not, however, understand the part of Chapter 8 on regression, and that is perhaps what started all this. While statistics might not tell you much of anything about how people function, I still suspect that they *can* help in certain types of predictions about individuals.

You say: "If you propose that each of five conditions is associated with high blood pressure, but have no model and no knowledge of the physical means by which each condition has its effect, you can only assume that each association is independent of the four others. There is no a priori reason to assume that testing high on one measure predicts testing high on another."

But if one has no model, why does that force one to assume independence among the four independent variables? In fact, we know in the behavioral sciences that everything often seems to be at least a little related to everything else, so why assume independence? Your "upshot" is suspect if the assumptions are suspect.

Regardless of train riding practices, correlations, as I understand them, work both ways. If there is a 0.7 correlation between percent body fat and blood pressure, then there is a 0.7 correlation between blood pressure and body fat. Now, the regression line (and equation) will be different depending on which way you go, but that is only because the variances of the two variables are not likely to be equal.

Bill, you talk about causality; I'm only talking prediction. Why do we need causality for prediction? There is probably a positive correlation between shoe size and reading ability among elementary school children. This doesn't mean that kids use their feet to read; the causal factor is more likely to be something like age (but even this alone will not cause better reading skills). But as long as there is a nonzero correlation between shoe size and reading ability, I can use shoe size to make a prediction about reading ability that is better than a prediction made without knowledge of shoe size. Being ignorant of shoe size, I can only predict the mean of the group with a standard error of estimate equal to the standard deviation of the reading scores. With shoe size, I can reduce this error of prediction so that it is *less* than the standard deviation of the reading scores. And if I have a perfect correlation, there is no error at all. Why I do I need to find causal factors to make predictions? The daffodils coming out of the ground do not cause Easter. And yet when I see them growing, I can predict that Easter is not far away.

You also say that "through statistics... you learn *nothing* about the underlying processes that lead to individual behavior." I agree, but that still doesn't make it clear to me that statistics is useless for predicting aspects of individuals. Insurance companies would all probably go broke if they didn't use statistics for these purposes.

Let's try to keep away from the "understanding specimens" argument. Runkel does this well, and anybody can read his book. However, if we can effectively dismantle the individual prediction rationale for statistics, this will really pull the rug out from under the social (including medical) sciences, and this would indeed be great fun. I'm really on your side (I think), but I'm not yet convinced. Please be patient.

Mark Olson: Bill, like Gary, I understand that we want to keep away from an "understanding specimens" argument, and that the idea in question is whether statistics has any predictive value. Gary's argument makes complete sense to me, so I am anxiously awaiting your rebuttal, and like Gary, I hope you are right. If I may make a trivial request, could you stick with the shoe size and reading ability example—this is the example I use in my educational psychology class to teach the concept of correlation—the train example confuses things. Thanks.

Chuck Tucker: The important point Runkel makes that can get lost in these discussions is not that statistics is bad or dumb or worthless, but that it is a tool that can be used for some specific purposes but not for others. Statistics is a very weak tool to make sense out of what people do—some statistics make sense or are useful, but others are not as useful. It is like using a hammer to put a screw into wood—you can do it, but it will mess up the screw head and the wood and probably won't hold very well. This is the case with most of statistics *if* you are concerned with how the human being works; its use is very limited and might in fact be harmful to your understanding. The argument is pragmatic in the best sense of the word.

Rick Marken: Bill says: "You learn through statistics what masses of people actually do, but you learn *nothing* about the underlying processes that lead to individual behavior." Gary replies: "1 agree, but that still doesn't make it clear to me that statistics is useless for predicting aspects of individuals. Insurance companies would all probably go broke if they didn't use statistics for these purposes."

I think we are getting philosophical here—so I'll jump in blindly. I think there is nothing harder for people to understand than the point you guys are trying to make. People make individual decisions based on mass data all the time, and they consider it very reasonable. In other words, they are predicting aspects of individuals (themselves) based on statistical data. Lots of behavior is done solely because the statistics imply that you, as an individual, are more likely to be X rather than Y if you do Z. Even a somewhat rational person like me bases some individual decisions on what the statistics say.

Gary is right about prediction and statistics—my prediction that a person will have value X on a particular dimension is better (smaller RMS error over predictions) if I know some predictor variables and the equation relating them to values on the dimension of concern. But Bill is right because this kind of prediction is of no use for an individual. Accuracy is defined over prediction occasions, and an individual is just one occasion. So it is perfectly reasonable, I think, for an insurance company to charge me more for life insurance if I smoke. But it is silly

for me not to smoke based on statistical data. I am not a likelihood. I'm just me, once. I can only base my attempts to control things (and that is what you are trying to do when you base life decisions on statistical data) on what is happening now, not on what might happen on repeated samples of my life. I can control my insurance premium, my attractiveness to those I care for, and other things by not smoking. But I have no way of controlling how long I live or whether I get cancer. Those things only happen once, and there is no evidence that they can be reliably controlled by individuals' variations in their smoking behavior (individually—I know that, statistically, non-smokers do better on these things, but this is irrelevant to me individually).

Maybe control is the operative concept here (not statistical control, but perceptual control). Statistical evidence gives no evidence of an individual's ability to control variables. Statistics on smoking tell me nothing about how I, individually, can control cancer in myself. People often point out the individual irrelevance of smoking statistics by pointing to folks like George Burns. This irrelevance does not mean that smoking might not be bad for many people—eating candy is bad for some people, too. Also, there are probably perceptual consequences of smoking that can be controlled by cutting down or stopping. If people want to control these consequences, then controlling their smoking might be tried. But trying to control variables by basing individual actions on statistical data is just silly. People can only control perception; controlling imagination doesn't help anything. In fact, spending a lot of effort controlling imagination is called neurosis, isn't it? The applicability of statistical data to any particular individual is imaginary, so controlling individual behavior based on its imagined statistical consequences seems to me like neurosis.

Joel Judd: I got the impression from Gary's last comments that he was looking for some logico-mathematical reasoning for arguing against inferential statistics, instead of the "specimens" argument. But it seems that all one needs to do when contemplating the use of a tool—e.g., statistics—is ask, "What do I want to use this tool for?" One doesn't have to delve into the physics and whatnot of screws and screwdrivers and hammers to figure out that a hammer doesn't put in screws well (Chuck's example). Every statistical tool has some mathematical assumption(s) underlying it, delimiting its use. What else should one have to say when defending a perspective such as Runkel's? I want to know *why* someone does X. Group statistics can't tell me.

Mark Olson: Rick, I think I follow your smoking/cancer example. But I first need a distinction to be made before I feel I truly understand. We say that smoking and cancer are correlated. We also say that children's feet size and reading ability are correlated. Yet I see these as being correlated for very different reasons. In the former example, smoking "could" cause cancer, while in the latter example, size and ability cannot be causally related. It seems that this difference should have some importance in this whole issue, and I can't quite seem to articulate what that might be any insights?

Bill Powers: Rick says: "Lots of behavior is done solely because the statistics imply that you, as an individual, are more likely to be X rather than Y if you do Z. Even a somewhat rational person like me bases some individual decisions on what the statistics say."

Statisticians like to point out that people who use informal statistical analysis as a basis for choosing behavior don't do very well at it. I bought two lottery tickets because the pot was $60 million on Wednesday. A rational analysis shows that if I had bought *all* of the tickets, I would have been *certain* to lose something like $20 million (or some big number). So the optimum number to buy, considering that the $2 could have been spent on a hamburger which would certainly do me some good, was zero.

But Rick's point is well taken. It reminds us of what statistics is all about: trying to make predictions about what will happen on the basis of what has happened. This is all people could do prior to science: they didn't know how to figure out the underlying processes so they could predict what is going to happen without having to remember and analyze what has happened. Once you have a workable idea of the inner organization of any system, you can predict what it will do even under circumstances that have never happened before. Of course, you have to study what happens in the world in order to find a good model. But once you have the model, you predict from it, not from average past behavior. The record of physics and chemistry shows that this approach is far superior to merely watching behavior and assuming that the future will be like the past.

When your motorcycle starts making a funny tapping sound, there are two ways to fix it. One is to try to remember what the mechanic found the last time that sound happened and replace the same part. The other is to understand how the engine works, inside, and figure out that *this* time it's the tappets. What was wrong the *last* time is then irrelevant. Of course, if the previous trouble was also the tappet adjustments, then this time you should *not* merely adjust the tappets. First, you should figure out why the setting isn't holding. You have a different problem, and the tappet maladjustment is only a symptom of it.

Tom Bourbon: Concerning the recent discussion about statistical predictions, there was an observation that there is a difference between

correlations such as the one between smoking and lung cancer, and the one between shoe size and reading skill. That is true. A correlation between two sets of numbers means nothing more than that the positions of individual cases on one measurement scale resemble their positions on another scale. The equations used to calculate the degree of correlation could care less where the numbers came from or what they mean. That is as it should be, and that is one reason statistical analyses alone cannot reveal information about individuals.

However, when used in the context of research driven by a theory that makes bold predictions (i.e., specific, quantitative, falsifiable predictions), correlations can provide strong evidence about causal relationships. In the case of correlations found in control behavior, however, the correlations go counter to what most behavioral scientists have come to expect. For example, if a person is controlling a variable that is subject to independent disturbances, the actions of the person will be essentially *uncorrelated* with the value of the variable the person is controlling, but will be *highly negatively* correlated with the net disturbances acting on the controlled variable. To an uninformed observer, the person's actions will appear random, and the person's control over the perhaps unchanging controlled variable will go unnoticed.

In tracking studies such as those used by some of us who do control-theory modeling, the correlations between 1800 pairs of values of positions of a control handle and of values of the net disturbance on a controlled cursor are as high as -0.998. Of course, with $n = 1800$, no test of statistical significance is needed to know that the person moved the handle to negate the effect of the net disturbance. To do a statistical test of significance on data such as those would be utterly ridiculous.

In tracking data, the correlation between positions of the cursor and of the handle varies around 0.0, but it can be as much as +0.2 or -0.2. With $n = 1800$, those correlations are highly statistically significant; but of course they are totally meaningless.

In more traditional psychological research, correlations can provide some grounds for prediction, but only if the assumptions and requirements of the statistical procedures are met. That was one of Phil Runkel's major points in his book. Phil did not reject the "method of relative frequencies," *as* he identified traditional research designs and statistical analyses. But he did rightfully and masterfully show that those methods cannot work if one uses them to gather information that lets one make firm statements about individuals.

An excellent example of the problems encountered when people try to use statistical evidence to make statements about individuals can be found in R. M. Dar, D. Faust, and P. E. Meehl, "Clinical vs. Actuarial Judgment," *Science 243,* 1989, 1668-1674. The authors summarize the now sizeable literature which reveals that nearly any simple-minded actuarial procedure can out-diagnose nearly any practitioner who relies on "clinical judgment." Those results are telling. But the authors make another major point: even the best actuarial procedures are not very good. The actuarial procedures produce validity coefficients a few per cent higher than those produced by clinicians acting on professional judgment alone. The correlations between diagnoses and confirmed "pathology" are in the 0.20-0.50 range, which is the range one typically sees in the literature for the behavioral sciences. It appears that the clinical psychologists, burdened as they are with the "scientist-practitioner" model under which they train, do about as well as the behavioral scientists when it comes to identifying relationships—and neither group does very well.

Dar, Faust, and Meehl also draw a distinction between the state of affairs in clinical diagnostics and that in science, where access to a strong, corroborated model gives the edge to the scientist over actuarial procedures. The reason, of course, is that the scientist has an understanding of *causes.* Those who rely on actuarial procedures labor under the handicap of ignorance about causes—or else they act as though they understand causes, as when they assume causal relationships among the variables that enter into a multiple regression equation.

Gary Cziko: Reading some of Tom's comments, I get the feeling that the issue we are discussing here all reduces to the notion of individual differences in reference levels (internal standards). If everyone in *a* population had the same reference level for some perception, then we would get nice group correlations between disturbances (which would look like stimuli) and behavior which (it seems to me) *would* tell us something about the workings of individuals. However, individual differences cloud this relationship, so the only way to get at it is to examine individuals separately and then see what the invariances are at a more abstract level.

As far as I know, all strips of copper or containers of oxygen are basically alike. We can push and pull on them and send electrical currents through them and see how they behave without worrying about differing internal standards. And this is what traditional psychological methods do with people. Maybe psychology has forgotten why people in experiments were originally (and are still today) called "subjects." For the type of research usually done in the behavioral/social sciences, aren't they really treated as objects?

Tom Bourbon Gary Cziko has remarked that the behavioral and social sciences treat people like objects. That is true, not just of their treatment of people, but of living things in general. It is as though behavioral and social scientists expect living mice to "obey" the same causal

laws as the obliging "creatures" whose tails plug into computers, and who jump at our merest touch.

Nestled among the ever-increasing contents of my CST bookshelf is Lewis Carroll's *Alice's Adventures in Wonderland & Through the Looking-Glass.* Carroll understood the distinction and expressed it eloquently in the chapter on "The Queen's Croquet-Ground." I believe Carroll's message is one every control theorist understands—one every behavioral and life scientist should learn:

> Alice thought she had never seen such a curious croquet-ground in her life: it was all ridges and furrows; the croquet balls were live hedgehogs, and the mallets live flamingoes, and the soldiers had to double themselves up and stand on their hands and feet, to make the arches.
>
> The chief difficulty Alice found at first was in managing her flamingo: she succeeded in getting its body tucked away, comfortably enough, under her arm, with its legs hanging down, but generally, just as she had got its neck nicely straightened out, and was going to give the hedgehog a blow with its head, it *would* twist itself round and look up into her face, with such a puzzled expression that she could not help bursting out laughing; and, when she had got its head down, and was going to begin again, it was very provoking to find that the hedgehog had unrolled itself, and was in the act of crawling away: besides *all* this, there was generally a ridge or a furrow in the way wherever she wanted to send the hedgehog to, and, as the doubled-up soldiers were always getting up and walking off to other parts of the ground, Alice soon came to the conclusion that it was a very difficult game indeed.

That's life!

Mark Olson: Tom said that it is true that we can't compare correlations of smoking and cancer to correlations of feet size and reading ability. But this didn't answer my question about what *is* that difference between these two examples. What Tom wrote was helpful, but it didn't answer my question (at least not directly). Any comments?

Tim Cutmore: Would we say that smoking causes cancer if it were found that all (or perhaps just "most" would do) people who smoke also were exposed to Z-rays when children, and the Z-ray exposure induced the degree of desire to smoke? *And* it was also noted that Z-rays have a dose-related latent effect in causing cancer (amounting to accounting for 99 per cent of the variance in lung cancer!)?

In this case, we would have a superordinate variable which caused both smoking and cancer (vis-a-vis age reading experience -4 reading ability and age → foot size; age is the superordinate variable). The difference in what we believe appears to depend on perceiving the relations of the dependent variable (reading ability or cancer) to a superordinate variable (or not).

Izhak Bar-Kana: As the name says, a correlation only shows that some relation apparently exists between two different things, for example when one is large, the other is mostly large, etc. It doesn't say if one is the cause of the other, if one precedes the other, or not. The difference between the smoking and cancer vs. feet size and reading ability examples is only in the *additional* knowledge or assumptions involved. People have assumed for a long time that smoking might lead to cancer, and the correlation shows that, statistically, there might be something here. If the correlation is all you have, you might assume that cancer *is* the cause of smoking, or that both have some common cause.

In the second case, one only starts measuring and finds some statistical relationship between feet size and reading, and then tries to make something out of it. But one then needs more: assumptions, revelations, or some discovery that would prove/disprove that the statistical result is relevant.

Tom Bourbon: Mark has convinced me that I did not make my point clearly. One may assert that *any* two (or more) sets of correlations are comparable. Nothing in the procedures for calculating correlations rules out any use to which a person might put the results of the calculations. As I understand it—and I am not a skilled mathematician—computational procedures of all kinds are blind as to the origins of, and the meanings of, the numbers that are fed into them. And they are equally blind to the meaning of the results. Meaning and significance are in the eyes of those who behold the results, not in the results.

That is why Tim is free to tell us that his hypothetical Z-rays really do explain the variance in occurrence of lung cancer, and that the putative association with smoking should be put aside. For some reason, I doubt that Tim would do that, not because of anything in the rules by which one plays the correlation game, but because such an argument would not sound plausible to the professional community. Too many other things people believe they already know would be in jeopardy—and I do not mean that in a trivial sense. The assertion of as-yet-unrecorded rays that can play a major role in a prevalent medical problem would stretch at the boundaries of science. (Goodness knows, the boundaries need stretching from time to time—ask any control theorist who tries to publish!) Unless Tim could offer clear evidence that passed the scrutiny of scientists, and, more importantly, of good professional magicians, his

assertion would sound too much like the N-rays that Blondlett and his associates could see in France, early in the century. (Heard much about N-rays, lately?)

Which is merely another way of saying what I did in my last post: the smoking-cancer association *seems* more plausible than the shoe size-reading ability one. It is all in the sense of how the assertions fair with (fit with, form a nice figure with) the other things we know. And that has nothing to do with the numbers, per se.

Wayne Hershberger: Tom, your reference to the article by Dar, Faust, and Meehl reminds me that Meehl published an article within the last three years—in one of the APA journals, I think—comparing the methodologies of the hard and the life sciences. His arguments are consistent with, if not identical to, Bill's emphasis on "model building" and Phil's concern with "testing specimens."

Bill Powers: It seems to me that there are three topics concerning statistics needing separate discussion here. One is the question of causality; another is the question of applying a statistically obtained regression line to individuals; the third is the quality of the data on which the analysis is based.

On causality: I think we are all agreed that correlations do not reveal causation. Causation could run backward to the intuitively assumed direction (incipient cancer causes a desire to smoke), could result from a superordinate cause (Z-rays cause both a desire to smoke and cancer), or could be symptoms of some other process (smoking is a normally successful attempt by the system to counteract the onset of cancer—what percentage of smokers don't get cancer?). No information about these possibilities or any other comes out of a statistical study.

On the application of statistical relationships to individuals: Large studies involving many individuals yield a scatter of data. The common assumption is that this scatter is due to uncontrolled environmental variables. But an even stronger assumption is that measuring many individuals under varying conditions is the same as measuring *one* individual under varying conditions: in other words, all individuals in the population are alike and interchangeable.

Even granting an underlying justification for associating a statistical relationship with a causal relationship (for example, having a model whose properties agree with the statistical results), the statistical relationship (regression line) for a population might have nothing to do with the quantitative relationships inside each individual that link individual behavior to the independent variable(s). I showed in my *American Behavioral Scientist* paper that individual differences can account for the slope of a population regression line, while inside each individual the relation of behavior to the independent variables has a slope opposite to that of the population.

Also, confidence levels do not apply to individual measures. If p is less than 0.05, this means only that there is less than one chance in 20 that the correlation observed in the aggregate data is due to a chance fluctuation in variables that are actually unrelated. If the entire study were repeated 20 times, only once would the correlation measure zero. Is there any way to calculate the chance that an individual deviation from the mean is due to random departure from the population mean effect rather than a random departure from the condition of no relationship? It seems to me that this would be like the effect of an individual not actually being from the same population (where a population is defined as people with identical properties). What is the chance that an individual is not a member of the assumed population? Isn't it the product of the probabilities that the person will test positive on each indicator of population membership?

On the quality of the data: I've said that a correlation of 0.8 looks terrible on a scatter plot. By this, I mean that if you take the regression equation $y = ax + b$ as a prediction of the value of the dependent variable y from a known value of x, the mean error seems to be very large in relation to the range of predicted values of y. Can someone who is fluent with statistical calculations figure out the general relationship here? Given such-and-such correlation and a Gaussian distribution of errors, what is the RMS error of prediction of a single measure from a regression line?

There's another way to view data: in terms of signal-to-noise ratio. This is the ratio of peak-to-peak fluctuations of a signal to RMS noise, where signal and noise are defined in different frequency bands. For ordinary purposes of transmitting quantitative analogue data such as an audio waveform, a signal-to-noise ratio of 6 to 1 is barely tolerable; for high-fidelity purposes, it should be at least 80 decibels, which is a ratio of 10,000 to 1 in amplitude terms. Ordinary meter readings useful for diagnosing electrical system problems need a signal-to-noise ratio of 30:1 or greater (3 per cent accuracy). This latter signal-to-noise ratio is about what we get in tracking experiments for the prediction error using a control-system model. The corresponding correlations are around -0.995. So a correlation of -0.995 implies the lower limit of acceptable noise in a physical measurement or prediction.

Of course, we sometimes have to accept worse signal-to-noise ratios, but the worse the ratio, the less believable is any statement that the theoretical model "predicts" the data. The question is, how bad a fit are we willing to accept while still claiming that the theory has any scientific usefulness?

I think that to claim scientific respectability, we have to insist on very

good fits of theory to data. The reason isn't aesthetics, but the need to be able to make deductions from multiple premises. When a scientific deduction depends on the truth-value of each of several premises that all have to be true for the conclusion to be true, the truth-value of the conclusion is the product of the truth-values of the premises. Four premises *and*ed together to create a conclusion, each premise having an 80 per cent chance of being true, result in a conclusion that has a probability of truth of 0.41. Sad but true.

Any science is built on a foundation of premises that have individually been checked experimentally and found to be acceptably true. A grown-up science is a large structure of logically related statements describing facts of nature. But what kind of science can you have when you can't string together four premises and come up with a conclusion that is probably true? The answer is: a very fragmentary one. You end up with isolated observations that have some small chance of being true in a narrow range of circumstances, but which have to remain isolated because the quality of the data is too low to permit building anything like a complex structure of knowledge.

My chief objection to the way data are analyzed and used in many of the life sciences is that observations of very low precision and repeatability are used just as if they were as precise and repeatable as those of physics. Deductions from premises are made just as if each premise had a truth-value of 1.0. There is an enormous gulf between the achievements of the physical sciences and those of the behavioral sciences. It directly reflects, I think, the difference between a model-based approach to nature, in which very high standards are set, and a statistical approach that provides an excuse for setting very low standards concerning what will be accepted as a true statement.

I have a feeling that we're starting to preach to the converted about statistics. Maybe there is some further point in doing this, and if so, why not? But I'm starting to get the itch to see control theory applied to some real problems some more. There are probably lots of people out there who are searching for applications pertinent to their interests, and who didn't intend to do statistical studies anyway. Of course a lot of participants on this net are in the position of having to develop an interface between control theory and conventional approaches, so maybe that's really what we're doing right now. As we're rejecting 90 per cent of the work being done by hundreds of thousands of well-funded investigators with loads of clout, however, it might be optimistic to think that these arguments are going to sway anyone who doesn't already accept them. There are limits to the vaunted open-mindedness of scientists, no matter what Carl Sagan says in *Parade.* We'll probably get furthest in the end by keeping our noses to our own grindstone as we've been doing for lo, these many years, welcoming those who are interested in joining forces with us, and otherwise ignoring the stuff we no longer believe.

Here is something I worked out, with the help of a mathematics manual, right after I wrote that I was tired of statistics.

Let X be the independent variable (for example, a disturbance acting on a controlled variable). Let Y be the dependent variable (a measure of the action that opposes the disturbance). Let r be the correlation coefficient calculated from N samples of X and Y. The regression equation is then Y = r * (sigy/sigx) * (X - Xbar) + Ybar, where sigx and sigy are the standard deviations of X and Y, and Xbar and Ybar are the average values of X and Y.

The ratio of standard deviations, output/input, is sigy/sigx. This is the scaling factor that represents the average amplification factor applied to the input to produce the output. That ratio takes care of any overall scaling needed to convert X into Y. The correlation coefficient can then range from -1 to 1, indicating the match in waveforms of X and Y (considering them to be time functions).

The standard error of an estimate of Y from X, according to my manual, is given by Sy = sigy sqrt(1 - r^2), or Sy/sigy = sqrt(1 - r^2).

The ratio Sy/sigy is the RMS discrepancy between the predicted and actual values of Y divided by the RMS variation in Y. Because we have pre-scaled the predicted value according to the ratio of sigy/sigx, a complete failure of prediction would make the standard error of the estimate equal to the RMS variations in Y: in other words, Sy/sigy = 1 means complete failure. A perfect prediction would give Sy/sigy = 0. I thus call this measure the "coefficient of failure."

We can now construct a table showing the relationship between the measured correlation of X and Y and the coefficient of failure defined as Sy/sigy.

Per Cent Prediction Failure	|Correlation Coefficient|
0	1.0
3	0.9995
5	0.9987
10	0.995
30	0.954
44	0.900
50	0.86
60	0.80
70	0.71
80	0.60
90	0.43
98	0.20
100	0.0

This is not like an error bar, because the average ratio of Y to X (RMS) has been removed in the calculation of r. A prediction error of 100 per cent is the maximum possible error, representing complete failure. At the low end, the prediction error is approximately the normal proportional error of prediction.

We can see that very high correlations, indeed, are needed to achieve prediction errors of only a few per cent. The error rises drastically as the correlation coefficient falls from 1.0 to 0.8. At a correlation of 0.6, there is an 80 per cent failure of prediction, and at 0.2, a 98 per cent failure (almost total failure).

The "failure of prediction" here is precisely the failure to predict the value of a single point using the regression equation obtained from all of the data points: in other words, the error in predicting individual behavior from the behavior of the aggregate. The significance of the larger errors must be judged not as if on a linear scale, but with the realization that a failure coefficient of 100 per cent means the ultimate degree of failure.

I think that this vindicates my informal estimate that correlations below 0.95 (failure coefficient 0.30) indicate that the model is too far off the mark to use in predicting individual behavior. An individual could actually show the opposite effect at this level of failure, over a significant range of values of the independent variable, with a probability of 50 per cent.

A more sophisticated treatment than I can produce would be needed to show the relationship between the failure coefficient and probabilities of various predictions. But I think the general picture is clear enough.

David Goldstein, I believe, told me that thinking of a regression line as a predictive model is not the normal way to use statistical results. But when mass statistics is used to predict individual behavior, that is exactly how the regression equation is being used. Isn't it?

Gary Cziko: Bill, you provided a very interesting table relating correlation coefficients to your "coefficient of failure." I've never seen this coefficient used before to give an idea of the error involved in predicting individuals based a group correlation coefficient; it would have been an ideal companion to Jimmy Carter's "misery index."

This coefficient is simply the ratio of the standard error of estimate (i.e., the typical amount of error for an individual prediction) compared to how much you would be off just using the mean value of the predicted variable in the sample. Simple enough. But to make sure you weren't pulling a fast one, I worked out a concrete example to convince myself. Perhaps others will find this useful as well, but it is really quite mundane stuff, and those of you who are wise about statistics should probably stop here.

To give a concrete example, I often get a correlation of about 0.60 between height and weight for the ca. 60 students in my (you guessed it) introductory statistics class. Imagine that the mean weight (X) of the class is 60 kg (132 lb) with a standard deviation (SD) of 5 kg, and the mean height (Y) is 160 cm (5′, 3″), with an SD of 10 cm. These figures, along with the correlation coefficient of 0.6, give a regression equation of height = 1.2 * (weight) + 88, so that someone weighing 60 kg would be predicted to be 160 cm tall (that makes sense—someone of average weight is predicted to be of average height).

Now, you say using this regression equation will give a whopping 80% error. Let's see how. Recall that the SD of height is 10 cm. Using the formula for the standard error of estimate (Sy), we get 10 * sqrt(1 - r^2), which, with r = 0.6, gives us Sy = 8 cm. This means that by using this regression, we will typically be off by 8 cm in making our predictions. Not using the regression equation at all, i.e., just using our knowledge of the group mean height (with no knowledge of weight), will give us an error of 10 cm (which is the SD of height). So it looks like you're right in that our typical error in using the regression equation is 80 per cent of what it would be if it were not used at all. Or, we could say that a correlation coefficient of 0.6 reduces error by only 20 per cent (should this be called the "coefficient of success"?).

Now, this example is a bit silly, because if I have both the height and weight of my students, and I want to know their height, I will not use a regression equation to predict their height—I will just look at the height I have already measured. If I were to be brave and predict the heights of my *next* class based on just their weights, my predictions would most likely be significantly worse than the original 80 per cent error, even if they were from the same population, whatever that means. Hmm.

Only two problems remain. First, why is it that statisticians always talk about r-square, the misnamed "coefficient of determination"? They would take my r = 0.6, square it to get 0.36, and then say that variation in weight explains 36 per cent of the variation in height. This 36 per cent is not great, but it does look better than a coefficient of failure of 80 per cent or coefficient of success of 20 per cent. I've yet to figure out how r-square relates to these two new quite pessimistic indices of the predictive power of regression equations.

Second, you have been arguing that adding in more predictors makes the error even worse. But typically, adding more predictors does increase the absolute value of the correlation coefficient (multiple r), which, by your own table, *reduces* the coefficient of failure. I can't see how your argument holds, unless you get into problems of sampling and cross-sample validation.

Mark Olson: I just wanted to thank those of you who explained the difference between the smoking/cancer and reading/feet situations. I think the statement that "there is no difference between the two except the assumptions one brings to each" is what "enlightened" me.

Gary Cziko: As a follow-up to my last post, I just discovered that Bill Powers' "coefficient of failure" does appear in one of my statistics books, where it is called the "coefficient of alienation" and is calculated as k = sqrt(1 - r^2). It would be interesting to see how many statistics books even mention this coefficient.

I would prefer to call it the coefficient of "uselessness," since it tells how useless a predictor (or group of predictors in multiple regression) is in predicting the Y of an individual.

I recently had a colleague give a presentation showing how, using all sorts of measures in the right combination, he can obtain a multiple r of 0.5 in predicting children's adjustment/happiness in school. He justified this by saying that this is about the best you can get in the social sciences. I wish I had been able to tell him that his findings were 86 per cent useless in predicting the adjustment/happiness of individual children.

Finally, it occurs to me that r-square looks better than k because the former does not depend upon making predictions for individuals but uses the rather more abstract concept of "shared" or "explained" variance.

Bill Powers: Gary, if I understand Phil Runkel's argument, what you gain by adding more predictors is more than offset by the smaller N in each group. If you had started with only one predictor (weight predicts height) in your class of 60, the N is 60. If you now add, say, grip strength as a second indicator of height, you now have at least four combinations of independent variables instead of one: high-high, high-low, low-low, and low-high. Each subgroup now has only 15 students in it. One-fourth the N means twice the standard error. Now, in order to fit the prediction, a person not only has to be heavier than average and taller than average, but also stronger than average. All you've done is to eliminate some of the heavier people who are taller. Even if the N in the high-high group is larger than in the other three groups, I think you always lose some predictivity. If you don't add any new people to increase N, it seems to me that you've just cut down the number of people who fit all the criteria: instead of just heavier and taller, they have to be heaver, stronger, and taller. I think that this is what Phil Runkel calls fine-slicing.

I don't know how to work this out mathematically. Can you do something analogous to what I did with the one-dimensional case?

My hunch is that the higher correlations found in multiple regressions are offset by the increased standard error, or more than offset. Higher correlation, but higher uselessness index—maybe.

As to "explained variance," individual measures don't have any variance, do they?

Gary Cziko: This continues the discussion about how group statistics are not very useful for making decisions about individuals.

Effect sizes have become a commonly used metric in educational research to describe the difference between an experimental group (e.g., new way of teaching math) and a control group (e.g., old way of teaching math). The effect size is the difference in means divided by the standard deviation. So if the standard deviation of the math test is 10, and the experimental group mean after treatment is 55 compared to the control group at 50, there is a 0.5 effect size.

For some reason, an effect size of at least 0.5 has become accepted as indicating that there is a practically significant difference between the two groups, hence the new method is better than the old. I wouldn't be surprised if a similar standard has become adopted in other areas, for example in medical research. One positive consequence of using effect sizes is that it gets around the problem of tiny differences being "highly statistically significant" simply because one has used large samples.

But let's see just how exciting an effect size of 0.5 really is. With two normal distributions whose means are separated by 0.5 standard deviation, we find that 31 per cent (almost one-third) of the individuals in the low group are actually higher than the mean of the high group. Also, an additional 38 per cent of low-group individuals will not be more than one standard deviation below the mean of the high group. This gives us a total of 69 per cent of low-group individuals which are either higher than the mean of the high group or not more than one standard deviation below the high mean. The same, of course, could be said conversely of the high-group individuals (69 per cent are lower or not more than one standard deviation above the mean of the low group).

An effect size of 0.5 does not seem very impressive in making predictions about individuals.

Chuck Tucker: The discussion on statistics is wonderful. I hope that all of you who teach statistics will incorporate these ideas in your courses and make it a point to catch those who claim they are not interested in individuals (that is the retort in my sociology department) when they try to use statistics to talk about them.

Bill Powers: Gary, I hadn't heard about "effect sizes." Half a standard deviation? Surely you jest. Do people ever actually replicate studies of

this sort? I approve of getting rid of statistical significance that's based mainly on large N, but is it an improvement to accept smaller N and also relax the meaning of significance even further ("practical significance")?

You say: "One positive consequence of using effect sizes is that it gets around the problem of tiny differences being 'highly statistically significant' simply because one has used large samples." Now you can get significance with tiny differences, even without using a large sample. It seems to me that someone is trying to recycle the garbage. How to do a bad experiment and still get it published?

Rick Marken: I want to just say "bravo" to all those involved in the statistics discussion. I don't think any conventional psychologists will be converted from the statistical to the modeling game, but it's nice to point out the problems for posterity, and for the unconverted who could contribute to the development of a science of life.

Martin Taylor: Gary defines "effect size" as the difference between the means of two distributions measured in units of the standard deviation. In psychophysics, this measure is called d′ ("d-prime"), and a d′ of 1 is taken as roughly what people mean when they say that there is a "threshold" effect. A subject will usually not claim to have detected an individual signal at a level giving a d′ much less than unity, but will usually claim to have detected an individual signal at a level giving a d′ appreciably greater than unity. Gary says that in educational research, an effect size of 0.5 is taken as practically significant, and he thinks the same is true of other areas. In psychophysics, the usual equivalent is an effect size of unity, which seems appropriate, given that the subjects in an experiment *are* working with individuals, and unity is roughly the d′ that separates conscious detection from non-detection.

Gary Cziko: Martin, could you provide a bit more information about what the psychophysical "effect size" d′ is as used in psychophysics? You say: "A subject will usually not claim to have detected an individual signal at a level giving a d′ much less than unity, but will usually claim to have detected an individual signal at a level giving a d′ appreciably greater than unity." Are you referring to a type of signal-to-noise ratio here? If this is analogous to the effect size in educational research, what are your two means, and what is your standard deviation based on? I suppose a simple example would help us non-psychophysicists to understand this.

I would guess that psychophysics should be of some interest to control theorists, since, as I understand it, it uses the method of specimens (one individual at a time to find invariant laws) in much the same way that control theory does.

As a follow-up to my previous post, I have constructed a table to show how various effect sizes can be used to make predictions about individuals in low" and "high" groups. The table assumes Normal distributions. I wouldn't be surprised if I made some typos or calculation errors here, but the numbers all go in the right direction, so there are no obvious errors.

In the definitions below, the words "low," "lower," and "below" can be interchanged with "high," "higher," and "above," respectively.

A = Effect size, (Xbar-Ybar)/SD
B = proportion of low scores higher than mean of high group ("surprises")
C = proportion of low group no more than 1 SD lower than mean of high group (low group scores as close to high mean as typical high group score is to high mean)
D = total of B and C (total proportion of low group scores easily construed as being part of high group)

A	B	C	D
0.50	0.3085	0.3830	0.6915
0.75	0.2266	0.3721	0.5519
1.00	0.1587	0.3413	0.5000
1.25	0.1056	0.2954	0.4010
1.50	0.0668	0.2417	0.3085
1.75	0.0401	0.1865	0.2266
2.00	0.0228	0.1359	0.1587
2.25	0.0122	0.0934	0.1054
2.50	0.0062	0.0606	0.0668
2.75	0.0030	0.0371	0.0401
3.00	0.0013	0.0215	0.0228

Column D is most informative (and most damaging) because it gives the total proportion of individuals in the low group who would not be out of place in the high group (or vice versa).

Note that at the "practically significant" (in educational research, anyway) ES of 0.5, more than two-thirds of the low group fit nicely into the high group (and vice versa). Even at a "whopping" ES of 1.00 (equivalent to a difference in mean IQ of 16 points, for example), this is still the case for half the individuals in each group. It is only when we reach a "mammoth" ES of close to 1.75 that this proportion drops to less than 0.25. An ES of 2.75 is nice, since then the proportion is less than 0.05. Has anybody ever seen one this big in the social sciences? Perhaps the difference in height between Pygmies and Dinkas in Africa.

Of course, all this looks even worse when we try to use findings like these to make predictions about *new* individuals who were not part of the original data, and who might or might not be considered part of the same population (whatever that means).

Bill Powers: Gary, if you think about publishing this sort of analysis, I hope you'll make the paper a comparison of what's good for education as opposed to what's good for the student. What's good for education is, of course, a good track record. What's good for each student is to be evaluated accurately, to be treated appropriately, and to learn successfully. What we've been doing in these posts is developing a way to show that the goals of educators can be met, while, in significant numbers of cases, those of students are not. It's no good to point out, as defenders of the present methods will do, that substantial numbers of students are treated properly. We have to focus on those who are misjudged by the statistics. Even with two standard deviations between group means, one student in six will be treated as if he or she belongs in the wrong group, according to your chart. In a class of 30, that's five people about whom the teacher will get the wrong idea. I don't think that this kind of misevaluation is harmless. It ought to be actionable on the basis of an implied warranty.

All this would be more convincing if we could come up with a way to apply control theory in teaching or testing that would work better than the present methods. Let's talk about it.

What I am hostile to is the misuse of group statistics. If you want to compare two methods or two tests to see which is "better" with respect to producing or measuring some aggregate phenomenon, statistics works fine. Just don't make the mistake of using the methods or the tests to evaluate individuals. Not unless your correlations are running 0.99 or better.

Gary Cziko: Bill says: "What I am hostile to is the misuse of group statistics. If you want to compare two methods or two tests to see which is 'better' with respect to producing or measuring some aggregate phenomenon, statistics works fine." But even this idea seems based on a linear, one-way view of causality which does not seem compatible with control theory. Much (if not most) of quantitative educational research is determined to show that certain combinations of inputs ("independent" variables) will give you certain outputs ("dependent" variables), and of course group statistics is used to try to do just this. Results have been rather dismal so far, but that just means that not enough variables were taken into account, or the measures were not reliable/valid enough, or the statistical analyses were not abstruse enough (structural equation modeling using a program called LISREL is the latest trend in statistical analysis). This is done, of course, in the hope that once the input-to-output links are known, teachers and administrators can better control the behavior (i.e., success, achievement, drop-out rate, motivation, etc.) of their students. It seems that even your statement seems to imply an input-to-output view.

Group statistics seem to be used in at least four ways in educational research:

(1) to tell us about the psychological processes/functioning of students;

(2) to make predictions about individuals;

(3) to find out what combinations of input variables (e.g., teaching method) cause certain patterns of output variables (e.g., mathematics achievement); and

(4) for polling (survey) research.

Runkel's book and your *American Behavioral Scientist* article do what I feel is a convincing job to debunk the first. Our recent discussion about individual predictions using correlations and effect sizes addresses what appear to be serious problems with the second. We are discussing the third now. It might be that only the fourth is a legitimate use (if we can figure out what a random sample is and don't worry too much about the problems that the Bayesians point out).

Fred Davidson: In response to the recent discussion of statistics, effect sizes, and what's-good-for-the-student (Cziko, Powers, and others), I recommend J. R. Frederiksen and A. Collins, "A Systems Approach to Educational Testing," *Educational Researcher 18(9),* 1989,27-32. There are many in educational testing who would love to see the downfall of norm-referenced epistemologies. Frederiksen and Collins propose an elegant new "validity" (= truth) of measurement: "systemic validity." They say: "Evidence for systemic validity would be an improvement in those skills [which the test claims to measure] after the test has been in place within the educational system for a long time." (p. 27)

In language testing, we call this "backwash"—the effect of testing on instruction. We backwashers believe that testing is the servant of successful learning. That's a concept that the quasi-scientific, clinical, detached, norm-referenced-measurement establishment seems to have forgotten. I like "systemic validity" better than "backwash," since the former elevates the concept to the level of a "validity"; there are about four validities taught in educational measurement courses: face, content, criterion (predictive and concurrent), and construct. Politically, that is a good idea.

Now to control theory: I suspect that control theory offers a way to further justify systemic validity/backwash. Isn't successful learning also a well-functioning control system?

Bill Powers: Gary, I said that group statistics can be used to compare methods or tests. You said: "But even this idea seems based on a linear, one-way view of causality which does not seem compatible with control theory. Much (if not most) of quantitative educational research is determined to show that certain combinations of inputs ('independent' variables) will give you certain outputs ('dependent' variables), and of course group statistics is used to try to do just this." We have to be careful about treating control theory as a dogma with which we must keep faith. If a lineal cause-effect model could predict individual behavior accurately, we would have to accept it as a contender against control theory. We don't really need to consider control theory when evaluating a cause-effect explanation of behavior. If we reject a cause-effect explanation, we should do so on the basis that it predicts poorly, not because it violates the precepts of control theory or because there's something that says cause-effect systems can't exist. This means we judge against standards of prediction. So where are we to set those standards? Is a measure that has a uselessness index of 60 per cent OK? Are we willing to accept the many wrong predictions that result from such a low standard? If so, then, as Rick Marken would say, go for it. It would certainly make life easy for those who need to publish regularly. But this isn't how you achieve real knowledge about nature.

What it all comes down to is a system concept. What kind of science do you want to mean when you call yourself a scientist?

Of course, I agree with you about the cause-effect approach. It isn't really even a model, because it tries to explain the output on the basis of the input without any idea at all of what goes on between them. That's truly just floundering around in the dark. You don't even know if the change of behavior isn't produced to counteract the effect of the input!

But I don't think that we've effectively debunked anything yet. How many conventional educators have called you up all weepy and apologetic and promised that they'll stop doing those bad things? I think we have to concentrate on finding something that works better, so it can be taught and used. That's the only thing we can offer that will change anyone's mind. Nobody will prefer a method that works worse over one that works better. Not for long.

Gary Cziko: Bill says: "If we reject a cause-effect explanation, we should do so on the basis that it predicts poorly, not because it violates the precepts of control theory or because there's something that says cause-effect systems can't exist." Yes, I basically agree with this, although I wonder what your reaction would be to someone who wants to show you a perpetual-motion machine (perhaps even one that can do work). I suppose you should ask to see if it works, although most of us wouldn't waste our time, since all we know about physics says such machines can't work. But, yes, control theory has nowhere near the status of the laws of thermodynamics, so we need to keep our eyes open to see what works.

Now, here's a concrete problem. I've been showing the "random" program which you describe in your article in *American Behavioral Scientist,* September/October 1990. One reaction I get is that a multiple regression (MR) could make good sense of these data if you included the reference level, cost, and wage variables. Something tells me that this is *not* the case, since this would still be an analysis of relative frequencies, not a test of individuals.

What I'd really like to do is to get the program to generate some data which I could try to analyze using MR (or better yet, give it to one of the many MR-whizzes around here) and see what could be done. So my two questions are:

1. Would it be possible and worthwhile to get a data matrix from this program for such an analysis?
2. Do you have any ideas about what MR analysis could reveal about such data? Could it find that reward is under fairly tight control and that costs and wages are disturbances?

I hope that those who are familiar with this article and know something about MR analysis will join in here.

Bill Powers: [In reply to a post by Peter Parzer, in the Department of Psychology at the University of Vienna.] It's now beginning to look as though we have been using the concept of correlation incorrectly in talking about our tracking experiments. When we speak of using a model to predict behavior, the independent variable used both for the model and for the real person is predetermined and exactly known (i.e., not a random variable). This implies that we shouldn't be talking about the "correlation" of the independent variable with the dependent one. Intuitively, we have realized that when you get correlations of 0.99 and up, correlation ceases to be a very useful measure and starts becoming a tool for making an impression on someone. The more useful measure is just the RMS error of prediction in proportion to the range of the expected value, which I have already referred to as the signal-to-noise ratio.

I'm not sure of this conclusion, however. Perhaps if I describe a basic experiment, Peter can tell us the right measure to use.

The task is for a person to use a control handle to keep a movable object on the screen aligned between two "target" marks. The position of the movable object (the "cursor") is determined by the sum of two numbers: one represents handle position relative to the midpoint, and the other is a time-varying disturbance generated by smoothing and

scaling a table of random numbers. When the target marks are stationary (the simplest case, "compensatory tracking"), accomplishing the task perfectly implies moving the handle in exact opposition to the disturbance, so the net effect on the cursor remains zero (which is the position between the target marks). The disturbance thus becomes an independent variable that predicts handle position.

The disturbing function itself is invisible, being applied inside the computer that runs the experiment. Stabilization of the cursor is not, of course, perfect; the cursor wobbles slightly up and down during a typical one-minute run. Its wobbles do not resemble the variations in the disturbance. The data consist of 1800 samples of cursor and handle position (the disturbance waveform is stored beforehand), or one set of samples every 1/30 second (more or less, depending on which computer is used).

The model used is that of a control system, which for this case is indistinguishable from a stimulus-response system except for the fact that the most obvious "stimulus," the cursor position, is continuously dependent on the "response," the handle position, as well as on the "independent variable," the disturbance waveform. In addition, all variables are continuous, instead of discrete as is usually assumed in stimulus-response analyses. The control-system model that we use most commonly also puts one time-integration into the output of the system. The output is a constant times the time integral of the deviation of the cursor from the target marks. For slow variations of the disturbance, this integrating model works only slightly better than a pure proportional model.

The subject and the model are both run with the same disturbing waveform. This enables us to find the value of the integrating constant (or gain of the control system for the proportional case) that makes the model fit the data best. Typical errors of fit are about three per cent RMS of the peak-to-peak excursions of the handle. Next, a new disturbing waveform is generated by the computer and the model is run using the parameters already obtained. This result is now a prediction of the way the subject will move the handle when the same new disturbance is applied during a "live" run. The errors of prediction are typically three per cent to five per cent of the handle excursion.

Predictions of the *cursor* position are not so accurate, because the cursor position represents the difference between the handle position and the optimal position called for by the magnitude of the disturbance at any given instant. For very slow disturbances, the cursor prediction error can be quite large—100 per cent RMS or more. But the more difficult the disturbance (so that stabilization errors become larger), the better the prediction, the RMS error dropping sometimes to 10 per cent of the cursor excursion.

Correlations of cursor position against handle position are probably meaningful, because unsystematic tracking errors are seen; these correlations are typically 0.2 or less (positive or negative), becoming smaller as the task gets easier.

We have also been calculating correlations between the momentary handle positions and the momentary magnitudes of disturbance. The disturbance variations, however, are accurately known, so this “independent variable” is not really random, although it is derived from a table of random numbers. In principle, because of the smoothing used to limit the speed of variation of the disturbance, some short-term prediction of the independent variable is possible (for this reason, some workers have proposed that control systems must contain predictors). Our model, however, does no predicting, and it works well enough that I don’t think we need to add such a feature to the model. But the question still remains as to whether the disturbance should be considered a random variable or a given variable. That’s what I’m asking Peter to think about, if this explanation of the experiments has given enough information to allow making a judgment.

[Following another post from Peter Panzer.] Before proceeding, I’d like to clear up the nature of the control-system model, as well as our way of using it for predictions. Let’s see if I can construct a diagram that will make the relations clearer:

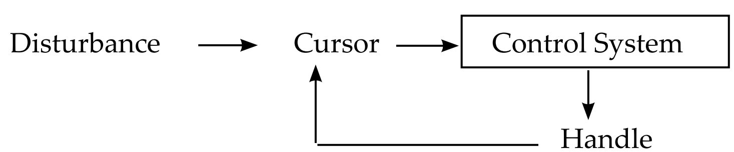

The effect of the disturbance on the cursor position occurs inside the computer; the disturbance itself cannot be seen by the subject except through its effects on the cursor. The handle also affects cursor position at the same time. So the input to the control system (visible cursor position) is not independent of the output (handle position). The true independent variable is the disturbance, a slowly and continuously varying waveform. The disturbance and the handle position affect the cursor at the same time, so cursor position depends jointly on the disturbance magnitude and the handle position. The behavior of the cursor does not reveal what either the disturbance alone or the handle alone is doing.

We can measure the cursor position within about one part in 350 to 480 of the maximum possible excursion on the computer screen (de-

pending on display resolution). We can measure the handle position (in my equipment) to one part in 4096 of the maximum possible handle excursion, give or take a per cent of nonlinearity. The disturbance values are known exactly. So we really aren’t talking about errors in measuring the input or the output, are we? We know what the input and output are with relatively high precision. The problem is to guess how the control system in the box is organized such that it produces the observed relationship.

If t represents the stationary target position (zero by definition), c represents the cursor position, and h represents the handle position, the simplest model that seems to predict well has the form: h’ = k * integral(c’ - t) * dt, where dt = about 1/30 second. The experimental apparatus is set up so that (exactly) c = h + d, where d is the current magnitude of the disturbance, and h is the current measured position of the handle.

The model is run by solving these two equations simultaneously via simulation, since d is not an analytical function of time. The variables c and h are given initial values, and then the disturbance is run through all its values while the values of c’ and h’ are computed over and over, yielding tables showing positions as a function of time. The subject is run by being put in the same relationship to the apparatus as the box labeled Control System, above.

The primes in the expressions designate the *predicted* values of c and h. Let *c* and h (without primes) represent the observed values (from a run with a real subject). We are then interested in the departure of c from *c’* and of h from h’. Generally, the RMS departures are enough larger than the errors of measurement of c and h that we can ignore those errors of measurement.

We can measure both the model’s and the subject’s handle positions with an accuracy of much less than one per cent. We take the subject’s handle position as the definition of zero error, and evaluate the model’s error of prediction by comparing its simulated handle positions with those of the subject over the course of the experimental run. It seems to me that this definition of prediction error is not arbitrary or model-dependent [as suggested by Peter].

What is arbitrary, of course, is the form of the model in the box labeled Control System. There is actually more in that box than is discussed here, because we have to be able to account for other cases—for example, the case in which the subject holds the cursor some fixed distance *away* from the target marks. We have picked the simplest model that accounts adequately for the data. More complex models can slightly improve the results. For example, by putting a time-delay of about 0.15 second into the model, we can halve the RMS prediction error. But it’s always possible that Mother Nature has put something else into the Control System box. All we can do is make our best guess and hope that more detailed data about the neuromuscular systems will help us to find a still better model. But as the simplest model leaves only about three to five per cent difference between model and reality, we aren't going to gain much more accuracy.

There are two steps in making a prediction. First we match the model to the behavior as well as possible by adjusting k in the equation above. Then we generate a new waveform for the disturbance (when we're fussy we require that it correlate less than 0.2 with the former one) and use that to make a predicted run, with the previously found value of k (the only adjustable parameter). The predicted handle waveform will be different from before because the disturbance waveform is different. Finally, the (same) subject's behavior with the new disturbance waveform is recorded and compared with the prediction. This latter step, in which the model is used first under new conditions, is what we call a true prediction. The RMS difference between model and real handle positions in the second step is typically three to five per cent. Tom Bourbon has shown that this same accuracy of prediction is found even with a lapse of *one* year between the prediction and the real run. The property represented by k thus appears quite stable over time, although it differs markedly (2:1) between individuals.

We have not said where the random errors come from in our model, but clearly they have to be coming from inside the subject, because our knowledge of d, c, and h is relatively exact.

I wonder if it still seems to Peter that there is no difference between the statistical and the model-based approaches (at least ours)? I have a suspicion that the way we are using the term "model" isn't quite the same as the way Peter is using it.

Gary, here is the part of the "random" program that generates the data:

```
for i := 0 to maxdata do begin
b := 1.5 + 3.5 * random; {for Hercules and EGA}
  k := 5.0;
  d := -random(40);
  r0 := 100 + random(200);
  effort := k * (r0 - d)/ (1.0 + k * b);
  reward := (b * k * r0 + d) / (1.0 + k * b);
  v2[i] := round(effort); v1[i] := round(reward);
ref[i] := r0;
end;
```


I set maxdata to 4000, but there's no need to go that far. The error sensitivity is fixed at 5.0 (k). The "cost" is d; the "wages" are b. The resulting effort and reward figures for each person are stored in two arrays: v1 (effort) and v2 (reward). The reference signals (amount of reward desired) are stored in the array ref. The entries in the reward and effort arrays amount to a single determination for each person.

In the article, I pointed out that in order to measure the reference signal for each person, it would be necessary to do a control-system type of experiment with every individual. You would have to vary the disturbance to find out what level of reward leads to zero effort *in each individual* (the definition of a measured reference level of a controlled quantity). As presented, the data do not show this: we know the internal reference setting for each person only because we know the correct model for each person. For an experimenter who does not know about reference signals, there is nothing to indicate their settings. The only externally observable variables are effort and reward.

I doubt that MR analysis would reveal the reference levels for each person. The concept of a reference level, a preferred level of input, is model-dependent, and here the model is that of a control system, not an input-output system. Similarly for the idea of error sensitivity (k). You can't measure k for an individual from a *single* observation. The loop gain of the system can't be seen unless you vary the disturbance and observe how much the disturbed variable, the reward in this case, changes. The loop gain would be the ratio of the disturbance magnitude to the change in reward relative to the no-disturbance value, minus 1. We know the external part of the loop gain (the wage) but must deduce the internal part, the error sensitivity k. I don't think any of these concepts are part of the model assumed under MR analysis.

The above program would be easy to implement in BASIC or any other language, or even on a spreadsheet. Rick Marken has done control systems on spreadsheets. Most statistics packages, I believe, can import data from spreadsheets.

You also said: "... I wonder what your reaction would be to someone who wants to show you a perpetual-motion machine (perhaps even one than can do work)." After all my experiences with control theory, I wouldn't reject a working perpetual motion machine on principle. But I would like to be alone with it for half an hour, with a few hand-tools.

Tom Bourbon: To Peter Parzer: I have enjoyed watching the dialogue between Bill Powers and you. You have certainly raised some important points concerning the nature of modeling. The most significant reminder you made for me is that the selection of variables and metrics is always in the hands of the modeler and can be done in various ways that can enhance the apparent success of the modeling enterprise.

As for the reliance on correlations in presentations of the results of modeling by control theorists, that selection was driven in part by a desire to have at least the index of performance be familiar to psychologists and other behavioral scientists, the majority of whom never work with continuous variables, and who never use other indices, such as RMS error.

Bill Powers: (In reply to a post by Peter Parzer.] It seems to me that the simplest comparison between the simulated handle position and the observed handle position would be a plot of the differences between them for the 1800 data points in a tracking run. We want to do this so that we can compare different models and see which predicts the results the best. We could simply look at two plots of prediction error against time and say, "Ah, the first one stays closer to zero over most of the points." Or, more likely, we would look for some measure that would be more reproducible over observers, such as the RMS error calculated for all the data points. As you imply, there isn't any "objectively right" way to measure overall error. But there are ways that are useful, simple, and reproducible.

Whether absolute or relative errors are used depends on the application. If you're talking about arithmetical calculation errors, absolute error is all that makes sense—after all, the relative error is always zero, in comparison with the range of values that numbers can take on (infinity). On the other hand, if you're judging how well a person steers a car, relative error makes sense, because what matters is how much the car wanders in relation to the width of its lane. I agree that there is a choice, but usually there's a pretty good reason for the choice. There's no one measure of error that suits all occasions.

In a tracking experiment, we have a record of 1800 positions of the handle. The model reproduces these positions with some error. But why should we assume that the errors we see are due to a random variable in the subject? Why shouldn't we assume that the model still does not capture all the properties of the real system correctly and that the remaining errors are systematic? Indeed, we find that when we refine the tracking model—for example, by putting in that time-lag I mentioned—the prediction errors become significantly smaller. In one experiment, the RMS errors of prediction dropped from 3 per cent to 1.5 per cent (noise-to-signal ratio). That tells us that at least half the error we obtained before was not random. Why should we assume that all of the remaining error is random? Of course, at some point we will run into what looks like a basic noise level, but the errors are already so small that they're approaching those of a physical measurement. When you speak of an "adequate" model, you have to ask "adequate for what purpose"? I think that in terms of predicting simple behavioral phenomena, the control-system model is adequately precise for any purpose we can now imagine. Our biggest problems now are in modeling more complex behavior.

The difference between models and statistical analysis really comes down to a difference in basic assumptions. I assume that prediction errors occur because although the person's behavior is completely systematic, the model is not yet exactly correct. It might not have been apparent, come to think of it, that when we speak of predicting handle movements in the tracking task, we mean predicting all details of movement with quantitative accuracy, not just comparing mean slopes or other average measures. The tracking model generates a trace of simulated handle movements that can be laid right over the trace of the real handle movements. It's hard to realize that the two simple equations I presented can do this, but they really can.

The other assumption would be that the model must be correct (for some philosophical reason), so the prediction errors are the organism's fault. Psychologists decided long ago that the variability of behavior was caused not by an inadequacy of their lineal cause-effect model, but by some inherent randomness of behavior. I have always felt that they gave up about 150 years too soon. We will surely have to give up trying to improve our models some day, but I would rather see that day come when "random" errors of prediction are in the 1 per cent range rather than the 100 per cent to 1000 per cent range.

In the models we use, not only the variables have empirical meaning, but the individual relationships between them have empirical meaning, or at least a proposed empirical meaning. We propose, for example, that an error signal results from neurally subtracting a perceptual signal from a reference signal. The subtraction process is part of the physical model. In the tracking experiment, d, c, and h have empirical meaning, but so does the relationship c = h + d. If we gave the handle twice as much effect on the cursor, the relationship would be c = 2 * h + d. This part of the model embodies known physical relationships. The other equation proposes physical relationships inside the control system. The behavior of the system grows out of the interaction of these two aspects of the model.

We use "generative" models. That is, they do not directly represent behavior, but propose an underlying physical organization that creates behavior because of its inputs and the way it treats signals internally. Such models predict not only the specific input-output relations observed in a single experiment, but a whole family of relations that can be seen under many different experimental conditions. The model I described for the tracking experiment, for example, predicts just as accurately when we make the target position a function of time, without any change in the parameter k (still applying a disturbance directly to the cursor as before), and when we halve or double the effect of a given handle movement on the cursor. Most experimental psychologists who actually try these experiments find the generality and accuracy of the models to be little short of uncanny—especially in comparison with what they're used to.

This is why I can't get too excited over just how we measure prediction errors. We're talking about errors an order of magnitude smaller than those that are usually seen in behavioral experiments (outside psychophysics).

Martin Taylor: Gary, one could indeed say d′ is a measure of signal-to-noise ratio in some abstract sense. Given an ideal observer under specified constraints on information gathering, one can determine the SNR that gives a specific d′. (Actually, it is signal energy rather than power that usually determines the d′, but the details always depend on the observing constraints). One asserts that there exists some perturbation of the observation (noise) that can move a non-signal observation to a more signal-like state, or a signal observation to a more noise-alone state. If the signal is weak enough, the distributions induced by the perturbations can overlap. One asserts furthermore that there is some criterion on which the observer makes a judgment as to whether a signal was present, and that "signal" is more likely the greater the value of the observation on this criterion. If the criterion axis can be transformed (squashed) so that the perturbation-induced distributions take on a Normal form, and particularly if the Normal distribution has the same variance whether or not a signal was present, then d′ is the distance between the means of the distributions in units of their common standard deviation. In more complex situations, the definition is different, but related. With common Normal distributions, it is exactly your "effect size," and unity is often taken to be the dividing line between "perceptually nonexistent" and "perceptually valid," though the subject sees each individual signal presentation as there or not, regardless of d′. The problem for the subject is that the signal might be perceptually there when none was presented, or not there when one was presented.

Perception *is* a problem of statistics, and treating it (properly, in my view) as a control problem will not make that go away.

Bill Powers: Martin, I agree that statistics can enter into perception, but I doubt that a properly designed "test for the controlled variable" (which identifies controlled perceptions, as nearly as we can) will leave us worrying about effect sizes and standard deviations in the way you suggest. When you've identified a controlled variable using control theory, it's pretty unequivocal.

In control theory, we seldom do experiments with perceptions at their lower limits of detection. The normal case, which I think represents the overwhelming majority of real cases, involves perceptual variables that are far above their thresholds of detection or discrimination, and neural signal frequencies that are comfortably above the levels where individual impulses have any appreciable effects. After we have models that function well in this middle range, we might want to explore behavior and perception near the limits of operation where noise becomes a significant consideration. But I don't think we've reached that point yet.

Rick Marken: Here is another thought I had about statistics—just to see if it can stir up some comment. The previous statistics discussion has dealt mainly with the problem of using group-level statistics to form conclusions about individual processes. This was approached in several ways—in particular, showing that even relatively high group-level correlations imply substantial error in individual prediction (the coefficient of failure).

But group-level statistics do work on groups. Lowering my cholesterol intake might not help me personally (indeed, it might kill me), but that does not diminish the fact that, at the group level, there is evidence of lowered heart disease with lowered cholesterol intake. This is "true" at the population level. On PBS last night, they reported that a government program to reduce dietary fat in Finland has led to a 30% decrease in heart disease. Ignoring the problems of attributing all of that 30% to the dietary change, this is evidence of a group-level change having a group-level effect. The same thing happens with seat belts. Death rates, at the group level, do (I believe) decrease substantially with mandatory seat belt laws—even though this is not necessarily the case individually. In fact, many people who might have survived an accident (like a burning car) were probably killed because they were wearing their seat belt. But overall, the death rate does go down.

That's the basis of my question. What do you folks think of this problem? Apparently, we can have some control over group data by doing things individually which might not be in our best interests. Apparently, we can influence the group-level rate of heart disease by collectively (but as individuals) reducing fat intake. We can do this even though some of us, individually, might actually be worse off as the result of taking that action (though we can't know that, of course, because we only have the poorly predictive group-level data to go on). This seems like a crazy paradox; and it seems to occur a lot in society. "Should I ignore the potential group-level good and continue to do what I want based on the extremely good argument that it is meaningless to base my individual actions on group-level data? Or should I cooperate with the statisticians in order to produce a beneficial group-level result by taking action that could possibly have negative individual consequences?"

If the data say "80% of people who take X get cancer," and (1) I like X, but (2) I don't want to get cancer, isn't it a good bet for me to avoid X? (Assume that I like X *far less* than I dislike cancer).

Gary Cziko: The answer to Rick's last question depends on how much he likes X and how much he dislikes cancer. This is the stuff of classical decision theory. A nice introduction to this kind of thinking can be found in Ronald Giere's *Explaining Science.* (But Bill Powers would probably add to this that it also depends on how similar you think you are to the 80% of people who get cancer doing X.)

Here are two quotes from J. G. Taylor, "Experimental Design: A Cloak for Intellectual Sterility," *British Journal of Psychology* 49, 1958, 106-116.

> If Newton had had at his disposal not a vast amount of detailed information about a single solar system but a much smaller number of facts about each of a thousand solar systems, collected by a thousand observatories, he might conceivably have developed statistical methods for organizing this material. He might have found correlations between such variables as the number of planets in the system, the average number of satellites per planet, the average distance of the planets from the sun, and the like. He would, by this means, have learned a good deal about solar systems in general, but he could not have calculated the time and place of the next eclipse of the sun, and he could not have arrived at an understanding of the laws of planetary motion. He would have learned a lot about the ways in which solar systems differ from one another, but nothing about the ways in which any one of them works. For this it was necessary to know as much as possible about one system. Fortunately Newton had no alternative, and the result of his labours was the construction of a theory that survived until the advent of Einstein's theory of relativity. (p. 109)

> Suppose that an investigator, knowing nothing about the construction of a motor car, decided to choose as his area of research the behaviour of the speedometer needle, and to this end took a series of readings in each of a hundred different models. Just to make the problem more like a real one we shall suppose that the speedometer dials are not provided with scales, but that the investigator can measure the angular deviation of the needle. Among the variables he might be expected to record are the distances of the accelerator and brake pedals from the floor, the position of the gear lever, the gradient of the road, the direction and velocity of the wind, and, of course, the speedometer reading. He takes a succession of simultaneous readings of all those variables in each car, and then proceeds to examine his data in the hope of solving the riddle of the speedometer needle. At first the material looks completely chaotic. There is no single independent variable that is functionally related to the dependent variable, and he decides to have recourse to statistical analysis. He finds negative correlations between the speedometer reading and (a) the distance of the accelerator pedal from the floor, and (b) the gradient of the road; and positive correlations with (c) the position of the gear lever, and (d) the distance of the brake pedal from the floor. He finds significant differences between the speedometer readings when the gear lever is in first, second, third, and fourth positions, but the distributions overlap extensively. He now decides to record additional data, such as the weight of the car and its consumption of petrol, but the riddle remains unsolved. Of course we know the answer. If our investigator will only take independent measurements of the speed of the car he will find that in each system (car) the speedometer reading is a function of speed, but not necessarily the same function in all systems. He will find, moreover, that he can now dispense with statistical methods and can examine each system, considered as a matrix of pointer readings representing the several recorded variables, to determine how it hangs together. He will discover that what he at first took to be evidence of arbitrariness or caprice in his data was actually an artifact arising from the simultaneous examination of pointer readings taken from a hundred different systems. He will find that the same general principles apply to all the systems, but each of them has its own specific set of parameters, with the result that, in Ashby's (1952) terminology, the lines of behaviour of all the systems are different. Continuing to use Ashby's terms, each system is regular and absolute. It is regular because whenever it starts from a given state and a primary operation is applied to it, such as an increase in the gradient of the road or a specific depression of the accelerator pedal, the system will change to another state, and always to the same state. It is absolute because this is true no matter how the given initial state was arrived at." (pp. 110-111)

I'm not sure that even Bill Powers or Phil Runkel could say it better than this.

Rick Marken: Thanks to Gary Cziko for his response to my little statistical question. I'll tell you why I asked. I had a discussion with my wife and daughter about the value of using statistical information for individual decisions. I took my typically extreme position, claiming it was useless. I, of course, was creamed in this discussion, not only because both of my opponents are orders of magnitude smarter than I am, but also because they made it personal. They asked if I would feel any different if my daughter were walking around at night in a statistically dangerous as opposed to a statistically safe neighborhood. Well, I'd rather she weren't walking around alone at night, period—but the fact is, I would rather she avoid the dangerous neighborhoods. We do base personal decisions on statistical data (in a decision-theoretic sort of way, as Gary pointed out). I suppose that we do so mostly when we can imagine a plausible causal relationship between what we do and the possible results. That's also why we don't stop listening to Bing Crosby when we find out that Bing listeners don't live as long as others; there is no plausible causal link that we can imagine doing anything about.

What I was looking for was a nice, clear, simple, and compelling way to justify ignoring group statistics if they really are irrelevant to individuals, and to show why and when this is the case. I think this is relatively important, because this is how medicine, social science, and most of the other life sciences work right now—they present group data as something that should be used as guidance for individual behavior. If this is a bad idea (and I feel somewhat that it is), then we should have a clear, crisp explanation of why this is so. I have been unable to clearly articulate that explanation.

I don't think it's often a problem, but I think many people actually do have serious conflicts (and control theorists should be interested in them) resulting from the fact that they are given group data suggesting that they should change their wants. In this sense, group statistics, which suggest ways to get "group-level improvements," can create individual conflicts.

Bill Powers: Rick, regarding your statistical question, if the indications are that 80 per cent of people like you are put at risk by taking X, you will only take X if you like it at least five times as much as you dislike getting cancer. But do you think that the numbers for any of these highly publicized risks are anything like 80 per cent? Consider this statement: "Among all people with clinically high cholesterol, p per cent of them die from heart attacks." Can anybody supply an actual number for p? Then consider this statement: "Among those who undergo a program designed to reduce their blood cholesterol, q per cent die of heart attacks." Again, can anybody tell us what q is?

With knowledge of p and q, you could then get a realistic picture of how worthwhile it is to try to reduce your blood cholesterol. My hunch

is that p is going to be a small number, and q is going to be only slightly smaller. The data for risks like these are never presented honestly; they're hyped up to create the most alarming numbers possible. They say, "People with high blood cholesterol are five (or whatever) times as susceptible to heart attacks as people with normal cholesterol." They don't tell you what the actual odds are, or how effective cholesterol-reduction programs are, because those numbers would be much less scarey or promising. In his book *Heart Failure,* Tom Moore pointed out that with the stroke of a pen, the Surgeon General declared 25 per cent of the population of the U.S. to have a medical condition (high cholesterol) demanding the immediate care of a physician. Drumming up business, that's what it was.

Gary, those quotations from J. G. Taylor show that, whether he intended it or not, he was helping to lay the foundations for a change to the method of modeling and the abandonment of statistics as a way of understanding human organization. Three cheers for Taylor. I'll even forgive him for citing Ashby and for overlooking invisible disturbances.

Joel Judd: Rick says: "I don't think it's often a problem, but I think many people actually do have serious conflicts (and control theorists should be interested in them) resulting from the fact that they are given group data suggesting that they should change their wants. In this sense, group statistics, which suggest ways to get 'group-level improvements,' can create individual conflicts." This strikes me as relating to cultural anthropology. Hunters and gatherers (to make a sweeping generalization) didn't have the *New England Journal of Medicine* giving them statistical data on what was safe to consume, etc. It's not simply a question of making decisions alone—we make them with regard to culture/society. We do not function in isolation. We do, however, make our own decisions. Hence the conflicts which often arise between what *we* want and what we *should* (is that a good way to put it?) want according to cultural institutions such as medicine, government, etc. Perhaps this gets back to the insidiousness of behaviorism—the propensity for those institutions that wield so much influence in our world to use behavioristic modes of thought to make decisions about what is right/wrong, good /bad, healthy/unhealthy—whether or not they do it explicitly. And so we are faced with dilemmas in making our decisions.

Rick Marken: Thanks to those who helped with my question about taking group statistics into consideration when making individual decisions. The solution seems simple: just ask how good the group statistics actually are (i.e., do 80% of people like me show the result, and do only 10% who are not like me not show it?); then, based on those data, decide if the result of changing to be not like yourself is worth it to you. It seems that, in most cases, the group results are so weak that it really isn't worth it at the individual level.

Tom Bourbon: In the many discussions about statistics, one issue we have neglected is that of the rates of occurrence of various conditions in the general population. An analysis of this issue goes to the heart of some of the more ridiculous abuses of statistics, and of the people to whom they are applied. This is a problem that even Phil Runkel misses in his delightful and devastating book.

An elegant recent example of how far thoughts can stray when scientists ignore base rates might be pertinent to Rick Marken's defeat in the conversation with his daughter and wife about crime, criminals, and "statistically crime-infested" neighborhoods. And this case shows how even the most sophisticated experimental procedures and analyses cannot save those who ignore base rates.

The study is A. Raines, P. H. Venables, and M. Williams, "Relationships between N1, P300, and Contingent Negative Variation Recorded at Age 15 and Criminal Behavior at Age 24," *Psychophysiology 27*, 1990, 567-574. (With a title like that, you know something good is in store! "Sliced and diced," a la Runkel's analysis.) N1, P300, and contingent negative variation (CNV) are measures of brain activity—in this case, electrical activity recorded from the scalp.

The study is predicated on previously published data showing that 16.2 per cent of boys who are not criminals at age 15 become criminals by age 24. The authors report the results of their work in which they recorded brain responses (ERPs), elicited by brief stimuli, from the scalps of 15-year-olds. They administered a variety of "psychological instruments" to the boys. At age 24, they determined how many of the 101 boys were criminals. Then they looked back at the ERP data and the psychological assessments and determined which of the many possible features of the ERPs correlate significantly with anything—test scores, criminal record, one another, etc. The results convince the authors that certain "cognitive components" of the ERPs predict criminality.

For example, there is a "highly significant" correlation between amplitude of N1 at 15 and "psychopathy" at 24. (They report $r = 0.73$, which means p(failure) = 0.68.) Another "highly significant" ($r = 0.65$, p(failure) = 0.76) correlation occurs between amplitude of CNV at 15 and "psychopathy" at 24. Now those results really tell me a lot about criminality! For Rick, I guess it means you might want to set up an evoked potential system by the front door, for testing your daughter's dates!

The reason for that is that of the 101 boys, 17 became criminals by age 24. (That means 84 did not.) And a discriminant function analysis using N1 amplitude and P300 latency (why *that* particular combination?!) at 15 as "predictors" of criminality status at 24 correctly identified 75 per cent of the budding crooks! That means ERPs correctly predicted 13 of the 17 who became criminals. Impressive, isn't it? It isn't! The same "predictors" incorrectly tapped 26 per cent of the innocent boys as future felons. That means 21 boys.

The authors attend to the *percentages,* within a limited sample; by doing that, they see that the ERPs correctly identify nearly three times as many criminals as they misidentify (75 per cent vs. 26 per cent). But if you look at the *numbers* of boys, nearly twice as many innocent boys are pegged as future criminals as are guilty ones.

Oblivious to that fact, the authors go on to talk about the use of ERP data as possibly playing a role in identifying potential criminals. What if they were to succeed in that goal? Imagine a major program designed to spot the little buggers and nip them in the bud. If they tested 1,000,000 15-year-old boys, and if everything worked as they report in their research, 162,000 boys would be criminals by age 24, and the ERPs would have spotted 121,500 of them. Now *that* is war on crime! But they would have misidentified 217,880 innocent boys.

Imagine what kind of world this would be if people really *believed* the stuff that comes out of behavioral research! Wouldn't it be nice if each editor of a journal in the behavioral sciences required that authors report the results of an analysis of base rates—the actual numbers of people *in* the population—who would be correctly and incorrectly identified by the procedures described by the authors? That policy, along with a requirement that no correlations be published below $r = 0.87$ (the 50-50 point for being right in a prediction), would reduce the literature to about one slim volume a year. A person could read it in an evening and could have faith that at least part of the material was worth even one evening.

Bill Powers: Base rates! I knew there must be a term for it. Thanks. Tom, why don't you work up all this material for a letter to *Science?* No doubt we would be dismissed by professional statistical types as amateurish, but if you could get a letter published, at least a discussion might be started, and we would be trying to do something about these atrocities. Maybe we could at least get p(failure) accepted as a necessary part of any report on statistical data.

Statistics is an excellent tool for evaluating data and even for seeing whether there is something to a new hypothesis. You can't (apparently) get along without it in quantum mechanics. We use statistical measures even in tracking experiments. And Rick Marken has used a statistical method for identifying controlled variables in situations where the reference level for the controlled variable is continually being changed by the subject. I envision many applications for statistical analysis in the control-theory approach to behavior.

What I insist on, however, is the proper use of statistics. A statistical measure should be used only for the population from which it came. Mass measures should *never* be used to evaluate individuals if the odds of a misevaluation are significant *in terms of the payoff for the individual.* There are legitimate uses for mass measures, but the most common uses do not properly take into account the potential (and very often actual) unfairness to individuals that results from mechanical applications of statistical facts. Too often, statistics is used as an easy way to get a publishable result, with (as Tom indicated in his post) a consequence of flooding the literature with meaningless garbage (not that I'm in favor of publishing meaningful garbage, either).

Statistics is really not a tool for prediction, because all predictions imply that we want to know the value of a variable at a particular time and under particular circumstances, whereas the statistical analysis is derived from many variables evaluated at many times under variable circumstances. If we understood the underlying principles that make one variable dependent on others, we would not have to use statistics except to judge the uncertainties of measurement. More importantly, the principles that relate variables in actual behavior can hardly ever be boiled down to a simple cause-effect relationship, nor should they be. Even when we know that a person reacts with fear to dogs 80 per cent of the time, we do not know why the person reacts to any one dog with fear. Reducing that person's fear-reactions to 10 per cent might do the person a terrible disservice, if there are pit bulls and attack-trained Dobermans in the environment. Knowing the particulars is always better than knowing generalities.

And never forget that *real* statistical results seldom give us probabilities anywhere near 80 per cent.

Cross-correlation is a valid statistical method, in fact the first method I used some 15 years ago to try to detect a transport lag. I based my initial opinion about the lack of *a* transport lag on the fact that a cross-correlation measure had a peak at zero delay. But it was also true that the cross-correlation function did not show a clear peak; it was very broad, too broad to discriminate well. I think I now understand the reason. The cross-correlation method deals only with the intact closed loop of control processes, so the variables (cursor position and handle position) are not really independent. Cursor movements are dependent on handle movements, as well as on the independent disturbance. I did not find any effect of a transport lag until I put it into a working model in the forward part of the loop (the person) and by trial and error found the value that minimized the RMS error between the model's handle behavior and that of the real person. The minimum in the prediction error is still very broad, but it occurs quite reliably at the same value, trial after trial, and that value is not zero.

Control theorists are often criticized for using single-subject data. But if I had tested this model for transport lag in the usual way, proposing a one-size-fits-all model and fitting it to pooled data from many subjects, I doubt that there would have been a significant result. The model parameters differ from person to person (although the best transport lag differs less than the other main parameter, integration factor). The use of a model applied to individual data is essential here; without it, the statistical results would mean very little.

So 1 believe in the use of statistics, but only when it is properly applied and subordinated to a model. Predictions should be made from a model tailored to the particular system being observed, not from statistical measures alone (which rest on too simple a model). There is no way to avoid studying individuals if you want to understand individual behavior. I believe that current attempts to understand mass behavior are mostly ineffective. I believe that once we have a decent model for individual behavior, we will be able to synthesize predictions of mass behavior that work far better. If we see any point in doing so.

Also, Tom, from your numbers, I take it that a total of 13 + 21 boys, or 34, were predicted to become criminals. Of the 17 who became criminals, four were predicted innocent, while among those who were innocent, 21 were predicted guilty. This means that 73 per cent of the predictions of criminality were wrong, doesn't it? The "coefficient of failure" is 0.68, so it's an underestimate in this case.

You mentioned two criteria: N1 and CNV both correlated with criminality. How many subjects showed *both* N1 and CNV, and what was the criminality rate for those showing both? This is pertinent to the discussion that Gary raised (which got us into all this) about using multiple criteria for evaluating risk. My contention was that multiple criteria would do even worse than any single one.

Tom Bourbon: Bill, I am working on a letter, or a short report, on this topic. If I include a few of the many other examples from different types of journals and on a selected range of topics (to show that no major area of the behavioral-social-life sciences is clean), it might be a bit long for *Science.* Another possibility is *American Psychologist.*

And yes, the multiple criteria did have a higher likelihood of being wrong! Another thing about that multiple-variable, discriminant function analysis is that the variables entered into it are not the same ones used to report on significant single-variable correlations with "psychopathy." For the simple correlations, the authors used "amplitude of N1" vs. "psychopathy" and "amplitude of contingent negative variation" vs. "psychopathy." (By the way, the "instruments" used to "assess" "psychopathy" are yet another grisly issue!) For the discriminant function analysis, the amplitude of N1 is still in, but CNV is replaced by the latency of P300. Now, why was that done? Of course, I do not have the details, and I do not wish to impute dishonorable motives to the authors. However, brain response data offer a wealth of conceivable "measures" to enter into analyses: the amplitudes and latencies of every distinguishable "event" in the data record, the ratios of any conceivable combination of measures of "events," and so on. The list is immense. So why do any two, or more, of those measures happen to "predict" in one study, but some other combination or combinations work in another? The answer is that none of the combinations predict, except in the trivial sense of meeting a criterion of statistical significance. And the many discussions, post hoc, of why that particular combination worked in an earlier study, but this combination worked this time, lead nowhere.

Joel Judd: Tom, do I detect a note of *cynicism?* Just to keep you a little wider awake at night, the "study" you mentioned reminds me of a CIA contract the psychophysiological lab here on campus was trying to get a couple of years back when I was attending lab meetings. The "shop" was dangling fat grants to labs which could produce a sure-fire ERP lie-detector test. Fortunately, I don't believe anything ever came of it, at least not here.

Tom Bourbon: Joel, you seem to share my concerns over the misapplication of "objective" physiological measures which correlate, however pitifully but significantly, with important behavioral and psychological processes. In the late '60s, I was asked by a company in the region to look at a proposal submitted to them by a neuroscientist-psychologist. He wanted the company to put up venture capital for the manufacture and distribution of his device for measuring the latency of one "component" of human auditory evoked potentials (EPs).

He claimed, in his proposal, in several publications, and in the reports submitted to federal funding agencies, that the latency of that one component correlated significantly with various full-scale and subscale measures of "intelligence" (with n = 566 children, he had r's from -0.04 to -0.35 between latency and various IQ scales and subscales; and, as he reported, with n = 566, Pearson r's of 0.16 are significant at p less than .0001).

The scientist went on to say that his "findings" (why does that word always remind me of "leavings"?) could have "considerable educational significance," principally via the use of the EPs for "objective, culturally independent biological assessment of mental potential useful in exploring possible racial differences in intelligence." And he went on to suggest that EPs recorded from fetuses might weigh heavily in decisions about whether a pregnancy should go to term or be aborted. All of that from correlations the best of which would lead to incorrect predictions at least 94 per cent of the time.

My report to the company was not received kindly. And the "real scientist" (who was I to question him?) took umbrage. By that time, his research was featured in various educational journals and magazines, and in offerings to school districts, which could purchase the system or the services of professionals who would administer the assessments.

This abomination vanished soon after. I like to think that my report helped it on its way. The episode marked my awakening from graduate training in which I had to virtually swear a solemn oath that the answers to psychological questions were to be found in physiological research.

The assumptions one makes about the causes of behavior and the data one accepts as supporting those assumptions are not matters of idle sport and speculation. When they work their way into decisions about policies that affect the lives of innocent people, the scientists who offer them ought to be held strictly accountable and responsible. All the more reason for us to insist on models that work at least in simple instances of behavior and on data that predict what actually happens, at least half of the time!

Gary Cziko: Tom has been providing some fascinating accounts of the misuse of statistics in predicting individuals. But I am having some difficulty understanding the way he is conveying information about correlation coefficients.

For example, he says: "... there is a 'highly significant' correlation between amplitude of N1 at 15 and 'psychopathy' at 24. (They report r = 0.73, which means p(failure) = 0.68.) Another 'highly significant' (r = 0.65, p(failure) = 0.76) correlation occurs between amplitude of CNV at 15 and 'psychopathy' at 24." And he also says: "All of that from correlations the best of which would lead to incorrect predictions at least 94 per cent of the time."

It seems in the first quote that Tom is saying a correlation of r = 0.73 gives a p(failure) of failure of 0.68. I don't think this is quite the way to put it, since, to me at least, p normally indicates a probability, which this isn't.

If we take 0.73, square it, subtract the squared value from one, and then take the square root of the difference, we will indeed have a value of 0.68, which I have seen referred to in at least one statistics text as k, the coefficient of alienation. That is, $k = \mathrm{sqrt}(1 - r^2)$. But k is no probability, it is rather the ratio of the standard error of estimate of using one variable to predict the other to the standard deviation of the criterion variable. So if 0.73 is the correlation between years of education and income, using education to predict income will give us 68 per cent (about two-thirds) of the error (difference between predicted and actual income) that we would get if we knew nothing about anyone's education and just used *the* mean income of the group to predict each individual's income. Or, subtracting 0.68 from one, we find that the correlation of 0.73 gives a 32 per cent improvement in predicting Y based on X over not knowing anything at all about X.

So it seems to me that the p(failure) notation is misleading if Tom is using p for probability. In fact, the probability of predicting someone's score exactly right on a continuous variable measured with infinite precision is actually zero (which is why statisticians don't like point estimates and use interval estimates instead).

Also note that correlations start to look better when you are trying to simply predict whether someone will be higher or lower than some predetermined criterion. If I simply want to know whether someone has an above average or below average IQ based on some predictor (e.g., some brain-wave measure), then the probability of correct predictions rises dramatically (I can give some tables if this is of interest). But then the question arises as to what average IQ is, how it is determined, and how just being above or below average correlates with some other variable of real interest (such as whether someone finishes high school or not). So I doubt that the predictive value is really much better even in this dichotomous case. (It might be better if the criterion variable were something clear-cut like sex, but there are probably easier ways to predict sex than by using brainwaves.)

Maybe the best way to talk about this new index we like so much is to subtract it from one, multiply the difference by 100, i.e., 100 * (1 - k), and call it something like "per cent improvement" (PCI). So in the above case of *r* = 0.73, PCI = 32 per cent, meaning that errors of prediction using the predictor variable are on average 32 per cent better (i.e., less) than just using the mean of the group to predict each individual's score in the group.

This is what Tom's interesting statement would look like using PCI: There is a "highly significant" correlation between amplitude of N1 at 15 and "psychopathy" at 24. (They report r = 0.73, which means PCI = 32 per cent.) Another "highly significant" (r = 0.65, PCI = 24 per cent) correlation occurs between amplitude of CNV at 15 and "psychopathy" at 24.

Hmm. After looking at this, I think I prefer the "uselessness" approach after all. Just like above, but don't subtract from one. That gives the "per cent uselessness" (PU; it even sounds right). Now the statement looks like this: There is a "highly significant" correlation between amplitude of N1 at 15 and "psychopathy' at 24. (They report r = 0.73,

PU = 68 per cent.) Another "highly significant" (r = 0.65, PU = 76 per cent) correlation occurs between amplitude of CNV at 15 and "psychopath)," at 24.

Yes, I like PU much better, since most of the correlations we find in social sciences research really do stink. Suggestions welcome. Vote for PCI or PU.

Tom Bourbon: Gary properly chastised me for saying that k might represent the probability of failure in predicting Y from X, given a correlation r. My initial interpretations of Bill Powers' remarks on k were to blame—the fault is mine, not Bill's.

My utter lack of familiarity with this index puzzled me: the coefficient comes directly from the calculations for Pearson's r, so why is it not discussed in statistics books with which I am familiar?

I did find one fleeting paragraph in a text from my student days, but it is in a section marked "not assigned." I just located a rather thorough discussion in a text from 1956 (before my university days): J.P. Guilford, *Fundamental Statistics in Education and Psychology,* McGraw-Hill, New York. On pages 375-379, he discusses "the correlation coefficient and accuracy of prediction." Guilford characterizes the relationship between r and k as follows: "Whereas r indicates the strength of relationship, ... k indicates the degree of *lack* of relationship.... If r is 0.50, k is not also 0.50 but 0.886. Where r is 050, then, the degree of relationship is less than the degree of lack of relationship. It is when r = 0.7071 that the relationship and *lack* of relationship are equal."

And, as Gary suggests, multiply k by 100 and: "Our margin of error in predicting *Y with* knowledge of X scores is (k * 100) per cent as great as the margin of error we should make *without* knowledge of X scores."

Guilford goes on to describe 100 * (1-k) as the "percentage reduction in error of prediction," also known (then) as the "index of forecasting efficiency, E." I wonder why all of this dropped out of the statistics texts?

I vote for PU, of course!

Chuck Tucker: I think that the comments on statistics on the net are clear, concise, well documented, and will disturb the social and behavioral scientists (sic) to no end. These comments question the "articles of faith" that support the social sciences. They should be published in some form, if nothing more than being sent in outline form to every electronic network in the country with members who are social scientists. I only have a few comments by way of refinement.

(1) We should not make the error that everyone else makes when using the word "group." A group is a set of people who at least interact

with each other. My criticism of sociologists is that they define their discipline as the "study of groups," but they only study individual characteristics—not the individual as a person, or even a personality. So the statistics we are talking about are numbers generated (how?) from individual characteristics and put in categories or other aggregate forms through various means of classification—we don't have group statistics. The closest we come to group statistics (which sociologists have completely ignored in their work) is to be found on the sports page of the newspaper and, to some extent, the business section. Most of the statistics that we are told about and find in our journals are *not from* groups.

(2) We should note very clearly that these statistical presentations have serious effects: many people, especially government officials, "control for" such numbers. There is very good evidence (yes, numbers) that most journal editors (Clark McPhail has a series of papers on this issue) and readers will not consider a paper suitable for publication without statistics. We have developed a nation of quantofanatics!

(3) I wonder if those who are critical of the use of aggregate statistical analyses being applied to individuals and also believe that extreme competition leads to many of the problems we have among people have abandoned the use of "curving" or distributions for deciding what grades students receive. I believe that one of the most serious problems of our public education system is the use of "curves" to determine a student's grade, rather than the use of a standard set by the instructor/teacher and understood by the student. When the "standard" is merely doing better in a statistical distribution than some others, students only have a minimal notion of what is "excellent work." When we have raised a generation of parents and teachers who have experienced such procedures and continue to pass them on, then we should expect a continual lowering of the statistical standards (by the way, this would be an excellent experiment to be done by those in education—does it lower standards?). The point: to be consistent with control theory, a teacher should set a standard, encourage students to use that standard, and judge students' performance by the standard set, without regard to any statistical distribution of an aggregate (a college class is not a group, either!). When this is done, all can get high, medium, or low grades. Students can study together; there is less conflict among them and between students and teachers (although I do get complaints when they don't get high grades—but I am the only teacher at my institution who approaches grading in this way, and a less-than-high grade is a disturbance).

Martin Taylor: I am just starting to read Bill Powers' 1973 book for the first time, and in talking about time-scales of response (page 54), I

come across the following quote: "Psychologists who believe that intermittent reinforcement is more effective than continuous reinforcement should give this whole speed-of-reaction problem serious thought—for a long enough time." I realize that this was written a long time ago, and might have been amended later in the book, but it does resonate with some of the threads that have been weaving through the net—statistics, in particular. So although it might be unfair, I will comment.

Intermittent reinforcement is not usually seen as "more effective," but as more resistant to extinction. And a statistical reason is not hard to find. In the laboratory, the animal is confronted with a situation in which it is sometimes rewarded for behavior A, but never for behavior B (or less often, perhaps). Now, if the experimenter decides to stop rewarding behavior A, how can the animal know that the world has changed its rules? Previously, failure of reward for A has been followed by further reward on a later occasion. It cannot know that this will no longer be true. Only by implicitly evaluating the statistics of the reinforcing event can it determine after a while that a long period of non-reinforcement would have been unlikely under the regime to which it had become accustomed. If you like, there is a "continuous" higher-order event—a statistical event—which occurs on a time-scale much longer than that of the single reinforcement.

In such an experiment, the experimenter tries to make sure that the animal has no access to information that might let it know which rule is in effect. Many experiments have been found to give results that depend on the animal hearing a click or something that the experimenter had not noticed, but that occurred only when reinforcement was going to be provided. The animal then has a context that turns the statistical event into a predictable event. It can know that the world has changed if it no longer hears the click.

It should be much easier to learn a behavior that has a perfectly predictable consequence, but normally we do not have access to all factors that influence the consequences of our behavior, and so we have to resort to statistics to determine how our behavior is influencing our perception. The control system can be fully determined in its behavior, but if we cannot tell the difference between a context in which behavior A leads to result P and one in which it leads to result Q, then all we can do is to take advantage of the best information we have; that is, for example, that A then P has happened 75 per cent of the time we did A, and A then Q has happened 25 per cent of the time. If we want P to happen, and it is not too bad if Q happens instead, then we would do A. But if Q would on this occasion be disastrous, we might try another way of getting P to happen rather than risking behavior A.

Life, even in a control-system view, is a statistical game.

Sorry if that's all too obvious to have been mentioned, but I have read so much trashing of statistics on the net that it seemed rather to be so obvious as to have been overlooked.

Bill Powers: Martin, you say: "Intermittent reinforcement is not usually seen as 'more effective,' but as more resistant to extinction. And a statistical reason is not hard to find." I agree in both regards. I was thinking in terms of "habit strength" and Skinner's "shaping" experiments when I said "more effective." Both are related to extinction. (Skinner found that by changing the schedule so as to deliver fewer reinforcements for the same behavior, he could *increase* the rate of responding. He cited this as an instance of the power of intermittent reinforcement, never realizing that this relationship is the opposite of the one he always assumed to hold between reinforcement rate and behavior rate.)

As to the statistical reason, there are many cases in which a statistical analysis comes out with the same results as a modeling analysis without statistics. Suppose that an animal has learned to perceive the rate at which some almost-rhythmic stimulus appears. Representation of this rate as a neural signal (by analogue means) would require a smoothed frequency detector. The smoothing is required to eliminate the individual instances of an input and produce a signal whose magnitude is proportional to the rate of appearance. The amount of smoothing used determines the range of input frequencies over which the signal magnitude usefully indicates input frequency (too long a smoothing time yields a maximum signal for all rates above a certain limit). Within the range of operation, the signal magnitude corresponds roughly to the probability that an input will occur within a given time interval, related to the smoothing time. So the analogue perceptual function can accomplish the same end as a probability calculation, but in a quite simple way. If we were choosing on the basis of simplicity of circuitry, I would pick the analogue method. Of course, we must ultimately pick the method that the nervous system actually uses.

Given the smoothing time, it will take a certain number of input events to bring the perceptual signal to a constant level, and this will determine about how fast the related control system can act. When the input events stop occurring, the perceptual signal will take the same length of time to decay, so the system will go on attempting to control the signal after the input events have actually stopped (the extinction curve). This is in fact how it works: if learning takes a long time, so does extinction, at least in certain learning experiments.

I believe that this analogue model gives about the same results as a statistical-perception model does. The analogue model works with inputs that have an average frequency with random variations. It does *not* work properly (and neither does the statistical model) when the

input frequency is perfectly regular. We notice the first tick of the clock that is missing or comes too soon or too late. So that sort of situation requires not an average rate detector, but a synchronized detector (I think I would put it at my "event" level of perception, whereas the other kind of rate detection would go one level lower, at the "transition" level).

Generally, I think that your analysis of intermittent reinforcement is correct. I'm only proposing an analogue method that does, in effect, the same computations but without requiring statistical calculations.

I'm not against statistics in general, or even against statistical explanations of neural functioning (at the appropriate level). When we consider noise in control systems, statistical methods help us appreciate its effects. What I "bash" with enthusiasm is the misapplication of statistical facts to individual occurrences. I've tried to make my criticisms specific to that case. That would seem to be a subject different from the one you are talking about.

I don't think we often get into situations where the environment is ambiguous or unpredictable. When you look around, you see a pretty noise-free visual field, with clear demarcations between objects, colors, sensations, relationships, and so on. When uncertainties do arise, we might sometimes use statistical methods to deal with them, by which I mean literally computing chances, but I think in many cases we simply smooth out our perceptions and operate on the basis of the artificially unambiguous result—often wrongly. Anyway, people don't seem to compute their behavior on very good statistical grounds, do they?

Just for fun, a poem by Maurice G. Kendall, originally published in *American Statistician* 13(5),1959, 23-24:

Hiawatha Designs an Experiment

1. Hiawatha, mighty hunter
He could shoot ten arrows upwards
Shoot them with such strength and swiftness
That the last had left the bowstring
Ere the first to earth descended.
This was commonly regarded
As a feat of skill and cunning.

2. One or two sarcastic spirits
Pointed out to him, however,
That it might be much more useful
If he sometimes hit the target.
Why not shoot a little straighter
And employ a smaller sample?

3. Hiawatha, who at college
Majored in applied statistics
Consequently felt entitled
To instruct his fellow men on
Any subject whatsoever,
Waxed exceedingly indignant
Talked about the law of error,
Talked about truncated normals,
Talked of loss of information,
Talked about his lack of bias
Pointed out that in the long run
Independent observations
Even though they missed the target
Had an average point of impact
Very near the spot he aimed at
(With the possible exception
Of a set of measure zero).

4. This, they said, was rather doubtful.
Anyway, it didn't matter
What resulted in the long run;
Either he must hit the target
Much more often than at present
Or himself would have to pay for
All the arrows that he wasted.

5. Hiawatha, in a temper
Quoted parts of R. A. Fisher
Quoted Yates and quoted Finney
Quoted yards of Oscar Kempthorne
Quoted reams of Cox and Cochran
Quoted Anderson and Bancroft
Practically in extenso
Trying to impress upon them
That what actually mattered
Was to estimate the error.

6. One or two of them admitted
Such a thing might have its uses
Still, they said, he might do better
If he shot a little straighter.

7. Hiawatha, to convince them
Organized a shooting contest

Laid out in the proper manner
Of designs experimental
Recommended in the textbooks
(Mainly used for tasting tea, but
Sometimes used in other cases)
Randomized his shooting order
In factorial arrangements
Used in the theory of Galois
Fields of ideal polynomials
Got a nicely balanced layout
And successfully confounded
Second-order interactions.

8. All the other tribal marksmen
Ignorant, benighted creatures,
Of experimental set-ups
Spent their time of preparation
Putting in a lot of practice
Merely shooting at a target.

9. Thus it happened in the contest
That their scores were most impressive
With one solitary exception
This (I hate to have to say it)
Was the score of Hiawatha,
Who, as usual, shot his arrows
Shot them with great strength and swiftness
Managing to be unbiased
Not, however, with his salvo
Managing to hit the target.

10. There, they said to Hiawatha,
This is what we all expected.

11. Hiawatha, nothing daunted,
Called for pen and called for paper
Did analyses of variance
Finally produced the figures
Showing beyond peradventure
Everybody else was biased
And the variance components
Did not differ from each other
Or from Hiawatha's
(This last point, one should acknowledge

Might have been much more convincing
If he hadn't been compelled to
Estimate his own component
From experimental plots in
Which the values all were missing.
Still, they didn't understand it
So they couldn't raise objections
This is what so often happens
With analyses of variance).

12. All the same, his fellow tribesmen
Ignorant, benighted heathens,
Took away his bow and arrows,
Said that though my Hiawatha
Was a brilliant statistician
He was useless as a bowman,
As for variance components
Several of the more outspoken
Made primeval observations
Hurtful of the finer feelings
Even of a statistician.

13. In a corner of the forest
Dwells alone my Hiawatha
Permanently cogitating
On the normal law of error
Wondering in idle moments
Whether an increased precision
Might perhaps be rather better
Even at the risk of bias
If thereby one, now and then, could
Register upon the target.

Tom Bourbon: Several of my colleagues are somewhat tolerant of me and of students who turn on to PCT, but others are not so open or supportive. The person who asked seniors in a statistics course to present a talk on some controversial topic concerning uses of statistics in psychology was not prepared to have one student give a reasoned discussion of the "coefficient of failure," as discussed on the net. Nor was he ready for another student who, by all accounts, gave an elegant review of Phil Runkel's critique of abuses of the method of relative frequencies.

My students are told by some people that they don't care what kind of evidence he (I) might present, PCT isn't right, and it isn't psychology (I believe that!). During a discussion with several students who invoked

PCT as part of a challenge to his pet theories, a faculty member blurted out, "What does he [II do to you people, brainwash you?"

Bruce Nevin: I think there is a confusion of statistical prediction with prediction for an individual control system.

One can predict that most middle-class children will get an education; one cannot predict that a particular one will, unless that one is controlling for getting an education (for whatever reason). Likewise for learning their native language (exceptions may be autistic, severely retarded, kept locked in a closet, etc.). One could predict that Bill would marry an intelligent person because that was what he was controlling for (among other things), but not because most engineers in the field of astronomy who are former psychology students marry intelligent people.

Joel Judd: Bruce, isn't this why, in a certain sense, prediction becomes trivial in PCT? The trick is to find out what someone is controlling for. Also of interest is what the person does to reduce error. This might also be why, historically, so many psychological and educational researchers haven't told us much about process and mechanism, so concerned are they with predicting the right damn outcome. Why the outcome occurs and *how* it occurs must be explained by that black box up there.

Bill Powers: Perceptual control theory is fundamentally a theory of individual organization. You get to statistical predictions for populations in a different way. First you study enough individuals to find how their control parameters are distributed. Knowing that, you can predict how a population of "similar" (oops) individuals will do the same sort of control task. You will also know better than to speak of the "average way of controlling in this task." Nobody controls that way.

If you have ways of measuring individuals' control parameters, wouldn't it usually be unnecessary to go through the population-study route? When you study populations, you get characteristics of the population, but you don't learn anything about an individual, except perhaps the outer limits of variation within which this person might be found—unless the person happens to be from a different population and your criteria for population membership just didn't happen to pick that up.

One point of using control theory is to get away from statistical studies in which experimenters are jubilant (typically) over correlations as low as 0.8. Facts that are determined statistically are true only of a population and are next to useless for predicting the performance of an individual. There is a tendency to elevate findings that are true only of a majority of a population (say, 60 per cent of subjects) so that they are assumed true of the whole population.

There are two ways to understand natural phenomena. One is like trying to figure out a system for winning at roulette. You observe and observe, and finally you get an idea: every time two blacks and a red show up in that sequence, an odd number between 11 and 27 will win, but if the sequence is black, odd, black, red, the best bet is a number ending in 5. This is "looking for rules." It is also the basis for statistics, because when you're testing a rule like that, you have to keep track of how often it worked. If it doesn't work often enough to be useful (i.e., to keep you from going broke), you go back to searching for more rules.

The problem, of course, is that even if a rule appears to work, you have to consider how many chances you had to find it, how many times it might have failed before you noticed it, and how often it will fail in the future. Even if the rule appears to work in all your tests, it might still have nothing to do with anything. Even if the rule works 20 times in a row, there is always the chance that it is irrelevant or will become irrelevant without advance warning.

In fact, all you need is one exception to show that the rule is irrelevant. If you can have one exception, then you can have two in a row, 10 in a row, 100 in a row, and go broke.

Of course, there's always the chance that the rule you found actually has some explanation; it might be a reflection of a real regularity in nature, so that the rule really has to work (even though you don't happen to know why) or sometimes has to fail (depending on occasional underlying circumstances you haven't discovered). This, of course, is what we hope for when we try to guess at the rules. This is the mode of research that I call "trying to get lucky." Getting lucky means stumbling across one rule among all the others that is an expression of an underlying mechanism.

If you get into the gambling hall after hours, you can look under the roulette table. When you see a little button where the croupier stands, you can immediately deduce a rule for betting that has some reason for working: bet (small) against the biggest betters. The game is rigged.

So this leads to the other way of understanding nature: look for the way in which the game is rigged. Don't waste too much time trying to guess at the rules just by watching phenomena. The only rules that actually work are those that work for an underlying reason. All the rest are illusions. If you just look for rules, you can't tell the illusory rules from the real ones. And the real rules don't work just because they work: they work because they have to work. The game is rigged that way. The system is organized that way.

Modeling is an attempt to see under the roulette table.

Rick Marken: The only time I have encountered anything approaching hostility to control theory is when the listener figures out that control theory is completely inconsistent with the whole experimental/statistical framework on which psychology is based. Most psychologists really believe in this model. They spend years learning statistics and experimental design. It is the core of the discipline: the basic foundation on which the search for psychological truth has been built. Control theory says: forget it. When you say that to the people who wrote the texts, taught the courses, labored in the statistics classes, and paid their dues running hundreds of subjects in complex factorial experiments, you don't get big cheers. Even if you carefully show why conventional statistics/experimental design seems to work but really reveals little if anything about the internal organization of living systems.

So my experience is that control theory has the biggest problems when it comes face to face with faith in the *scientific method* as articulated in the pages of the exalted textbooks of statistics and methodology that are the bedrock of *all* (cognitive, behavioral, ecological, etc.) psychological science.

Bill Powers: [Replying to a researcher in cognitive science.] Experimentation under the control-system model is aimed at the characterization of individual behavior. The only reason for using multiple subjects in a single experiment, other than checking for flukes, is to see how variable the individual measures are over a population. We would never average such measures together! Question: What is the average damping coefficient of the human arm control system? Answer: That's not a meaningful question, because the damping coefficient must be appropriate to the build and organization of each control system, if it's stable. Details of organization vary greatly from one person to another.

I don't think that statistical studies can hack it in the long run. They have their uses, but once you've seen how control-theoretic experiments go, you'll be spoiled for statistical work. I say that with fingers crossed, because actually nobody is doing systematic research on PCT at the cognitive levels where you work—this is by way of inviting you to learn the basic principles of PCT and be a pioneer. Doing so will earn you the distrust of your colleagues, difficulties in publishing, and experiments with clear-cut results that you know are right. And friends like us who give you a hard time. You have to weigh the costs and benefits yourself.

The question we always ask people who report statistical results is "How many subjects *didn't* show the effect, and how does your hypothesis explain *their* behavior?" I claim that if you have to use multivariate analysis to show that there is an effect, you haven't got an effect. Real effects stand out like sore thumbs. They aren't the results of causes, but of organization.

My biggest objection to most statistical analyses (I don't know about your analyses) is that almost uniformly they employ a cause-effect model of behavior. We can *prove* that's the wrong model. Organisms produce consistent outcomes by variable means. It's easy to demonstrate this principle in almost any context, at any level. Most experimenters carefully avoid disturbances that might interfere with output, not realizing that the same outcome would happen anyway. Of course, if they did introduce disturbances, and the outcome did repeat, this would completely screw up their experimental paradigms. Maybe that's why they don't do it.

[In reply to a post from Eileen Prince asking what PCT has to say about autism.] Control theory isn't like most other theories: it doesn't say that if X happens to people, Y will be the resulting effect on their behavior. It's about the way behavior works; it describes relationships of a very general nature between perception and action. At the same time, it is a theory of individual behavior: in order to apply it to an individual, one must determine what variables that individual is controlling, and with respect to what internally specified states, and the quality of that control. The hierarchical model suggests a nested stack of types of controlled variables that people seem to be able to control when all is well—but the particular examples of these types that an individual controls can be discovered only by studying that individual.

Control theory doesn't use categories such as "autism" to explain behavior. To say that a person is autistic is only to say that certain externally visible patterns of action have struck people as similar enough (and unusual enough) to be lumped into a "disease entity." This does not mean that the same defect exists in all autistic people, or that the symptoms arose from some common history, or that the same treatment will succeed with (and not harm) everyone included in this category. The conventional empirical approach to treating problems as "diseases" is simply to try something on people in a given category and see if it helps a statistically significant number of them. There is no attempt to analyze what has actually gone wrong—what the person can still do normally, and what the person can't do. There is no attempt to relate deficits to a model of internal functioning. I suppose the idea is that if you accumulate enough experience with treating people in arbitrary categories, you will eventually be able to look up the symptoms in a big book and read off the treatment that has been effective most often in the past. In my view, this approach is an ill-advised attempt to bypass understanding of the human system and find solutions by relying on guesswork and luck. Before the advent of science, it was all we had. Sometimes it works. But there has to be a better way.

Rick Marken: Here's my hypothesis about what variable conventional psychologists (of virtually all stripes) are trying to control: the perception that they are able to have relatively (statistically) predictable effects on what other organisms do. Not surprisingly, the behavior of other organisms is, from the point of view of a psychologist, a controlled (or potentially controllable) variable. This holds even in cognitive psychology, I think. I used to do some research on visual search. Nearly all of this work is aimed at trying to find factors that affect the rate of search -such as similarity of target to background, statistical properties of the background, and so on. You can find things that have pretty strong *statistical* effects on search rate. So you can control search rate (or at least the average rate) by messing around with the background. To the extent that you get the effects you want in your study (effects that match your reference) then you are happy. The experimenter is typically more concerned with his own ability to control what happens than in the organism's ability to do so.

Bruce Nevin: As regards control of perceptions relative to internal reference values, statistical measures are of little use. As regards the processes by which people set internal reference values of the "social convention" sort, measures are in order that correspond to the way individuals generalize across the outputs of other members of their population. This way of formulating the problem might suggest more apt ways of formulating statistical analysis, ways that can be modeled in control-theory terms.

Chuck Tucker: Here is a recent version of the statistics so frequently used in social science.

*Relationships Among Several Descriptive Statistics**

r	r2	k2	k	E (%)
1.00	1.00	0.00	0.00	100
0.9995	0.999	0.001	0.032	97
0.9987	0.997	0.003	0.054	95
0.995	0.99	0.01	0.099	90
0.954	0.91	0.09	0.299	70
0.90	0.81	0.19	0.435	56
0.87	0.756	0.244	0.493	51
0.865	0.748	0.252	0.50	50
0.80	0.64	0.36	0.60	40
0.71	0.50	0.50	0.70	30
0.60	0.36	0.64	0.80	20
0.50	0.25	0.75	0.87	13

r	r2	k2	k	E (%)
0.40	0.16	0.84	0.92	8
0.31	0.10	0.90	0.95	5
0.20	0.04	0.96	0.98	2
0.10	0.01	0.99	0.995	0
0.00	0.00	1.00	1.00	0

*Compiled by Chuck Tucker, with the encouragement and assistance of members of the CSG (especially Gary Cziko) and Jimy Sanders.

Definitions and Interpretations of the Above Statistics

All of these measures describe two variables (X and Y) within a particular sample. It should be stressed that these descriptions and interpretations, especially those involving "predictions," are limited to a particular sample; if another sample is not a random sample from the same population, then predictions about Y will be unpredictably worse.

r is a correlation (or coefficient of correlation) which describes the linear association of one variable with another. It can also be characterized as "... a relative measure of the degree of association between two series..." of values for two variables. It varies between 1 (perfect positive correlation) and -1 (perfect negative correlation). The closer this measure is to a perfect correlation, the more confidence one has in "predicting" the values of one variable from another variable.

r2 is a measure of "explained" variance (or coefficient of determination) which describes "shared" variation, or the amount of variance of one variable "explained" by the other variable, or the proportion of the sum of y2 that is dependent on the regression of Y on X. The larger the numerical value of this measure, the more confidence one has in "predicting" the values of one variable from another.

k2 is a measure of "unexplained" variance (or coefficient of nondetermination) which describes "unshared" variation, or the amount of variance of one variable *not* "explained" by the other variable, or the proportion of the sum of y2 that is independent of the regression of Y on X. The smaller the numerical value of this measure, the more confidence one has in "predicting" the values of one variable from another.

k is a measure (called coefficient of alienation) which describes the lack of linear association of one variable with another, or the ratio of the standard error of the estimate to the standard deviation of the variable. The smaller the numerical value of this measure, the more confidence one has in "predicting" the values of one variable from another.

E is computed as 100 * (1 - k) and is called the "index of forecasting efficiency" (Downie and Heath, 1965, p. 226). It indicates the "improvement" for a prediction by knowing the coefficient of correlation (r) for two variables, as contrasted with knowing nothing about the linear association of the two variables. For example, with a coefficient of correlation of 0.71, one can "predict" the values of one variable from another about 30 per cent better (on average) than one could "predict" those values *without* any knowledge of the relationship between the two variables; *or* one has decreased the size of the "error of prediction" by 30 per cent (on average) by knowing that the correlation of the two variables is 0.71.

References

Herbert Arkin and Raymond R. Colton, *Statistical Methods,* 4th edition, College Outline Series, 1956.

N. M. Downie and R. W. Heath, *Basic Statistical Methods,* 2nd edition, Harper and Row, New York, 1965.

Bill Powers: In the social sciences, the word "theory" is used to describe a proposed statement of relationship: people who have characteristic X exhibit a tendency toward behavior Y. I would call this a proposed fact: either X's show Y or they don't. If they do, we now have an observed relationship (never mind how reliable it is) that demands theoretical treatment. The corresponding theoretical statement would tack on "because..." to the observation, and propose a mechanism that accounts for the observed dependency.

Another way in which description is confused with explanation is through the manipulation of categories. A specific instance of behavior by a specific person (Joe opens a door and walks out of the room) is converted to an instance of a class of behaviors of a class of persons (a male college sophomore exits from an enclosed space). The specific antecedent conditions are also converted to a category: "the room contains 400 people" converts to "the population density in the enclosed space is more than two persons per square yard." Now the happening becomes: "A white male sophomore exits from an enclosed space when the population density exceeds two persons per square yard." This now looks like a more general statement that will apply to more people than just Joe, and more instances of crowding in larger and smaller rooms. In many branches of the social sciences, this is considered to be an explanation.

Of course, the statistical approach and the generalization approach are used together.

The theoretician has to take the point of view of the behaving system. When you imagine being a particular control system, you realize that the actual environment is almost irrelevant: all you can know about it is contained in the perceptual signal, and the relationship of the perceptual signal to external processes and entities depends entirely on how the input function is organized. So the control system can control only its perception; the effects it has on the external world while doing so are unknown to it.

The key is not so much being able to prove that the model is right, but simply understanding how to propose processes in such a way that they *could* be right. This amounts to appreciating what sort of thing has to be accomplished by the system in order for its externally observed behavior to be as it is. We might not know how to build a general configuration-perceiver, but at least we know that the input has to be a set of sensations, and the output has to be a signal that covaries with our own sense of configuration. If *we* can think of one mechanism capable of doing this in one instance, that is better than not knowing of any mechanism. And when we have one mechanism that works, we can try to find another one that works better, seeing how the first one fails. And so it goes until we have a good model.

But we can never know that we have accomplished something in the same way that an organism accomplishes it, in every detail. For that matter, we have no reason to think that every organism of a given species accomplishes its functions in the same way as other organisms of the same species. Judging from the very large differences in brain anatomy that exist from one person to another, in fact, it's unlikely that all people are internally organized in the same way even if they behave in roughly the same way. The brain is plastic, and its organization is influenced by the experiences of a single lifetime. Our modeling is fundamentally limited by this fact: no one model can ever reproduce to the last detail the inner functioning of all examples of any kind of higher organism, because the originals are not all designed in exactly the same way. We will always be limited to modeling the "general idea" behind an organism, because that is the limit of consistency in the originals. The method of modeling is primarily a method of understanding individuals, and only secondarily a way of saying general things about all individuals. Models must always contain parameters that can be adjusted to fit the "general idea" to a specific organism.

This, naturally, has some serious implications concerning the nature of scientific research into human nature. It's usually assumed that one is dealing with a standard instance of *Homo sapiens*—the very idea of assigning such a term to the whole human race is to assert that fundamentally we are all the same. In the psychology lab, great attention has been paid to using a standard animal model—the Sprague-Dawley rat, during my formative years. If you have a standard rat or a standard person, you should get standard responses to standard stimuli. If any

human being is as good an example of *Homo sapiens* as any other, you can study groups of people as interchangeable units, drawing generalizations from the data which you assume to be measures of common underlying properties fuzzed out by uncontrolled stimuli.

But what if, below some level of observation, there are no common underlying properties? Then the whole rationale of statistical studies of populations collapses.

A generative model is one that will reproduce the phenomenon of interest by operating strictly from the interplay of its own properties. A generative model of control behavior is a control system with an input function, a comparator, and an output function, in an environment that links output to input in a specific way. There is no component in a control-system model that "controls." Control is the result of operation of a system with these functions in it, connected as specified by the control-system model, and operating as dictated by the input-output properties of each component.

So, given inputs, constraints, and parameters, a generative model must always produce some kind of behavior. We can't necessarily anticipate what such a model will do, but whatever it does is rigidly set by the properties we have given it, and by the surroundings with which it interacts. We hope that the behavior of the model will resemble the phenomenon we're trying to explain. If it doesn't (and few models do, the first time they are set in motion), we have to modify the model. That's how models grow and improve.

Greg Williams: Many physicists make a living *describing* certain phenomena, just as many psychologists are experimentalists. But modern theoretical physicists eschew the "hypotheses non fingo" stuff. And make *extrapolative, explanatory* models. Unfortunately, the bulk (well, there really aren't all that many) of theoretical psychologists still persist in making *descriptive, nonexplanatory* models solely at the level of the phenomena—rather than *generative* models of *underlying* mechanisms. "If you do basically the same procedures again, the organism will do basically the same thing." The weasel word is "basically," because these folks cannot circumscribe its bounds. So, the turn toward statistics.

I claim that the only reasonable answer to Hume's inductive skepticism (i.e., why should the sun rise tomorrow?) is making generative models which "hang together." Hypotheses non fingo leaves open the possibility that matter might disappear at any moment, since it can't predict that it *will* disappear at a *particular* moment. Contemporary generative modeling in physics says there's no "disappearing at such-and-such-a-time" relation within its (modeled) structure, so give us a break from your concocted philosophical "possibility" tales, Hume!

Descriptions at the behavioral level don't explain behavior, and descriptions at the sociological level don't explain sociological observations. A description in the Skinnerian vein would be that people show certain behaviors which are correlated with certain outcomes which can be lumped into a class termed "rewards" (of course, this begs the question of why some outcomes end up in the class and others don't, which can only be answered by invoking structural constraints embodied in organismic physiology). However, such description at the individual behavioral level is, I claim, what counts as an explanation of sociological phenomena. It appears to me that people generally accept accounts such as the following as explanation: How come the voting turnout rates of the poor are much lower than those of the rich? Continuing in the Skinnerian vein, for argument's sake, it is because some individuals receive few rewards from voting (and reduce their voting rates), while other individuals receive many rewards from voting (and keep voting). That is pure description at the individual behavioral level. But it isn't an explanation of the sociological phenomenon, just yet.

What must be added to the description at the individual behavioral level, as given above, is description at the sociological level, to wit: most individuals belonging to the class "poor" *actually are* in the first (few rewards from voting) group described above, and most individuals belonging to the class "rich" are in the second (many rewards from voting) group. Now we can *deduce* that the poor will come to vote less frequently than the rich. We have a generative model at the individual behavioral level which, coupled with a description of certain conditions observed at the sociological level *(but not the phenomenon to be explained at that level),* results in an explanation of the sociological phenomenon in question.

In this example, pure faith (precisely as criticized by Hume!) is the *only* basis for believing that *tomorrow* the poor will continue to vote less frequently than the rich, since there is no basis except a belief that "what was, will be" for extending the "functional relationship" (Skinner's term) from past correlations between voting and reward to future frequencies of voting. Without limits on the generalizability of such relationships, which I claim can only be placed by generative models at the level below individual behavior, you're in free fall. One might call it the free fall of statistics—comfortable, until you meet a boundary. Then, *splat!*

Bill Powers: I think that Popper's idea of "falsification" is predicated on the prevailing view of theories as being primarily statistical. Statistical theories don't propose any models, so there is no positive way to verify a theoretical statement. All that significance does for us is to assure us that the experimental results probably didn't happen

by chance. There is no a priori or logical argument against the result being a chance occurrence; it is reasonable to admit the possibility that chance played a part. This negative conclusion doesn't tell us that the hypothesis is reasonable, connected to a systematic world, or useful in any context other than the original experimental conditions.

Models, on the other hand, are tested by changing the conditions and verifying that the model still behaves as the real system does under the new conditions. The model provides an a priori systematic reason for the system to behave in some new way under new conditions, and commits us to specifying exactly what that new way will be. When the real system does behave that way, this is a positive indication of the model's worth. Of course, one could argue that there is still a possibility that the real system behaved in the new way by chance, but if the standards for acceptance are set as high as they are in the physical sciences, this possibility goes beyond the bounds of reason: there's a qualitative difference between p less than 0.05 and p less than 0.0000000005. More likely is the possibility that the real system behaved in the new way for a reason other than the reason for which the model behaved that way. This does not involve chance; it says merely that the model needs to be modified, and that sooner or later circumstances will reveal the needed change. The modeling approach is fundamentally systematic, not statistical. Modelers assume that the underlying processes, whether we have correctly identified them or not, are systematic.

Thus, I would say that I use the criterion of *testability,* not falsifiability. Falsifiability is a subset of testability that considers only the possibility of rejection. Testability also demands that hypotheses that are not rejected be accompanied by quantitatively correct predictions of new behaviors in new circumstances. The kinds of theories Popper was thinking about never went that far.

A true science needs continuous measurement scales so that theories about the forms of relationships can be tested. This means that correlations have to be somewhere in the high 90s. True measurements, with normal measurement errors, require correlations of 0.99 upward. If this were universally understood among scientists, two things would happen. The first is that most statistical studies would end up in the wastebasket. The second is that the good studies would be done again and again, with successive refinements to reduce the scatter, until something of actual importance and usefulness was found.

One of my objections to the statistical approach to understanding behavior is that after the first significant statistical measure is found, the experimenter quits the investigation and publishes. If you get a correlation of 0.8, p less than 0.05, your next question should be, "Where is all that variance coming from?" If you set your sights on 0.95, p less than 0.0000001, you won't quit after the preliminary study, but will refine the hypothesis until you get real data.

Quantitative methods in conventional psychology don't deal with quantitative data, despite the tremendous sophistication of statistical techniques. When you consider the models used in physics, where the systems are claimed by some to be "simple" relative to organisms, you find extremely complex structures in these models, extending from simple algebra, through systems of hundreds of differential equations, to tensor calculus. When you look at the models used in psychology, you find basically y ax + b. Of course, in order to see whether this model represents any regularity in a data set, you might have to apply very complex techniques for extracting signal from noise, but the basic model being tested is elementary, if that. So if the subject matter of psychology is so complex, why do psychologists try to handle it with such simple models?

The place where psychology is the least quantitative is in the data-taking stage. Most data exist in the form of simple and artificial events, which either occur (1) or don't occur (0). The behaviors investigated are characterized in only the crudest qualitative ways; quantitative continuous measures of behavior almost never occur except in psychophysics.

When I read the psychology literature, I see almost nothing being investigated that strikes me as a real phenomenon. Even when something real-looking is investigated, I see no quantitative measurements being made. The *only* quantitative analysis that shows up in most articles is the statistics, which takes for granted that the data are about something and offers no explanations at all.

I think that the control-systems approach, which is fundamentally quantitative, offers the promise of handling even complex behavior in a way that is as clean as the methods of physics. I don't buy the idea that psychologists have the problems they have because of the complexity of the subject matter. I think their problems come from a primarily non-quantitative, idiosyncratic, and disorganized approach to observing human behavior, and the acceptance of very low standards for what will be considered a fact of nature. The latter bothers me the most. You can't base a science on facts that have only a 0.8 or 0.9 probability of being true. Such low-grade facts can't be put together into any kind of extended argument that requires half a dozen facts to be true at once. You need facts with probabilities of 0.9999 or better—if you want to build an intellectual structure that will hang together. I don't think that psychology has come anywhere near meeting that requirement, individual cases aside. I would argue that we do not yet have any *science* of psychology.

The Control Systems Group is a membership organization which supports the understanding of cybernetic control systems in organisms and their environments: *living control systems.* Academicians, clinicians, and other professionals in several disciplines, including biology, psychology, social work, economics, education, engineering, and philosophy, are members of the Group. Annual meetings have been held since 1985. The CSG Business Office is located at 73 Ridge Pl., CR 510, Durango, CO 81301; phone (303) 247-7986.

The CSG logo shows the generic structure of cybernetic control systems. A Comparator (C) computes the difference between a reference signal (represented by the arrow coming from above) and the output signal from Sensory (S) computation. The resulting difference signal is the input to the Gain generator (G). Disturbances (represented by the black box) alter the Gain generator output on the way to Sensory computation, where the negative-feedback loop is closed.

Back cover

Closed Loop

Threads from CSGnet

Fall 1992

Volume 2

Number 4

Front cover

Closed Loop

Threads from CSGnet

Fall 1992 Volume 2 Number 4

Edited by Greg Williams, 460 Black Lick Rd., Gravel Switch, KY 40328

CONTENTS

Members of the Control Systems Group receive *Closed Loop* quarterly.

For membership information, contact Ed Ford, 10209 N. 56th St., Scottsdale dale, AZ 85253; phone (602) 991-4860.

CSGnet, the electronic mail network for individuals interested in control theory as applied to living systems, is a lively forum for sharing ideas, asking questions, and learning more about the theory, its implications, and its problems. The "threads" in each issue of *Closed Loop*, stitched together from some of the net's many ongoing conversations, exemplify the rich interchanges among netters.

There are no sign-up or connect-time charges for participation on CSGnet. The Bitnet address is "CSG-L@UIUCVMD" (use no quotes in this and the following addresses); "CSG-L@VMD.CSO.UIUC.EDU" is the Internet address. Messages sent to CSGnet via these addresses are forwarded automatically to all participants. Via CompuServe, use the address ">INTERNET:CSG-L@VMD.CSO.UIUC.VMD" to reach the net.

Initially, you should send a note to the network manager, Gary Cziko, at the Internet address "G-CZIKO@UIUC.EDU" (Gary's voice phone number is (217) 333-4382).

Conflict, Belief, Standards: Part I

Ed Ford: I have been meeting with a group of my former students every month for the past year or so, and recently we decided to attempt to apply perceptual control theory (PCT) to their work, specifically to the way they organize their staffs, run their organizations, deal with people, etc. Anything they do, from running a staff meeting to a group meeting, from setting standards to dealing with individuals, is to be done with PCT in mind. As a first step, they have to think of others as living control systems. Each time we meet, everyone will say what they have tried at their various places of employment. We'll review what they have done, whether it has been effective, and whether their knowledge and application of PCT have helped. The more people on the net talked about modeling, the more I realized that I had to get some modeling going of my own.

The jobs held by this group are most interesting. One is a superintendent of schools for Arizona's juvenile residential correctional system. Another heads a residential treatment center for sexually abused 7- to 11-year-old girls. Another is in charge of counseling and training at a residential treatment center for teenaged boys. Another works with the toughest behavior problems in a school district. Another is an adult probation officer. Another works with the most violent people in a state mental hospital. Another runs groups for women with various types of problems.

These are the kinds of settings where the rubber hits the pavement. These people don't play games; they're serious about succeeding and doing a better job. Both the supervisors and staff in these kinds of settings generally are all looking for a better way to make their job easier, more efficient, and more satisfying.

One example we're trying is developing a way to get your staff to do a good job. I think you first have to get each member to explain what he / she has done successfully, what they are presently working on, where they need help, and from whom. Problems are brought up, and a person from the staff volunteers to take the responsibility for researching and bringing back the results to the group for a group decision, yet my experience is that there is an ongoing recognition of where the final decision rests. In this kind of system, people begin to perceive that they have some control for setting reference signals and providing input to the system. They become much more cooperative and much more willing to look for a better way. The staff members are each operating as individual control systems, but each finds they can, through a cooperative structure, provide inputs to what is happening, know what other systems are perceiving and controlling for, and be able to control more easily for what they want and for their specific tasks. When the commitment to some overall concept of the organization is missing, the inevitable result is internal conflict.

Each person in our group is attempting to teach his/her staff PCT. The real key is our working together at our monthly meetings to reflect on what each is doing (using PCT as a basis), giving our own inputs, and then watching the results. I guess this could be called modeling in the real world. I have a close friend who is a Catholic priest and was recently assigned his own parish. He already has invited me to give these ideas a try at his parish.

Bill Powers: Ed Ford mentioned his monthly meetings with people interested in using control theory in the real world. In a phone call, he asked me to say something that would help his group "make models." I replied that what his group can do is probably better called "testing models," so that is what I'm writing about here. I'll digress at the start to introduce some background on the concept of prediction. The first part of this development is intended to amuse experimenters; the second part gets to practical matters.

On Testing Models

Part 1

Some time ago I remarked that the most common model in psychology is a cause-effect model in the form of a regression equation. The hypothesis is that the effect depends on the cause linearly, as in y = ax + b. To test this model, you'd take the values of a and b determined from a formal study and try to predict new values of y from new observations of values of x.

David Goldstein commented that this concept of using a model for predictions is not the way such findings are used in psychology. Once the regression line is drawn through the data points, that's the end of it. The model equation describes the data, but isn't then used for predictions.

On thinking this over, I agree that no formal use is generally made of the regression equation, but the findings are certainly used to predict individual behavior. Suppose the dependent variable y is a clinical measure of depression, and the independent variable x is a depression-factor score on a personality test. In computing the correlation between the test score and the clinical measure (in a study of many

people), a regression equation of the form y = ax + b is the basic premise behind the correlation calculation. If the correlation is positive and statistically significant, the conclusion drawn is that depression is predicted by the test score. Then the test is administered to a new individual (presumably from the same population), and if the depression-factor score is high, the person is diagnosed as depressed.

This isn't a formal application of the regression equation: you don't say that a test score of exactly 7 predicts a depression of exactly 25 units on the clinical scale, even if that's what the regression equation says. But a person who measures 15 on the test score would be judged as more depressed than a person who measures only 3. So while the slope and intercept coefficients aren't explicitly used, the general trend is implicitly used, and there are semi-quantitative judgements made.

The scatter in data of this kind is so great, of course, that literal application of the regression equation would be silly. The prediction for any individual when the correlation is as low as 0.8 would be seriously wrong most of the time, often even getting the sign of the relationship wrong for one person. The only correct way to make a prediction would be to begin with another equally large sample of the population and do the whole study again. You would predict that the same regression coefficients would be found.

But there is an urge to predict for individuals, and the form of the urge follows the regression line: a higher clinical score ought to predict a more severe depression. While it is folly to give in to this urge when the data are so bad, the motive behind doing so is consistent with the principle of modeling.

If the principle of modeling were followed through formally, the regression line would indeed be used to predict behavior. If the line has the equation y = 3x + 5, and the depression-factor test score for a new individual is 4, the model predicts that a clinical evaluation of depression will come up with 17 on the clinical scale for that person. To follow the test through, one would then submit the person to the same clinical evaluation as used in setting up the model, and see what number actually results.

Suppose the actual depression measure is 12 on the clinical scale. This is a deviation of -5 units from the value of 17 predicted from the test score, for an error of -29 percent. Is that good, or is that bad? The answer depends on how important it is to get the evaluation exactly right.

Of course in this case we know the clinical measure of depression, and if we believe it, we can just ignore the test score and the prediction. But what if we want to make the diagnosis on the basis of the test score alone? Now the generally expected error for an individual prediction becomes relevant. If you're going to prescribe electroshock therapy that will most likely severely disturb the person's life for many years, maybe even permanently, you might decide that a 29 percent error is too large to allow. Perhaps even an error of 5 percent would be too large if the person is a borderline case. On the other hand, if you're going to prescribe a tranquilizer that won't do any permanent harm even if the person isn't really depressed, then perhaps you can allow errors as large as 29 percent.

I've gone through this to illustrate that prediction errors can't be judged as good or bad without taking the context into account. But what if the context is that of testing a general model of behavior? Now the actions taken as the result of a diagnosis are no longer in the picture. All we want to know is which theory is better. Now the errors of prediction under different models are judged not against practical standards, but against each other. The smaller the expected error, the better.

I've also tried to show that even in standard approaches, the method of modeling is there just beneath the surface. It's probably not mentioned much because the predictions made from literal application of the mode—the regression equation—are so poor. But the model is there. It's that model that we have to compare against the control theory model, and the way we do the comparison is through making quantitative predictions using the actual form of the model.

Let's look at the rubber-band experiment. Suppose we just measure the position of the experimenter's end of the rubber bands and of the subject's end, designating the positions as a and s. Let's confine the experiment to a line, so we consider only one dimension. The zero point on the line can be chosen arbitrarily, with all measurements made relative to that zero.

If we now measure the positions a and s over a long series of movements by the experimenter, we will obtain a data set consisting of pairs of values of a and s. We can do a correlation between a and s. From the normal calculations, we can derive a regression line.

The regression line will have the form s = ae + b. The position of the subject's end will depend on the position of the experimenter's end. If the rubber-bands are identical, the coefficient a will be very close to -1. Half of the intercept b will correspond to a position on the line. That position will be the average position of the ends of the rubber bands: with a = -1, we will have (s + e) = b, or (s+e)/2 = b/2.

In fact, half of the intercept b will turn out to be a position nearly underneath the knot where the rubber bands are connected. The knot, as it will turn out, remains very nearly at the position b/2 all during the experiment.

There's a moral to this story, but it's not quite obvious yet. The first part of it is that when you do a stimulus-response experiment in the usual way, to get a regression coefficient, you can *sometimes* translate it

directly into a control-system experiment. If you find that the intercept b corresponds to something in the experimental situation that's remaining nearly constant at that value, you've found a controlled variable —actually, by finding its reference level first.

The second part of the story concerns the accuracy of the prediction. The stimulus-response prediction will be accurate only if the two rubber bands have identical characteristics, or strictly proportional characteristics. If their characteristics are different, the correlation coefficient you derive from the data corrected for the different rubber-band properties will be very much higher than the one derived from the model s = ae + b, which assumes identical rubber bands.

Part 2

In testing the control-system model, the basic procedure is to assume that all behavior—without exception—is control behavior, predict behavior on that basis, compare the prediction with the appropriate data, and let the match or mismatch decide the issue. You can never prove that a particular control-system model is the only correct one, but you can show that it is incorrect.

Considering the low correlations that are found in stimulus-response experiments, it might seem hopeless to substitute a PCT model for the linear regression model. When the data are that noisy, how can any clear decision be made? This objection, however, assumes that the stimulus-response experiment has correctly represented the data. While we can't prove that *all* stimulus-response experiments could be translated into relatively noise-free PCT experiments, there are excellent reasons to think that this can be done in a significant number of instances, maybe even most instances. To do this, however, requires some changes in viewpoint that might be hard to achieve.

A stimulus-response "fact" is expressed as an effect of a cause. Doing something to a person results in that person's doing something else. If the relationship expressed in this "fact" isn't clearcut and quantitative, then the control theorist has to start asking questions about the data.

The basic question is: what is it that was affected by the "stimulus" that was also affected by the "response"? If you utter encouraging words to someone, and that someone then shows added efforts to achieve something, you have a stimulus-response relationship. Now you have to try to guess: what did the encouraging words affect that was affected *equally and oppositely* by the increased efforts?

Equally *and oppositely?* There's the rub. You would like to think that there is something you said that helped this person do better. But control theory says that if your words of encouragement had some regular effect on the person's behavior (apparently), that behavior was aimed at *counteracting* your influence. If this is true, then you don't have the control over the person's behavior that you thought you had, even for the good. You are seeing yourself as helping the other person to do better. The other person, however, is seeing the situation differently: you're disturbing something, and the other person is acting to cancel the effect of the disturbance.

This might not be true, but if you're going to test the PCT model honestly, you have to pretend it's true and try to make sense of it. You can't test a model if you don't follow its logic faithfully and literally as far as you can. You can't look ahead and think, "If PCT is right, then I haven't been helping people the way I thought I was—so PCT must be wrong." You have to be prepared to change your ideas about anything at all. Otherwise your reasoning is just a sham.

Let me give you a real example from my high-school days. We had a coach, named Coach, who was tremendously popular, a great guy. We all loved him and wanted his approval above anything else. Coach would say, "You can do better than that, I know it—just give it one more try and you'll make it." And by golly, we'd give it one more try, and we'd make it, sometimes.

Now it would seem that his encouragement and belief in us caused us to try a little harder than we thought we could, so we achieved something we couldn't do before (sometimes). I suppose that Coach looked at it that way, as any reasonable person would. But I can tell you that from inside at least one person (and at the time I guessed this was true of a lot of the others), it wasn't all that nice.

The basic problem was that Coach went around all the time saying to people, "What you're doing isn't good enough to please me." That's what "you can do better" says. I was already doing better than I thought I could, in number of pushups, speed of climbing a rope, time in the 40-yard dash, or whatever. And I was damned tired and hurting and not necessarily interested in doing any better. I liked physics a lot better than physical education. But there was Coach telling me that he didn't like what I was doing. That mattered to me. So I got myself together and made it really hurt, and I felt great—because then Coach wasn't displeased with me, not because I'd achieved something I wanted, but because I'd done something to counteract his disapproval.

From Coach's point of view, he had helped me put out that extra bit of effort to surpass my previous achievements. No doubt if I had continued to go along with this, worked out, built up a lot of strength, learned the football playbook by heart, and all of that, satisfying the coach more and more all the time, I might have achieved even more. I might have been a college football star; I might even have become a professional football player and ended up as a coach myself, by now. I might be bold, aggressive, commanding, and rich. But I certainly

wouldn't be writing this. I also wouldn't be the Bill Powers you know.

What actually happened was that many of us simply gave up on pleasing Coach because we didn't buy the goal. It wasn't pleasant to do that—to decide we were trying as hard as we cared to try toward that particular end, and that we would simply endure the disapproval. We still loved Coach, and we tried to fend off his disapproval by seeming to try harder. But the price was too high to really do it. When Coach was called into the Navy and left in 1944, there was a huge tearful farewell ceremony for him, and I'm sure that amid the sorrowful participants there were many hearts filled with relief.

To apply the PCT model, this is the sort of thing you have to think about. It's especially difficult when the hoped-for effect on a person is beneficial. There's an almost inescapable tendency to suppose that what you think of as beneficial is also considered beneficial by the other person; that what you consider harmful is also thought harmful by the other. Coach would have been completely baffled by the present discussion. He would have said "Well, you did try harder, didn't you? And you did do something you thought you couldn't do, didn't you? What's so bad about that?"

The stimulus-response viewpoint encourages this sort of naive projection of one's own goals onto the behavior of others. I shouldn't even call it the "stimulus-response" viewpoint. It's really this viewpoint, adopted innocently by well-meaning people who have never heard of stimuli and responses, that led naturally into stimulus-response theory.

To test the PCT model in real life, you have to be prepared to follow its logic all the way. Forget about whether the "response" is good or bad. The question is how to find the controlled variable, the thing that is disturbed by what is done to the person, and is protected against more disturbance by the action that the person takes. If you find such a controlled variable, you will understand that person far better than you did before. If you want to help that person, you might even find out what he or she really wants and figure out ways that person could get there.

It's possible that you won't find any such controlled variable in a given circumstance. But if you don't look for one, you will certainly not find one, even if it's there staring you in the face.

The basic message here is that to test PCT, you have to make predictions from it and from nothing else. You have to follow out the logic even when it seems to say things you don't believe. Then you have to look carefully to see whether, in fact, the prediction holds true. This requires being consciously open-minded and willing to take a chance. You simply have to trust that if the theory does predict correctly, you'll be better off knowing what it predicts than not knowing, letting the chips fall where they may.

Martin Taylor: In his interesting example of the kindly coach who induced him to try so hard that it hurt, but produced great performances, and induced him to hate the coach because of it, Bill argues that all attempts to control are just introducing disturbances that are resisted in maintaining reference percepts, and that inevitably such resistance is accompanied by resentment or other bad effects (I might exaggerate, but that's what I get out of the example).

I quite agree with the first half of the conclusion, but not with the second. There is no technical distinction between the alteration of the error signal in an elementary control system by changing the reference versus by changing the percept. Each results in a determinate error signal that results in behavior that reduces the error (assuming a well organized control-system hierarchy). The coach played on this by assuming that Bill had a reference to be liked and admired by the coach, and by causing Bill to perceive that this was not the situation, though it could be. The coach also presumably assumed Bill had another reference (shared by athletic overachievers) that he should do as well as his *body would* permit. Bill asserts that he did not share that reference. If he had, then the coach's behavior would have induced percepts that caused errors with respect to each reference that the same behavior would have satisfied. But since Bill did not have the "excellence in athletics" reference, the "hurting" behavior helped to satisfy only the "find favor with the coach" reference, and it conflicted with the reference most people hold: "feel good in my body."

I don't think it is necessarily true that this sort of conflict leads to resentment and bad feelings. The better a system is controlling, the lower the errors within it, almost by definition. Rick Marken has pointed out that the errors don't go to zero if there is non-orthogonality within a control hierarchy, so that, to some extent, behavior that helps reduce one error increases another. Such conflicts are almost inevitable in a complex hierarchy, especially one in which there are fewer final degrees of freedom for control than elementary control systems at any one level. The human muscular system provides a good example: some 400-800 muscles (I don't know the exact number, but that's the range) control around 125 degrees of freedom for joints, face, voice, and so forth. There are two ways of resolving the conflict: mutual control, such as in opponent muscle pairs (one zeros its control while the other works), or tension (each tries to achieve its reference, and a balance between then is achieved).

I think tension and conflict are desirable, if not overdone, enabling a control system to react promptly to changes in perceptual situations. This is analogous to the temperature of a thermodynamic system. Zero conflict means a system perfectly organized for the disturbances the environment presently provides—the system is frozen and will not necessarily be able to respond well to new types of disturbance. Some tension means two things: the system is ready to move fast in many directions, and, equally important, it is prepared to reorganize if Bill's notion about reorganization being driven by accumulated error is correct. So a system with tension and conflict will be more robust than one that is placidly content.

The end-point of this line of thought is that we should have evolved to be happier with some level of disturbance and internal conflict different from zero than with a bland, disturbance-free environment or an environment that we have totally under control. Bill's coach was right, but perhaps went too far. Mild social control of that kind is what we like. We want to do well for other people, but we do want to find that we can reach the reference level of satisfying them without at the same time working too hard (diverging from other reference levels). I suspect that many marriage problems arise from a perception of inability to satisfy the partner, despite excessive efforts (which might be in the wrong direction, demanding reorganization).

Thus, tension, conflict, and uncorrectable disturbance are good, but not in excess.

Bill Powers: Martin, you've sort of taken off at right angles to the line of thought I was developing. The "Coach" example was meant to illustrate how an apparent stimulus-response relationship (encouragement results in doing better) can lead to quite a different interpretation when explored from the viewpoint of control theory. I wasn't trying to generalize from the particular way I and (probably) others dealt with Coach's urging us to overachieve. With another person or in another circumstance, a similar encouraging remark leading to improved performance could work in a different way. But it will never be a cause-effect way. My point was that to test control theory, you have to think of possibilities other than the surface appearances.

Since I'm into high-school stories, I remember another instance with a mathematics teacher. I didn't much like or dislike this teacher—he knew his stuff but wasn't strong on making things clear. The class was doing an exercise, each person trying to prove a trigonometric identity. I was stuck—something was wrong, and I didn't know if I was even getting close. The teacher was going around the room seeing how everyone was doing. When he got to me, he said, "That's fine, you're almost there." This told me that I hadn't made any mistakes so far and was headed in the right direction. So I stopped worrying and went ahead and finished the proof, my first one. That felt nice. The 60th proof didn't feel so nice.

Apparent stimulus-response relationship: he said what he said, I then went ahead to reach the goal. Cause and effect? No. Information. I wanted to know if I'd made some stupid mistake, and he told me (in effect) that I hadn't. With that information, I could stop looking for a mistake and devote my efforts to something more productive. I didn't finish because I liked the teacher or in order to please him. I finished because I wanted to be able to prove the identity. His remark wasn't a disturbance of something I was trying to control; it provided a missing perception so I could get unstuck from looking for a nonexistent error.

My Coach example was one in which the apparent stimulus actually did disturb something I was controlling for, and my response opposed the effect of the disturbance. The result was to put a very different light on what seemed like a simple S-R chain. That's all I was trying to show —not that there's something inherently bad about encouragement, or that being as pushy as Coach was necessarily leads to resentment and bad feelings. In fact I never resented Coach; not many did. He was a nice guy. I just resisted him. I regretted not wanting to live up to his expectations, but not enough to change my mind.

Regarding your comments on conflict: Conflict doesn't "lead to" anything in particular. What it leads to depends on how you resolve it or fail to resolve it. Most conflicts are unimportant; we just shrug and turn to something else, or we go into a little fit of reorganizing and think of a different way out. This happens all the time; we have natural machinery for resolving inner conflicts, and it usually works very well.

The degrees-of-freedom problem doesn't normally cause conflict because we've learned to use only those control systems that are compatible when working at the same time. The balancing of reference signals contributed by many higher-level systems isn't a conflict unless one of the higher systems is unable to keep its own error reasonably small because of the interference of other systems at the same level. The usual case is that all active higher-level systems keep their errors small, despite the fact that no one lower-order system's reference signal is the exclusive property of one higher-order system. The systems just find the analog solution of the simultaneous equations, and they all are successful.

When opposing muscles are used to control limb position, there's no conflict. In fact, there are two controlled variables that are independently adjustable: for the tendon reflex, one is the difference between the tensions in the two muscles, the other is the sum. The sum-of-tensions signal is controlled to produce a specific muscle tone. The difference signal controls the net applied force. Because the muscle is highly nonlinear, the sum (muscle tone) signal effectively alters the spring constant of the combined muscles near the zero-error condition, thus adjusting the static loop gain of the tension control system (and also the stretch control system).

Conflict is a problem only when it concerns some variable important

to the organism, is severe, and goes unresolved for a long time. That's what brings the clients to the therapist or counselor. Serious conflict destroys control or reduces its effective range to the point where it's not sufficient for the purposes normally served by the control systems.

A control system that keeps its error very small isn't likely to be "placid and content." It's able to keep the error small because it has a very high loop gain. This means that even the smallest disturbance will evoke an opposing effort, and that opposition will keep the controlled variable nailed to its reference condition. When you're driving a car along a mountain road with a washout on the cliff side, you tighten up that control system so the car stays precisely on the path you've picked to squeeze past the danger point. I don't think that "placid and content" describes that control system. But it's not in conflict, either: if it is, you have a problem because you won't be able to move the wheel as much as if there weren't any conflict.

There's a problem with your suggestion that "a system with tension and conflict will be more robust than one that is placidly content." The problem is that reorganization will start because of the chronic conflict. As a result, precise control will become impossible: the parameters of the control systems are going to be changing at random. What you get is a jittery and unpredictable control system that could literally do *anything* without warning.

Just because of neural response curves, I can believe that some slight amount of tension would help with rapidity of response to disturbances, because near zero signal, the slopes of the functions will be very low, and the loop gain will be low. But this is relevant only when the control point is set to zero and there are no disturbances. Most reference signals specify values of perceptual signals that are far from zero—somewhere in the normal range between zero and maximum. And there's normally some amount of disturbance to raise the error signals above zero, if only gravity. In those cases, there's no advantage to conflict, because conflict won't raise the sensitivity or speed of the system and will only reduce its range of control. I think that the best state to be in for possible action is one of alertness and calm. You should feel just a little zingy, but you certainly shouldn't be in white-knuckle conflict with yourself. You want everything working in the same direction.

So I guess I agree with your concluding remark: tension, conflict, and uncorrectable disturbance are good, but not in excess. I would figure something like five percent of the range of control. The rest of your reserve you would want to save for affecting the environment.

Some uncalled-for remarks on social conflict: In the background, I suspect, is an idea that competition is good for us. Up to a point, while it's fun, I agree. We like to set problems for ourselves and solve them, and get better at solving them. But competition as a way of life doesn't work that way, except for a few winners. A social system based on serious competition is just a step from violence (in the United States, a very short step). The losers vastly outnumber the winners: we end up with a society of losers, winners being anomalies. In situations where the terms of the game determine that only a few can win, chronic losers can get very nasty; in fact, they tend to abandon whatever social principles there might be that make civilization better than life in the jungle. I don't think that the price is right. Competition—interpersonal conflict—is the lowest level of social intelligence. I don't like to admit that even a little conflict can be a good thing, because we've accepted a *huge* amount of conflict as good and natural for far too long. It's time to get smarter.

David Goldstein: Bill, can we follow through on your Coach example? What was the controlled variable for you, the degree to which Coach was pleased? Coach's comments led to your perception that he was not pleased. The increased effort was designed to increase the degree to which he was pleased. However, by making more of an effort, you caused yourself some physical discomfort and took some time away from activities which you enjoyed more. So you were really in conflict: you wanted Coach to be pleased; but you didn't want Coach to be pleased, because that meant discomfort and time away from more interesting activities. Suppose that you said to Coach, "I am doing as well as I want to for myself. If that really bothers you, I will be glad to quit." The Coach could say, "I want you to quit, it really bothers me." Or the Coach could say, "I don't want you to quit, it doesn't really bother me." In either case, you would be pleasing the Coach.

What would you have done with each answer? I don't believe that you would have quit, even if Coach gave the first answer. This suggests to me that you played football for several reasons other than to please Coach. I think you would have felt more relaxed about not putting out more effort if he gave the second answer. But you might have tried a little bit harder just to show Coach that you cared about his opinion, even after you made your remark.

What was the controlled variable for Coach? I guess it was a principle-level generalization, something like "never accept the initial effort, always prod the players to do better, and accept whatever additional efforts they make." Underlying Coach's comments, perhaps, was the thought: "I think very highly of you. I can see potential in you which you don't see in yourself." Coach was probably controlling for increased efforts beyond the ones players could make comfortably. No pain, no gain. Effortfulness. Commitment. Coach's comments led to your raising the gain. When a player did this, his performance was probably close to his potential. Coach wanted each player to do the very best that he could.

Bill Powers: David, a conflict is expressed at the level where there are different goals for the same thing. I wasn't in conflict about liking Coach and wanting to please him, or about wanting to be doing more interesting activities. I even had some personal goals about getting stronger and getting better at athletic things. The conflict came when Coach pushed me *toward* one of my goals, and in fact *past* it. To be more exact, he made his approval contingent on my trying for a goal of physical achievement that was a lot higher than my own goal for physical achievement. To please myself and fit my athletics time in with all of the other things I had goals about, I made a certain amount of effort for a certain amount of time, and I was satisfied that I was doing pretty much what I hoped to do. Then Coach, for his own reasons, decided that it would be good for me to try harder, spend more time in the gym, become not just a social football player but a dedicated one, and so on. Maybe he saw some physical talents there and felt they should be developed more (of course that was what he was hired for and what he thought worthwhile in life).

The net result was that in order to maintain a good relationship with Coach, which meant mainly that this admirable guy would express approval of what I was doing, I had to reset my goal for athletic prowess at a level higher than what my own values recommended-at least temporarily. But doing that resulted in errors in my social life (too much time at the gym and football practice, more physical discomfort than I was willing to experience, a shift in self-image that didn't fit with my picture of me as a physicist, etc.). So I wanted to try harder and become better in order to please Coach (and get whatever other benefits would come from going that way, like being a football hero, scaring off people I was afraid of, etc.), and I wanted to try less hard in order to be more comfortable, have more time for my girlfriend, be with my other friends, tinker around with "scientific" projects in my room, and so on. It all came down to wanting to try harder (for one set of reasons) and wanting not to try harder (for another set of reasons). That was the focus of the conflict: I wanted to try harder, and I wanted to not try harder. I couldn't do both. My solution was probably a typical adolescent solution. I gave the appearance of trying harder without trying harder, and let Coach believe (or so I thought) that I just didn't have the talent he thought I had. At least I was convinced that he believed it, and so my conflict was resolved.

I could have quit football, but not physical education, which was required. And don't forget that one reason for going out for football (among several) was to please Coach! I didn't want to please him just to get him off my case. I liked him and admired him. The only problem was that he wasn't satisfied by that—he didn't just say, "Glad to have you on the team." He did say that, to my great pleasure, but then he went on to demand more of me than I was willing to give. A great way to turn people off is to "encourage" them to do more than they want to do.

This certainly wasn't the only area of my adolescent life in which there was a conflict that rested on wanting to be approved of and liked, a conflict that led me to do things that caused errors in my self-image and self-esteem, but satisfied (or would have satisfied if there had been no conflict) other desires. I was always aware of these conflicts, but I didn't really have any good ways of resolving them. Finding those ways took me another 30 or 40 years of messing around at random.

One of the things I was trying to get across with my example (aside from the main one, which was reinterpreting an apparently beneficial cause-effect situation in PCT terms) is that "helping" people doesn't always help a whole lot. Like the adolescent me, most people are already in the middle of trying to fit their various goals into one coherent structure. When you try to force them toward what seems like a worthy goal, you inevitably cause conflict with other goals. You become part of the conflict situation. Of course you're trying to help, but you're forcing the person in a direction that person has probably already tried to go, or in which that person has gone far enough to meet the goal. If the person hasn't spontaneously gone farther in that direction, it's because doing so would violate other goals. If you really want to help, you'll help the person find out what is keeping that person from achieving all goals that look attractive, not urge trying for any particular goal just because, in your life, it has proven to be worth pursuing for you or for others. And helping doesn't mean urging people to go *past* their goals.

My life has been full of well-meaning people who just knew that I could achieve great things of *the kind they thought worth achieving.* If I'd gone along with all of them, I would have been a physicist, a writer, an athlete, a biologist, a neurologist, a cop, a teacher, a debater, a poet, a gardener, an engineer, a playwright, and so on-but just one of these things and nothing else. Of course, I was smart, so I showed a little aptitude in all of these directions. But I was never into heavy competition—I didn't want to be *the greatest* in any of those fields, or in general. All you have to do is show a little interest in someone's field, and that person becomes convinced that you share the same obsession and wants to do you the favor of helping you achieve fame and fortune in that field. People are really very generous in this way. But they aren't really "helping." They're really trying to validate themselves, their own choices of goals.

You say, "Coach was probably controlling for increased efforts beyond the ones players could make comfortably. No pain, no gain.

Effortfulness. Commitment. Coach's comments led to your raising the gain. When a player did this, his performance was probably close to his potential. Coach wanted each player to do the very best that he could." Probably something like that. It's a common viewpoint. It's also a narrow one, because what's a person's "potential"? Potentially, I could have been a great criminal. Potentially, I could have become a Hulk with deltoids like balloons. I could have become a pro football player. I could have become one of the world's great atomic physicists (at least one of my classmates did). People who see "potential" in you aren't considering your values, but theirs. They're also communicating, in a not so subtle way, that they don't think much of what you've done already. David, you have great potential as a psychotherapist (if you'd only just try a little harder). How's that grab you?

Rick Marken: Martin alluded to the potential value of a moderate level of conflict. Bill Powers agreed that some small amount of conflict might help in some situations. Bill says: "So I guess I agree with your concluding remark: tension, conflict, and uncorrectable disturbance are good, but not in excess. I would figure something like five percent of the range of control. The rest of your reserve you would want to save for affecting the environment." I'd like to point out that Bill's "five per cent" figure is based on experimental evidence. Nearly two years ago, I stumbled on the fact that people can control better when the disturbance to a controlled variable is caused by the output of another control system than when it is simply the result of causal processes. I had subjects do a tracking task where the disturbance (d(t)) was the output of a low-gain control system that was trying to keep the cursor at the center of the screen. This control system was in conflict with the subject, who tried to keep the cursor at another ("target") location on the screen. The subject always "won" the conflict, because the opposing control system had such low gain. What I wanted to show was that the output of the opposing control system would be dealt with by the human subject just as a disturbance -as though it were simply drawn from a table of numbers in the computer, as usual. So I did one tracking session with d(t) generated by the opposing control system. I also saved this d(t) in memory. Then I did a second run using the d(t) from memory as the disturbance—the *same* sequence of numbers that had been the disturbance during the first run. Performance (measured as RMS error or stability or whatever) was *always* poorer with the replayed (or non-actively generated) disturbance. This was a *very* surprising finding; it was dubbed the "Marken effect"—which made my kids very proud.

Bill Powers discovered the explanation of the Marken effect. It turns out that it requires no changes in the PCT model, just the recognition that there are transport lags in control systems (we rarely build transport lags into our simulations, but we should). The "actively generated" disturbance (from the conflicting control system) acts a bit like a spring, allowing dynamic stability and, thus, better control. Once d(t) is generated and replayed, there is no possibility of moment-to-moment adaptation to the subject's dynamics by the opposing disturbance. Bill (and I) confirmed that a control system with a transport lag (I forget the value—I think 100 msec) exhibits the Marken effect, just as subjects do.

Bill suggested (and I confirmed) that you might be able to get improved control in a tracking task if you add the output of a conflicting control system to the "inanimate" disturbance in a tracking task. The gain of the conflicting control system must be low, of course, and the optimal value of the gain produces output that contributes about five percent of the total variance of the effective disturbance to the controlled variable. That is, if q = h + d (where q is the cursor, h is subject output, and d is disturbance), then in the "improved control" situation, d = de + dc, where de is the regular environmental disturbance and do is the added effect of the active output of a conflicting control system. Adding do to de *improves* control if do contributes only about five percent of the variance to the variance of d.

So conflict can help people control—but the gain of the conflicting system must be very, very low. If the conflicting control system were a person, he / she would be very unhappy, because he / she would *always be losing*—he / she would not have any control of the variable he / she would be trying to control.

So I heartily agree with Bill (again) that it's probably best not to harp too much on the presumed value of conflict; there is *far* too much interpersonal conflict already, and .the kind of conflict that seems to be of any value (like the kind in the Marken effect) requires that the gain of one system be so low that people would never want to be that system themselves; weak artificial control systems would. Be best in that role.

Martin Taylor: Bill, I guess I went in an orthogonal direction from what you had intended with your "Coach" story because it triggered things I had wanted to get onto for some time. But your response is also orthogonal to what I had in mind.

You have talked about reorganization as being a consequence of continued, sufficiently bad "intrinsic error." As I understand "intrinsic error," that would make reorganization a whole-system thing. But we had got so far last month as to agree that it had to be modular, and I was working on the presumption that reorganization within a module (a fuzzy module, not one with clear boundaries) would be occasioned by the continued sufficiently bad failure of the module to satisfy its various references. Under those conditions, we don't get a jittery and unpredictable system that could do anything without warning, at least

unless the modules concerned are quite high-level. Most of the hierarchy will still be quite stable.

I totally agree about the problems of social competition. We have far too much of it, and it is an article of faith for many in North America that competition is good. And I do believe that some level of competition is good. Without it, we would have a super-stable non-evolving society such as perhaps might have been in Europe before the Black Death, or in Egypt under the middle Pharaohs, or in China for millennia under the stifling civil service aristocracy. Such a society is not robust against new challenges and does not react quickly to disturbances, any more than does an undisturbed control system-I note your comment about high-gain control in a tense situation.

Rick Marken: A private post from Martin Taylor made me realize that you folks out there are too smart to let me get away with a mistake I made in my description of the Marken effect. I said that the conflicting low-gain controller was trying to keep the cursor in the renter of the screen. This was not correct (although it was correct for the "improved control" situation, where the output of a conflicting controller is added to an environmental disturbance). What I really did was have a subject try to keep the cursor on target (near the middle of the screen) while the conflicting controller tried to move the cursor back and forth randomly. I made the reference input to the conflicting controller a smoothed, time-varying random variable-just like the one that we ordinarily use for the disturbance itself.

In his private post Martin said that he didn't understand why the conflicting controller created a disturbance. Hopefully, my explanation about the "varying reference" in the conflicting controller explains this. Martin, you are right-if the conflicting controller had a fixed reference, then it would not be contributing a varying disturbance for the subject to counteract. By varying the reference of the conflicting control system, the system varies its output to try to get the cursor to match the reference—and since it has low gain, it does a poor job of it. But its varying output provides a nice disturbance to the efforts of the human controller. When this disturbance is replayed, the subject's control is poorer than it was when the *same* disturbance was actively generated, often by a factor of two or more.

When the output of a conflicting controller was added to the disturbance in a tracking task (d = de + dc, where de is the regular disturbance and do is the disturbing output of the conflicting controller), then I had the reference of the conflicting controller fixed. I just realized that this means that the variance of d will be slightly less than the variance of de alone, so any improvement in control using d rather than de could be attributed to the reduced variance of the disturbance. I'll have to do some more research to show that the improvement is due to the addition of de. I think it is a result of adding dc, not just the result of lower variance of d. I think so because if you make the gain of the conflicting controller too high, then there is the expected degradation of performance that comes from being in conflict with another control system—and when the gain of the conflicting control system is high, that control system is acting to reduce the variance of d considerably. So there is a performance decrement, even though variance of d is reduced. But I should do some more research on this. I should be able to do the necessary studies this weekend. I'll let you know how it comes out if you are interested. This is why I need graduate students, darn it.

Bill Powers: Martin, I think we agree that some low background level of reorganization is a good idea. To that, I could add that at the higher levels, where we are on the leading edge of evolutionary development, reorganization might be one of the main ways of groping for control. When reorganization shuts down at the highest level, creative life is finished. I suppose I harp too much on conflict (for reasons with which you evidently agree). We shouldn't forget that control can fail for other reasons, such as confusion, or lack of skill or knowledge. Simply developing the hierarchy is a massive job of reorganization.

Rick Marken: There is no escaping the fact that when the big guy created life, he placed us squarely in the middle of a frustrating paradox—we live by controlling, but we cannot control what is living. Because we are control systems, we cannot be controlled; and because we are control systems, we cannot help trying to control.

PCT is a tough sell because people want to understand things so that they can control better. It is difficult to convince people that things will go better (with other control systems) if they don't control (or, at least, control with a bit less skill). Still, while PCT is a hard sell, I am now convinced that it is very important to, if not sell it, at least make it available to those who might profit most from understanding it; i.e., everybody. Some people will resist these ideas-and even become rude and unpleasant in their efforts to remove the disturbance. But I think it is our responsibility to at least put these ideas out in front of purple, in as clear and convincing a way as possible, without compromising in order to "sell" it. Just give it a chance and understand that nasty replies or reviews are not personal attacks, but the understandable efforts of other control systems to protect principles and system concepts that they consider important.

More than ever in my lifetime, it scams that the world is bound and determined to solve its problems by controlling people. It seems even more insidious now, because this strategy is less obvious than it once

was—when we had clear-cut dictators like Hitler and Stalin using this strategy to the chagrin of most civilized people. Now, our enlightened society thinks its problems come from the fact that we have let people get out of control. So the proposed solutions are more laws, more police, more jails, more regulations, more death penalties, and stricter moral codes—control, control, control. The idea that it might be this orientation toward control that is causing the problem does not even seem to occur to most people. I hear very little serious talk about programs that would "empower people," helping to *give* control: education, work training, child care, cooperative work programs, community centers, insured medical care, etc., etc.

The only objections I hear to solutions that involve controlling others come when controls are suggested for limiting competition; the goal seems to be to have control over other people, unless this control limits conflict. This is a "kinder, gentler" society?

I think it's worth it to try to help people understand their own nature as control systems. If people don't want to understand it, then, fine, we are no worse off than before. But I think that the potential benefits of understanding PCT outweigh the potential unpleasantness associated with trying to teach it.

Joel Judd: There is still a feeling that religion tells people what to do; there is a lot of prescriptivism to the Bible and other scripture. People in general, not just PCTers, often have an aversion to being told what to *do,* even when it might save a life or prevent injury, for example. However, I think that there can be some divergence at lower levels, but convergence at higher levels. In religion, this would relate to getting to "Heaven." But let's use a more mundane example. I'm getting a degree in Education. So is the guy down the hall. But his five years have been spent learning about and practicing counseling psychology, while mine have been spent studying neuropsychology and teaching English as a Second Language. Yet we are both getting Education degrees. We both had to enroll, pass preliminary and final exams, submit a dissertation, pay the fees, etc. Yet none would say that we *did* the same things. There are requirements that *everyone* must fulfill, yet much leeway in how they are fulfilled.

Returning to religion, I once read a comparison made by a church leader between the seemingly "rigid" requirements of religion (Christian, in this case) and natural phenomena. People balk at the idea of "requirements" to get to a higher place (or perhaps they balk at the idea that a *man* purports to know what these are-that's another problem). Anyway, he said that we shouldn't be surprised that a God would place requirements on us, as we can see limits placed on things all around us. For example, water boils at 212°F (assuming we're not trying to make cocoa on Everest, of course). We can heat the water any way we want, as long as the water's temperature reaches 212°F. Dancing around the pot and chanting doesn't work.

My basic question is-and this sounds familiar-why not high-level convergence, and low-level divergence (ignoring for the moment the "who's going to decide which high-level values" problem)?

Rick Marken: What I was trying to say in my previous post is that there is no way to *make* society better—that is, to control it—other than by recognizing that society is made out of individual control systems that work best (and, I believe, work together best-this is my guess) when they are all able to control what they need to control; that is, when the set of 250,000,000 simultaneous equations (for the U.S.) can be solved for all of the unknowns (each equation's controlled variables) simultaneously. The PCT orientation is to help people *control-and not* judge whether or not you think it is something they should control. (Of course, you can't help making that judgement if their efforts to control interfere with your efforts to control; when that happens you are probably running into the degrees-of-freedom problem—not enough resources available to allow everyone to control. The PCT solution to the degrees-of-freedom problem is not very original: population stabilization and non-piggy resource usage.)

Greg Williams: Joel Judd says: "However, I think that there can be some divergence at lower levels, but convergence at higher levels. In religion, this would relate to getting to 'Heaven."' But, speaking purely empirically rather than normatively, it appears that there are wide divergences in notions of "Heavens" and about whether getting to one of them is desirable.

Bill Powers: The religious thing seems to be coming up again, with the usual sniping between the True Believers and the Unbelievers. It's obvious that the Unbelievers are not suddenly going to be converted to Control Theory for Christ, and that the True Believers are not going to switch from *being* believers to *studying* believers. I don't think that railing against a belief is going to advance PCT much, nor is blindly defending any particular belief going to win the day. Perhaps what we might more profitably do is examine belief as a phenomenon.

Belief is a phenomenon worth studying, quite aside from what is believed. What is most interesting is not just a single belief-there will be a sunrise tomorrow—but a *system* of belief. A single belief is usually defended for rather simple reasons: it's hard to find an alternative. But a system of beliefs is an elaborate thing that has the power to take over the mind and shape every aspect of experience to fit it—perceptions,

goals, and actions.

In Perceptual Control Theory (PCT), and even more in Hierarchical Perceptual Control Theory (HPCT), we attempt to build up a concept of how individual human systems work. In trying to learn and improve this theoretical system, we have all come up against our own beliefs; those who have spent years in conventional disciplines have often found their private confrontations of the new with the old unsettling, painful, and even costly.

It seems that simply growing up in a normal educational system, devoting oneself to study, learning what others have found, and meeting the demands of one's mentors is enough to allow systems of belief—or of unbelief—to get a grip that is hard to loosen. Consider the biologist's resistance to the concept of inner purpose. When children who are to become biologists do things on purpose, they take their own intentions, hopes, wishes, and goals for granted: the main problem is how to satisfy them. But put them through the series of educational courses that produces professional biologists, and they come out of it knowing in their hearts that organisms are just biochemical mechanisms with no purposes at all but survival to the age of reproduction. And not only do they "know" this, they *believe* it. To say they believe it means that they now consider their beliefs to be self-evident aspects of the world—not beliefs, but facts. They consider it their duty to inform the world of this truth, to reinterpret the descriptions offered by the misinformed so they properly acknowledge the purposelessness of life, and to deal with other people and more particularly animals as if they had no inner goals of their own. And, of course, they conscientiously interpret their own experiences so they fit the belief that purpose is an outmoded illusion—in their speech, as least, if not in their actions.

This phenomenon of belief isn't confined to biology. People arrive at firmly fixed belief systems about electron flow, quarks, continental drift, natural selection, grammar, etiquette, construction practices, and proper forms of music, art, poetry, and dancing. If you challenge their beliefs, they will defend them. In most cases having to do with less material beliefs, the ultimate defense is "I was raised to think that..." —and of course that is true, although it doesn't make the belief true.

Repudiating or even examining beliefs or unbeliefs is as much a social as a personal problem. To examine a belief or unbelief closely is already to devalue it slightly. To doubt it is to doubt all the circumstances that led one to adopt it in the first place. It is to question people whom one has admired, respected, submitted to, and loved. In effect, it is to see the truthtellers of one's formative years as liars, although of course they were telling what they believed to be the truth.

To question beliefs or unbeliefs is also to question the reasons for which one adopted, or once-and-for-all rejected, a belief. A belief in the ability of one person to control another is not just an article of faith adopted because of love for the teacher, or rejected because the teacher was unpleasant. Believing in the ability to control others suggests all kinds of interesting possibilities if one sees the chance of becoming one of the controllers, and all kinds of horrifying possibilities if it looks as though one will be among the controlled. Beliefs are adopted or denied in part because of what they imply about one's ability to achieve other goals. They are, or at least certain details of them are, expedient in furthering one's own interests.

And finally (although not exhaustively), belief systems are intertwined with one's self-esteem. A scientist who believes in science above all doesn't hold this as an abstract belief. Along with it goes the consciousness that *I am a scientist*, that science is the best of all possible approaches to life, and being a scientist is the best of all possible ways to be. And, of course, those who reject science and choose some other belief system feel that they are among those living some other best-of-all-possible lives, while scientists are either neutral or the worst of all possible people.

The most serious conflicts that take place between people, and the most difficult to resolve, are those that originate at the highest levels of organization. It is not systematic belief per se, nor systematic unbelief, that produces the conflict, but the inability to step back and re-examine a belief when it is confronted by a contrary one. If the Jews and the Moslems come into conflict over their divine destinies, the productive thing for the Jew to do would be to say, "Wait a minute—my beliefs say that this land is historically mine, and you seem to believe it isn't, or that it's yours just as much as mine. How strange—these beliefs can't both be true. What's going on here?" Of course, that isn't what happens, because to most people a fundamental system of belief is to be defended, not examined. The defense, however, guarantees conflict to the limits of brutality.

At the level of systematic belief, both principles and reasoning become subservient to preservation of the belief system. If you look at the arguments against purposiveness in behavior that were advanced—and thought rather clever—in the early parts of this century, you find elementary logical errors and straw-man arguments that wouldn't convince a schoolchild if the subject were something else. You find abandonment of principles of scientific detachment and objective argument in favor of emotional attacks and innuendo. The belief system justifies these alternative uses of principle and reason, because, above all, the belief has to remain true. When you are defending something that is above logic and principle, logic and principle must be bent to the higher purpose.

I count belief and unbelief together as system concepts. There is nothing inherently wrong with either—if there were, we wouldn't have evolved the capacity to form beliefs or unbeliefs. What goes wrong at this level of organization is loss of the ability to alter the organization of one's belief systems to achieve harmony among all the different belief systems necessary to a complete life—different belief systems inside oneself, and different belief systems among different people. I have not identified yet a higher level of organization than system concepts, but this might be entirely due to the fact that the currently highest level of consciousness is never itself an object of awareness; one must occupy a higher viewpoint to see that level as a level, an object of awareness, and a subject for potential modification. Even to speak of belief systems as belief systems, rather than as truths, implies, intellectually, that one is looking from a higher-level viewpoint. But there reason speaks; if there is no still higher level to which one can retreat, as there evidently isn't for me, the viewpoint can only be experienced as a ghostly sense of something just outside the range of peripheral vision that eludes the attempt to see it directly.

As I believe on all the evidence that I am not unique, I can only recommend that others who want to see belief systems as objects of study try to see them that way, thus occupying, if not being able to describe, this viewpoint from which one sees belief systems without identifying with them. To see them this way is not to accept or reject them, or to make them seem less than what they are. It is only to see them far what they are.

Bruce Nevin: Job's paradox, as paraphrased in Archibald MacLeish's verse play *J.B.:* "If God is God, he is not good; / If God is good, he is not God." Any superior intelligence, be it God or visitors in UFOs, cannot control humans and humanity, but can only influence. Bang! Right away, there goes occasion for fear of God or of any truly superior intelligence.

The principal method of influence is by suggestion. An important form of suggestion, whether explicitly in hypnosis or otherwise, is by nonverbal example. A possible PCT-paraphrase of the famous prayer taken up by the 12-step groups: Let me have the reference perceptions for controlling what I can control, for not trying to control what I can at most influence, and for discerning the difference.

Another very important form of influence is by presuppositions riding stowaway on agreements reached by more overt means, such as use of language. Sales techniques depend on this. So does socially institutionalized prejudice. So do most social conventions; only a small, visible minority of social conventions are normally available for conscious attention, those shibboleths that are overtly enforced.

Important among these conventions are those out of which we weave the fabric of personality and self-image. I described some time back research done in which a few practiced speakers recorded the same text repeatedly, varying parameters of delivery and voice quality such as nasality, pitch variation, speed, orotundity, etc. Subjects evaluated these on graded scales for polar adjectives like fat / thin, honest / dishonest, and intelligent / unintelligent. All subjects perceived them as different people, and there was near unanimity in the judgements of the personality attributes of those imagined people.

It is my belief that humans unconsciously control for such variables in constructing a self-image for presentation to others. Certainly this must be so for choice of linguistic dialect; it appears to be so for a very great deal more. We unconsciously control our behavioral outputs in ways that are consistently interpreted by others. Some of this is social convention; some of it is probably biologically innate (smiling when pleased, as a family of gestures encompassing a range of such details). The forced “toothpaste” smile of a model in some ads reads false and might register pain and anger. The Madison Avenue appeal might then actually be to misogyny—whether the ad people know it or not.

We drop bait in the water and keep a watchful (but not consciously acknowledged) eye out for rubbles. We do this by deliberate ambiguity. There is a socially sanctioned interpretation of the interaction that is admitted to awareness. The other levels of meaning are available for awareness, but we choose to ignore them. I propose that this is the real function of patterns such as those Eric Berne and his students describe (games people play, games alcoholics play, scripts people live, etc.). They’re not just to reduce anxiety by structuring time, as Berne suggests. They’re auditions, means of trying one another out for roles in unresolved psychodramas. They’re also opportunities for influence, because they are marvelously suited for re-framing at various levels. I suspect that any competent and experienced therapist does just this at least sometimes.

When I was at Penn in the ’60s, I heard a story about someone coming across the Walt Whitman Bridge and paying for the car behind. Maybe the car following him got sidetracked to a different lane and the driver paying extra didn’t notice, or maybe something more was intended. The toll-taker might have started it out of pocket on a whim. In any case, the next driver who came up holding out his fare, on being told, “The guy in front paid for you,” said, “What the heck, I’ll pay for the guy behind me.” This reportedly went on for several hours before someone put his money back in his pocket. This could be an urban legend, though I have not heard it since.

Joel Judd: Bill says: “It’s obvious that the unbelievers are not suddenly going to be converted to Control Theory for Christ, and that the True

Believers are not going to switch from *being* believers to *studying* believers.” In my case, I agree with the first clause but not the second. I am interested in *believing,* which is why I tried to formulate the convergence / divergence question *about* belief, and why I framed it in terms of my own experience. The go-around last spring made it clear why someone is not going to suddenly switch belief systems, and that’s fine. But that polemic ended with the call to be more “scientific” and find ways of understanding *concepts.* Regardless of the belief, there should be characteristics of control of *concepts* which can be examined just as control of other perceptions.

Bill also says: “...a system of beliefs is an elaborate thing that has the power to take over the mind and shape every aspect of experience to fit it... In trying to learn and improve... we have all come up against our own beliefs...” All the more reason to understand them. Would it be fair to say that even such objective topics as PCT are understood according to these belief systems? Obviously, I’ve been trying to fit it into mine, and judging from the comments, so have / are others.

I think what I would find useful is the development of an efficient way to get a handle on one’s beliefs and their influence on one’s actions —a sort of placement test, if you will. The method of levels has been discussed previously as a way of getting at higher-level goals, at least as far as one can recognize and verbalize them. What about going the “other way”? Supposing that one’s belief about the nature and purpose of language is X, Y, and Z, how does one begin to be aware of how that system influences linguistic principles and the syntactic quirks one controls for, and so on, in a way that can be useful both for potential teachers and learners? Ed Ford has explained several times how he uses a procedure to help people become aware of what they’re controlling for, and how this helps empower them to improve important relationships and resolve conflict. I am thinking that learning some things requires even more detail in terms of the perceptual hierarchy-another language, or adult literacy-there’s a lot involved in making such changes in one’s life. Obviously, such changes can be made. But how might we go about explaining such change in more detail?

Ed Ford: Bill says: “... I can only recommend that others who want to see belief systems as objects of study try to see them that way... To see them this way is not to accept or reject them, or to make them seem less than what they are. It is only to see them for what they are.” The problem for me is that to be properly studied, understood, and fully tested, a belief system has to be checked out through experience. As a Roman Catholic, I have found great internal satisfaction over the years from the standards I’ve set and the decisions I’ve made which have flowed from my religious beliefs. I know others who have left my church and established other beliefs. Some have found satisfaction in their lives, some have not. I think the standards I've set based on my systems concept, the choices I've made which reflected those standards, and, most important of all, the satisfaction that comes from achieving the various things for which I have controlled, are the real tests of systems of belief. It is pretty hard to see this system as an object of belief if, in order to validate it, you have to live it to test it. I think a valid test of any systems concept is this: Does it respect the rights and beliefs of other living control systems? Is it enough to judge a system of beliefs just by how others live it or by what it claims?

Perceptual Control Theory is a good example. Much of the understanding I have of PCT comes from my application of it within my own life, through my dealings with others, and through the success others have made of their lives through their understanding and application of PCT. It has given others a whole new way of looking at their fellow human beings and of respecting the worlds they know little about. I had to immerse myself in the concept and actively live it to really understand it.

Finally, we all have a belief system. It would be hard for my own view or systems of beliefs not to get in the way of those systems I'm trying to study. To me, the real test is when it is given a try, when the rubber hits the road. I guess it's the same as when many of you create a model of what you are thinking. A model for me is when people with whom I work attempt to find satisfaction by using a particular systems concept, *and,* through using this system, they are able to deal successfully with conflict. When people are functioning effectively, then whatever they're using to drive their system should be given respect.

Rick Marken: Rick here, from riot central [Los Angeles, after the Rodney King trial verdict was announced]. I spent the day at home today —work cancelled due to "civic unrest." Boy, are you social psychologists (and sociologists) missing some interesting interactions between living control systems.

I am motivated to begin another thread on social control—but frankly, I'm a bit shaken now. Suffice it to say that I want to talk about the fact that people don't think they are controlling other people when they are. For example, I have heard it said that it is not a control strategy to give people the option of working or living in poverty-it's their choice. I think this is disingenuous, and ultimately hurtful. But it does sound fair and humane—not like control. Just like operant conditioning, really: you can press the bar or starve, it's your choice. We even can be nicer and give you many ways not to starve besides pressing the bar; what could be fairer?

Dag Forssell: Bill says: "The religious thing seems to be coming up again... Perhaps what we might more profitably do is examine belief as a phenomenon." He is suggesting that we go up a level. I agree with Bill that there is *no* difference between *belief* in what we label religious areas and *understanding* in what we label secular areas. I find Bill's post lucid and indisputable—it hooks nicely into my system of understanding, that is.

Bill also says: "... I can only recommend that others who want to see belief systems as objects of study try to see them that way... To see them this way is not to accept or reject them, or to make them seem less than what they are. It is only to see them *for* what they are." Ed then says: "The problem for me is that to be properly studied, understood, and fully tested, a belief system has to be checked out through experience." Ed, as I interpret your comment, you do not mean to object to Bill's statement as such, but to emphasize the practical difficulty of passing judgement on some specific systems concept. You clearly recognize that both PCT and Roman Catholicism are systems concepts. You appear to me to support Bill's post, but you also appear to go beyond it.

You bring up issues of testing and validation of a set of systems concepts. In this, you express a point of view that I think is a good subject for discussion. This systems concept debate will not go away, because it is of great interest to many. We are each attached to our individual set of systems concepts. It illustrates the upper reaches of HPCT, which are of great concern to you and me and any others who try to learn from HPCT how to better teach or lead or counsel people.

Ed, you also say: "I think the standards I've set based on my systems concept, the choices I've made which reflected those standards, and, most important of all, the satisfaction that comes from achieving the various things for which I have controlled are the real tests of systems of belief." In my first reading of this, I understood you to say that systems concepts imply standards, and since the standards work and yield a satisfying life, this validates the systems concepts. However, I believe this last part to be a mistake.

You might not mean the second part the way I interpreted it at first, since you go on to say: "When people are functioning effectively, then what ever they're using to drive their system should be given respect." I think the interpretation that systems concepts validate standards (or "My standards work, therefore my systems concepts must be *true*") is an unexamined assumption behind most of the systems concepts strife we see in the world around us.

I want to focus this post on the standards. Perhaps in that I am "going down a level."

I would argue that the notion of validating or testing systems concepts is a mistake. It is not necessary, as you indicate in your last quote above. I respect you as a thoroughly decent human being. I can never study, understand, and check out your belief system without living your life from its beginning. I do not want to, and it is not necessary.

To think that the standards validate a systems concept implies that those *standards* that do the validation are *unique* to that *systems concept package* (read "religion"). This is the implication I perceive and am debating. Perhaps I am punching a big hole in the air. That's OK, too.

I sincerely believe that if there are five billion people on this earth, there are also five billion systems concepts (of God and everything else). To a PCTer, it is obvious that systems concepts are individually designed by each person.

Just as we PCTers recognize that a diverse set of objects can with some advantage be categorized as "chairs," so a diverse set of umpteen million individual systems concepts with some common, perhaps even superficial, characteristics are called 'Roman Catholicism." Other sets are called "Mormonism," "Islam," "Hinduism;" "Secular Humanism," "Atheism," etc. This is good enough for wars.

It seems impossible to understand another individual's systems concepts in anything more than the most cursory categorization, and then we know that we really don't understand very much.

The point I want to make is that many systems concept packages support the *same* standards. Therefore, it does *not* follow that your systems concept package is validated by the success of your standards. I would be content to say (I think) that your systems concepts are validated by the simple fact that they are yours. Your systems concepts are yours, and that is *enough.*

It *does* make sense to advocate religious freedom and to declare that any religious notion is acceptable, as long as it does not violate important standards that have been agreed upon after more or less public debate over tens of thousands of years (often in the form of wars). If indeed principles / standards / values are what count, and most people on reflection and discussion will arrive at a similar set, it will not be surprising that there is a great uniformity in that area between all religions. In the course of history, many creative thinkers and founders of religions have postulated different systems concept packages on top of them. The *(same)* principles / standards / values used to create a particular systems concept structure logically could be expected to be derived from it.

It is also possible that a principle taught or experienced "on the way up" is remembered and used "on the way down" without being explicitly recognized as part of a system of concepts. We experience a lot as we grow up in our families, which stays with us as principles / values / standards without deliberate connection with, reflection on, or support by our religious beliefs. The idea that systems concepts imply standards does require a deliberate effort to think things through. This should not be taken for granted!

It seems to me that the common inclination (if there is one) to validate your own particular systems concept package by the effectiveness of the (common) standards leads to some very unfortunate side effects.

The idea that the systems concept package is validated to be (rigid, objective) *truth* sets the stage for fruitless discussion, fights, and wars, since anyone who looks can see that the *other guy's* systems concept package is *false.* (Heretic is the word, I guess. Death to heretics!!!)

Religions as systems concept packages typically include whole superstructures of baggage in the form of miracles and explanations which at one time probably were designed to sell the packages to illiterate, ignorant people and keep them in check. Some of this creates unfortunate standards which prevent people from functioning well.

I have my systems concepts which flavor my interpretations. If a God created the Big Bang (today's news), fine with me. I do not recognize a God that can hear me. I think a pastor who tells people from the pulpit that if they can pray together in *His* name to put Jello gelatin "salad" to good use in their bodies—and they *believe* it—is doing these ignorant people a great disservice. Of course, they can pray for healing on Sunday. I have heard enough of this, as our family attended church regularly a few years back. We no longer attend. To me, this is part of the baggage that I personally object to as creating misleading and damaging standards. But then, as Ed says: "It would be hard for my own view or systems of beliefs not to get in the way of those systems I'm trying to study."

These packages may include some principles / values / standards that are not only misleading, but deny people rights we as Westerners take for granted. As Ed puts it: "Does it respect the rights and beliefs of other living control systems?" Consider women's rights under Islam. Since Islam is *true,* validated by the satisfaction of Muslim men, how can you question those things? By going "down a level" and recognizing that the systems concept is nothing more than a construct in your mind. It is not *truth.* There is no *truth* to be had anywhere. It is *all* subjective systems concepts.

I have bared a little of my prejudices here. Everyone has their own. The point is that, as I see it, the debate on creation has *nothing* to do with standards; miracles don't matter. A lot of the things we fight over in religion, between religions, against religion, and for religion do not matter; they are not essential to justify the *principles/values/standards* that *do matter.*

While I looked at my bookshelf of Thomas Jefferson materials, I was reminded of *The Five Thousand Year Leap,* by W. Cleon Skousen. This book by a constitutional lawyer and scholar spells out the 28 *principles* which the American founding fathers considered as they formed our government (a systems concept!!!). It is very clear from this book that the American constitution is based in large measure on the political writings of Marcus Cicero, which were well-known to our founding fathers, *not* on the Judeo-Christian tradition, as we are told often by some religion salesmen.

A nasty thought crosses my mind in regard to some of these salesmen. To paraphrase Hitler's information minister Goebbels: "A point of view repeated often enough becomes the truth." Perhaps Goebbels is another historic figure who clearly anticipated William T. Powers. But Solomon said, "There is nothing new under the sun."

About systems concepts: PCT shows us plainly that all of our behavior is designed to create or (much more often) recreate perceptions we want. From the lowest motor control perceptions to the highest systems concept perceptions. We perceive that which we want to perceive. At the systems concept level, you can rephrase that to say we make come *true* that which we want to be *true.*

Five billion people are controlling to confirm that what they already individually *know* to be *true* continues to be *true.* Progress takes place only when people experience an error signal with regard to a system concept, where it fails to explain or satisfy. Then, a person is open to consider alternative principles which will adjust the existing system of concepts to a new, revised one.

It has been a few centuries since one person claimed to have and have read all books; to know all knowledge. Today, it is impossible to know it all. Ignorance is the rule. The only question is one of degree and area. I am comfortable knowing that I am ignorant in vast areas of knowledge. This recognition makes for a sense of wonder and makes it easier to be open to new information in all areas.

Ed Ford: Dag says: "... many systems concept packages support the *same* standards. Therefore, it does *not* follow that your systems concept package is validated by the success of your standards. I would be content to say (I think) that your systems concepts are validated by the simple fact that they are yours. Your systems concepts are *yours,* and that is *enough."* I think you are looking at this in a linear way. My systems concepts level is my highest level, out of which I create a set of standards, criteria, or principles which form the guidelines for the decisions I make. So far, this is all theoretical. The real test for anything is when the rubber hits the road.

When I teach, I believe all of my students should be treated as fairly as I humanly can. At the same time, I have established a standard within that "fairly" framework that limits the time for individual explanation or debate with one student during classroom time, which, if lengthy,

would deprive other students of needed instruction or role play experience. The decisions I make and the consequent actions I take with individual students are constantly monitored by me as I compare what I want to the variable I'm trying to control, namely the student/teacher interaction variable. So it isn't the standards as such that are or are not successful, but rather the entire behavioral process within my system as it evolves during my class. So it isn't whether the standards in and of themselves (or as they relate to the systems concepts) are successful, for they can't be measured independently of the entire behavioral structure that is the operational living control system. Rather, it is our whole system operating as a continuous process. This involves a whole bunch of things that are all interlaced, interactive, and interrelated, each being a part of the whole process. I might have to adjust my systems concepts (as when I learned PCT), or change a few standards, or alter specific goals or decisions, or change my approach to controlling the variable, perhaps by dealing in a more effective way with the various obvious and sometimes unforeseen disturbances. Systems concepts are validated not because they are mine, but because, over a period of time, I have found satisfaction and fulfillment through controlling and closing perceptual errors using specific systems concepts as reference signals. This is the real test of any systems concepts, I would think. This is where real success is measured. Establishing systems concepts, setting standards, and making decisions is only a part of this process. It also involves being able to control for the right variable at the right time, dealing with both foreseen and unforeseen disturbances, and learning to "listen to and deal with" our reorganization system, while at the same time contending with other conflicting reference signals and principles, both within our own system and in the various systems around us.

You also say: "Religions as systems concept packages typically include a whole superstructures of baggage in the form of miracles and explanations which at one time probably were designed to sell the packages to illiterate, ignorant people and keep them in check. Some of this creates unfortunate standards which prevent people from functioning well." Concerning the use of my own faith as an example, I promise you, I'll not do so again. As a person who, at the tender age of 65, believes in a personal and loving God, in prayer, in miracles (I actually witnessed one), and in spiritual growth, I can assure you my standards have not prevented this illiterate and ignorant person from functioning well. As to keeping me in check, my wife Hester and my children have been trying to for years, but with very little success.

Rick Marken: I think I have been making the mistake of sounding like I believe that people *can* control other people—and shouldn't. What I mean is that people *try* to get other people to act as they (the would-be controllers) want. Of course, the controller is not *really* controlling; but the controller is acting as though he / she can control (and it looks enough like control so that people imagine that it can be done). The fact is, of course, that if you *really* try to control someone (make them do behavior X, no matter what), then you are simply placing yourself in conflict with that other control system. Most of what passes for social control is just social "influence" (manipulating a side-effect of control, for example by disturbing a controlled variable). When the controller becomes implacable (because the controllee fails to continue being influenced), then you get problems.

Bruce Nevin: Rick says: "Of course, the controller is not *really* controlling; but the controller is acting as though he / she can control (and it looks enough like control so that people imagine that it can be done)." This is complicated by the fact that people try to make and maintain social arrangements for cooperative action. This has the effect of people acting as if they were being controlled. The precursors of this are pretty basic in animal behavior, I think. Act in a predictable way around animals, and they get used to you. Act unpredictably, and they go on alert and can get quite upset. Social arrangements for cooperative action require predictable behavioral outputs of the participants, as though the participants were being controlled by one another or by the social arrangement itself.

On another tack, the other day I saw some books by Georges Bataille. In a pair of books with a title something like "the unbearable share," he (says the cover blurb) develops the notion that the converse of utility is at the root of social arrangements and culture. First, the paradox: on a utilitarian theory, in which X is justified by its utility for the sake of Y, the whole must be ultimately based on something that is useless. This neatly parallels the lack of reference perceptions (I almost typed "reverence perceptions") above the highest observable level of the perceptual hierarchy. He builds up his theory on the notion that useless things like potlatch, conspicuous consumption, and eroticism are more fundamental to culture and history than control of the means of production, etc.

Greg Williams: As I've said before, PCT isn't a single-edged cutting (through the crap') implement. Mapping out others' control structures using PCT techniques (particularly the Test for Controlled Variables) can be preliminary to manipulating the activities of those structures, as well as to "empowering" them. I hope some other PCTers will admit how effective PCT tools could be in the hands of the "predict-and-control" folks, and quit burying their heads in the comforting sands of

verbalisms like "there are no social control systems." (True, but not very comforting when you realize that Big Brother might prefer to let you go on controlling as you wish, but with *skewed* premises. And how does Big Brother decide on which premises to skew? One efficient way is to learn about parts of your control structure by applying the Test for the Controlled Variable.)

The last time I brought up this issue, Bill Powers suggested that such manipulations in the light of (partial) knowledge of what others tend to control for are doomed to be "short-term" only. But Bill said that "short-term" could mean many years. Ulp!

Rick says: "Me fact is, of course, that if you *really* try to control someone (make them do behavior X, no matter what), then you are simply placing yourself in conflict with that other control system." What I am getting at above is that it is possible (to a degree, and certainly within limits), by using the Test, to reduce conflict with another's control structure while manipulating that structure to want what the controller wants and not what the structure would have (hypothetically) wanted in the absence of the controller's manipulations.

Dag Forssell: Ed, you say: "So it isn't the standards as such... for they can't be measured independently of the entire behavioral structure that is the operational living control system." I agree with you. The standards certainly fit in a framework. They are at the 10th of 11 levels in the HPCT structure, as presently defined.

You say: "Rather, it is our whole system operating as a continuous process. This involves a whole bunch of things that are all interlaced, interactive, and interrelated, each being a part of the whole process." No argument here.

You say: "I might have to adjust my systems concepts (as when I learned PCT), or change a few standards, or alter specific goals or decisions, or change my approach to controlling the variable, perhaps by dealing in a more effective way with the various obvious and sometimes unforeseen disturbances." You are describing the HPCT hierarchy and noting that you carefully consider how it all ties together in order to function well. We are in perfect agreement. The careful consideration is an important point.

You say: "Establishing systems concepts, setting standards, and making decisions is only a part of this process." Yes, only the three highest levels.

And you say: "It also involves being able to control for the right variable at the right time, dealing with both foreseen and unforeseen disturbances, and learning to listen to and deal with' our reorganization system, while at the same time contending with other conflicting reference signals and principles, both within our own system and in the various systems around us." As I read you, you are describing the essence of the "behavior of perception" in a dynamic environment, and noting how reorganization fits into the picture when normal operation is not enough to control the error signals.

As near as I can tell, we are in perfect agreement—in part because I have learned from you. Since each of us has our individual construct of HPCT in our own head, we will never have quite the same concept of HPCT or anything else, or the same way to explain or think of it.

I still feel that it is more fruitful for human interaction to focus on principles / values / standards as a *subject of discussion*, and I would like to point out that unless I have misunderstood you, this is precisely what you do when you ask a counseling client, "What are your priorities?" You don't ask, "What is your understanding about life?" or "What are your beliefs?" or "What is the meaning of it all?" The systems concepts are a very large network of understandings. It is unmanageable to question systems concepts directly in therapy. You would get trapped in a labyrinth and never get out. The standards are both more relevant and more accessible.

I grant you that the person might look into his / her systems concepts to answer the question, "What are your priorities?" But perhaps not; the problem might be that the person has not spent much time to integrate a set of systems concepts, depending instead on fragments of principles / values / standards as taught by and absorbed without deliberation from parents, peers, siblings, teachers, etc. Perhaps your question about standards requires the patient to think about the systems concepts deliberately for the first time in a long time and create some. You teach PCT, which provides a good framework for that process, without being (or appearing to be) offensive to whatever pre-existing systems concepts the person might have.

I read into your post another aspect of your therapy: If the person does not know how to solve a problem (program and sequence levels), even with newly considered (reasonable) standards, the system does not work. It is an integrated whole! Then you have to teach how to solve a problem, starting with one that has a chance of success. Eventually (hopefully) the person learns to function better at all of the (integrated) levels.

Many things have come together to shape my systems concepts.

Ever since Luther gave Gustavus Vasa an excuse to grab all of the Catholic gold in Sweden in 1523, Sweden has had a Lutheran state church. From first grade through junior college in the public school system, I had two lessons a week in "Christianity." In the later years, it amounted to "comparative religion." I was introduced to the basic tenets of all the major world religions. This is conducive to thinking of them all as systems concepts—with malice toward none, with charity

for all—and seeing that one of the major purposes of religious teaching down through the ages has been character education: teaching standards, so that people can function well.

In science and engineering, I have understood since high school biology that the *only* way into the human nervous system is through the nerve endings of the various senses. With this perspective, it is clear to me that it's all perception. I did not need Bill Powers to make that a part of my systems concept. PCT suggests one way to imagine the specifics. Whether it is done on one level in one massive neural network or in 99 levels of hierarchy is immaterial to the basic premise: it's all perception.

In the past year, I have read Thomas S. Kuhn's *The Structure of Scientific Revolutions.* It could just as well also be titled *The Structure of Religious Revolutions.* Kuhn makes it abundantly clear that to understand a system of concepts, you must internalize that particular set of concepts. When you have done that, you will see and understand the world through the eyes of those rules, that "paradigm." If it works for you (at least reasonably well), you make it your *truth* and defend it against all comers.

I have a tape by Marilyn VanDerbur which includes a quote from Joan of Arc. Joan has been offered her life and liberty if she will only take back what she has said; deny what she believes in. Says Joan: "The world can use these words, I know this now. Every man gives his life for what he believes. Every woman gives her life for what she believes. Sometimes people believe in little or nothing, and yet they give their lives to that little or nothing. One life is all we have. And we live it as we believe in living it, and then it is gone. But to live without belief and purpose, to me is more tragic than dying. Even more tragic than dying young."

A few years ago, I read Bertrand Russell's A *History of Western Philosophy* and enjoyed the TV series "The Day the Universe Changed," by James Burke. It is clear to me that *many* systems concepts, explaining the world around us, have been used, lived by, and died for down through the ages in Western Civilization. It is also my perception that many of these still are in use, handed down through different religions, cultures, and oral traditions.

I think that to say, as Ed says, "... we discover... the true outcome of being human in a real universe" is another way of saying that our systems concepts (the creation of realities in the right way) are validated by our ability to function well, which is Ed's point in the first place. If we develop a reasonable set of systems concepts and reasonable standards to go with them, then we will function well in the Boss Reality. To wit: If we have adopted standards for a good diet, we have a better chance of maintaining health than if we depend on Jello and prayer.

Let me mention that I am in no way against prayer. I think, rather, that it is the atheist who refuses to engage in introspection and quiet dialogue with himself / herself, as an anti-religious posture, who loses out on that deal. It is the ignorant dependence on Jello that saddens me, and that is a question both of systems concepts in regard to the understanding of nutrition, and of standards in applying the knowledge.

To say that it is *all* perception seems ridiculous to a person eating breakfast. The world is real enough. Indeed, in millions of experiments since we came of age, we have hardly ever failed to touch an object as intended. The reality is palpable. We grab the cup. The coffee is hot.

A few months ago, Gary Cziko posted an experiment, which I have adopted. (Thanks, Gary!) Ask a person, while seated, to cover one eye and push on the other while gazing across the room. All that happens is that the image moves sideways a little. Then ask the person to stand up on one leg. Challenge the person to remain standing. Repeat the experiment. The point is that our senses are so well calibrated that we fail to notice the difference between the actual and the perception of the actual. But the moment we push on the eye—a sensing instrument—the difference becomes obvious.

At a higher level, I have adopted Ed Ford's discussion of the concept of wife. It is quite fun to tell the story of how Christine and I met in a whirlwind of fun, and, after three weeks, I said: "I love you, do you want to be my wife?" She answered: "I love you, I want to be your wife!" My concept of wife was based on seeing my mother slave away in the kitchen, taking care of six kids. Christine's concept of wife was based on seeing her mother shopping in London once a week, with the household handled by six servants. How long was the marriage likely to last?

So far, we have shown that it's all perception at the lowest levels and at the intermediate levels in the hierarchy. Why should anything be more than perception at the highest level? How could you *possibly* build certain truth on a foundation of uncertain perceptions? No, it's all perception; all the way up.

Since the dawn of human experience, people have no doubt tried to make sense of their experience, to suggest systems concepts which can explain. In the realm of human behavior, among those many concepts are (1) that God makes us do what we do; (2) that our Soul makes us do what we do; (3) that impressions of our environment (accumulated and presently impinging on us) make us do what we do. Then there is (4) HPCT, which says that our purposes in comparison with the environment make us do what we do.

Through loud shouting matches on this net, we know quite clearly that HPCT is not compatible with the environmental behaviorism S-R. S-R is purely a machine concept, directly at odds with the notion of

God or Soul. We do not mention that PCT is also not compatible with the idea of any one particular concept of God or Soul as *objective truth* or Boss *Reality*. It's all perception. However, the concept of God or Soul is quite compatible, I think, and perfectly respectable as an individual person's personal systems concept. All that is required for compatibility in every direction is for an individual to recognize and acknowledge that it's all perception.

As organisms, we learn *only* from experience. Our *only* source of information is the intensity (or energy) signals we experience from our nerve endings. With a head start in the structure our genes have instructed for the biological machinery, we construct an understanding of those experiences in our nervous systems. One advantage we as humans have is the spoken and written language. By way of language, we can share the experiences of others and thus accelerate and multiply our individual experiences. Still, this all has to enter through nerve endings.

Ed, I do not mean to pick on you, but by way of your own example: If Ed has read or been told about a miracle, that is a perceived experience. If Ed has personally witnessed a miracle, this is a perceived experience just the same, subject to Ed's perceptual capability and interpretation. Ed does the perceiving through nerve endings and construction of an understanding in Ed's mind in either case, and both are subjectively real to Ed. No one has any business questioning Ed's reality. It is his. As I said in my previous post, I think it is obvious that there have to be five billion individually constructed systems concepts among five billion people.

PCT requires a lot of reorganization and takes a long time to grasp, because it provides a complete perspective which is not really compatible with many of the systems concepts people have used with various success since time began. Things will be much easier 50 years from now, when PCT is taught in elementary school and all the way up. (Unless fundamentalists catch on and object, of course). When that happens, the world will be a better place for our grandchildren. That is worth living, working, and dying for!

In the meantime, I believe that discussion of particular systems concept elements as *truths* is pointless, but that it can be very fruitful to focus on the standards which have a much greater universality and direct impact on the functioning of an individual control system. (They are, after all, one level closer to where the rubber hits the road).

It's all perception!

Rick Marken: I think this discussion could be cleared up for me a bit if Ed or Dag could tell me what the word "standards" means in this context. I think of standards as specifications—so, for me, "standard" is a synonym for reference level for perceptual variables. Ed and Dag seem to be using the word "standard" to refer to a type of perceptual variable (like a principle or system concept). What do you mean by "standards"?

Dag says: "... one of the major purposes of religious teaching down through the ages has been character education: teaching standards, so that people can function well." Are they teaching you how to perceive "standards"—like "thou shalt not X, and thou shalt not Y are examples of standards, kiddies"? Or are they teaching you where to set your references (standards) for certain variables that the church assumes you can already perceive—like "I know you can perceive many different gods, but you better set your reference for perceiving YHWH as numero uno, or fry, bubby"? I think that Dag meant that religions teach standards in the second sense: "set your reference for these perceptions here, or else." Is this correct?

I would suggest that religions do try to teach people where to set their references for certain perceptions. I think this is not a good way to help people function well—in fact, it is just about the worst thing you can do to many people. It would only help if (1) everybody perceives the words in the same way, (2) everybody uses words exactly the same way in describing those perceptions (so that everybody knows an "abomination" when they see it), and (3) everybody lives in a world that produces exactly the same disturbances for everyone, so that certain reference settings are always the right way to correct for disturbances of higher-level perceptual variables. I think it's safe to say that the probability of any one of these conditions being met is close to zero. The probability of all three being met is thus zero times zero times zero equals zero. This is my estimate of the probability of religion being a reasonable solution to the real-life problems of any individual living control system.

But it's worth a try.

Dag goes on to say: "If we develop a reasonable set of systems concepts and reasonable standards to go with them, then we will function well in the Boss Reality." I would rather say that what we develop to function well in Boss Reality are *control systems*. We develop means of perceiving and of influencing the perceptions such that they are controllable. Unquestionably, there are ways to perceive and act that make control impossible; the solutions we develop for controlling our perceptions are constrained by boss reality. I must, for example, learn to exert forces on the steering wheel that bear a particular relationship to my perception of the angle between my car and the road's center line in order to control that angle. But there is not a "right" way to set the references for that force, since the amount I exert depends on continuously varying disturbances acting on the car.

Dag's statement implies that there are "reasonable" ways to set references (if standards mean references) for perceptual variables. If this is what he means, then I must disagree. Reference settings depend on the goals of higher-level systems *and* on disturbances to the variables controlled by those systems—there is no one "reasonable" setting for references at any level of the hierarchy. There can't be—and imagining that this is so can lead to internal conflict, interpersonal conflict, or self-destruction (I think that's what happened to Joan of Arc—lack of willingness to adjust a reference to control another variable. She imagined that there are absolute references. That's her choice, of course, but as for me, give me liberty or let me outta here).

Dag says: "To wit: If we have adopted standards for a good diet, we have a better chance of maintaining health than if we depend on Jello and prayer." Maybe, maybe not. "It's all perception," and all you do is control perception. If you can control the perceptions you need to control with Jello consumption, then it's fine—chance has nothing to do with control. You either control the perception or you don't-and you reorganize. If prayer works to control the perception you are trying to control, then great—if not, not. No chance involved.

If one eats vegetables to increase their chance of living longer, then I think they are controlling an imagined perception. If one eats vegetables to feel better—and they feel better when they do eat veggies and worse when they don't—then they are controlling some perception or other, and that's fine. Some people eat steaks and wash it down with a whiskey to successfully control the same perception. There are many ways that can (and, because of Boss Reality, sometimes must) be used to control the same perception. I think it's just important to be sure one is controlling perceptions and not just imaginations, because the perceptions could be getting out of control behind one's back.

Yes, Dag, it is all perception. But we have to live with the fact that we want some of those perceptions to be a certain way. We want to control them. And to do that, we have to be able to develop systems that will take into account the constraints of our own nature (the fact that we are controlling many perceptions at the same time) and the constraints of Boss Reality. And a control system only works (controls) if it can vary its output to compensate for disturbances to the controlled perceptual variables. These outputs are often references for lower-level perceptual variables; so the last thing you need in an effective control system is a "pegged" output-one that does not vary. A control system that believes that there is only one reasonable output (reference) value for another control system is, to my way of thinking, nothing but a big problem-whether that control system exists within our own hierarchy or in someone else's hierarchy. Control systems that think that there is just one "right" reference value for a perceptual variable are the control systems that really need to learn PCT!

Dag Forssell: Rick asks: 'What do you mean by 'standards'?" I have tried very hard to make the connection to the principle level. The word "value" is in there too. To Ed, it is at the level of understanding and belief, if I understand him correctly. I think it belongs at the principles / standards level, if it belongs at all. It signifies a judgement as to what is important among the things you understand / believe.

Rick says: "I would suggest that religions do try to teach people where to set their references for certain perceptions. I think this is not a good way to help people to function well—in fact, it is just about the worst thing you can do to many people." As I have talked to you and have read your posts for a long time, I have gotten the impression that you think that "people will do what they will do," regardless, and that, as a fellow human being, you have no business influencing them. You did admit to me once that you just might have influenced your kids along the way. How? Did you, perchance, teach them where they might profitably set their reference perceptions, so that they might function better? Dr. Spock told a generation of parents to leave their kids alone, and let them do whatever they pleased. I suppose those kids earn the highest incomes and have the happiest marriages now. Surely they must function well,, since no one tried to "control" them when they were little. I think that just about the only thing that separates humans from other animals is the ability to suggest reference perceptions which the young can adopt because they choose to.

Greg Williams recently commented on the tendency of PCTers to bury their heads in the sand when it comes to "social control." Influence is a form of social control, for sure. Why be afraid of it? Influence is for real, and it is important. The world is not populated only by well behaved, adult, PCT-academics, who object to being "controlled" by others. To pretend that positive influence through teaching "standards" or "principles" is (1) impossible or (2) bad is a cop out. Parenting, management, teaching, leadership, and counseling are about that.

When you make an earnest effort to help people manage themselves better, because they have hired you for that or because they are your kids, you are faced with the real question of how to influence them positively and effectively. You cannot duck and talk theory alone, but it sure helps to have a good one. You cannot afford the time and confusion of dealing with everything all at once. You have to figure out a good place to start. I know of no better application of PCT and set of suggestions on that subject than in Ed Ford's book, *Freedom from Stress.* Have you read it? Ed shows how to question people so that they will reason with themselves, but he also suggests and teaches. Ed is a master of positive influence.

I have wanted to try on the net my thought that the level of principles is key, and the suggestion that there are some well-defined, universally acceptable reference perceptions or "standards" that have worked well for a lot of people over time. Character education is, I think, a very useful form of "social control" that is vitally important, no matter where it comes from. Of course, it is also important that this same character education is not misused, as historically has been the case in many times, religions, places, and cultures. Greg might call it a double-edged sword. But the total absence is a disaster, for sure. That is why I think it makes a good subject for discussion.

Joel Judd: Rick says, regarding religions telling people where to set their references for certain perceptions: "I think this is not a good way to help people to function well... It would only help if (1) everybody perceives the words in the same way, (2) everybody uses words exactly the same way... (3) everybody lives in a world that produces exactly the same disturbances..." I think it functions very well if a religion has a "do all you can for others but be responsible for yourself" ethic at its roots. In this way, you try to point out to someone what kinds of things have worked for you and others, but you do not force them to act in your image.

The thing about principles that I think gets confusing sometimes is the distinction between how we label the principle and what we *do* that we interpret as reflecting it. I don't think there is anything wrong with telling someone, "Don't lie! It's bad." But there is always someone (invariably someone older and "wiser") who asks, "But what if the Gestapo is knocking at my door asking if there are any Jews in my basement?" Here we have a particular experience—not a common one, by the way—where I have no problem telling the officer 'No." But that doesn't make lying good! And my three-year-old certainly doesn't understand when I tell her, "Look, mistruths are generally not good, and telling them will contribute to a type of character most people don't appreciate, so you should always tell the truth, except when your mother asks you if her green hair is beautiful, or your friend asks you if her dying Mom is going to get better, or the Gestapo knocks at your door..." Besides proposing standards for people to follow, religions also usually provide guidelines against which to check your personal interpretation of the standards.

We *all* teach standards to others, whether we consciously recognize it or not. Being grown-ups and knowing so much about everything, it's sometimes tempting to let the benevolence in us make us reluctant to teach the things that really do bring happiness to people's lives, in the name of not infringing on their "rights" or "freedoms" or "autonomy" or whatever.

Bill Powers: Rick says: "Control systems that think that there is just one 'right' reference value for a perceptual variable are the control systems that really need to learn PCT!" And Dag says: "Greg Williams recently commented on the tendency of PCTers to bury their heads in the sand when it comes to 'social control.' Influence is a form of social control, for sure. Why be afraid of it? Influence is for real, and it is important." Influence is not control unless you (1) insist that your influence have a particular effect on the other person, and if it does not, (2) apply whatever means is necessary to make sure it does have that effect.

Influences should be thought of as disturbances. That is, you can perform an act that by itself would alter the other's perceptual world if it were the only influence. But you realize that you can't determine the *outcome* of that act in the other person. We tend to use the same word, influence, for the act we perform and for its effect, just as we do with the word "disturbance." Setting an example is an influence, in that it presents a situation to another's perceptions. But it doesn't necessarily *have* an influence, in the sense of altering the other's way of behaving. Even if it does alter the other's behavior, that change might be simply a way of counteracting the influence, and will disappear as soon as the influencing act ceases. Of course, what we hope for is a more or less permanent change in the other's way of doing things—but that result comes from the other person's way of dealing with and understanding the influence. We can't make it happen from outside that person. So it's important in using the term influence to distinguish between the act we perform that's intended to have an effect, and the effect that actually results or doesn't result.

Parents influence their children by (for example) advice, commands, example, demonstration, and story-telling. Children generally being eager for new experiences and not being very sure of themselves in situations beyond their capacities, they normally latch on to these influences and adopt from them whatever fits their growing organizations.

If, however, they don't adopt some of them, or reject some of them, the parents might then resort to punishments and withholdings as ways of trying to make their influences have the desired (by the parent) effects. Then we get all the ills that result from concerted attempts to control other control systems. The children learn, in protecting themselves from direct external control, how to satisfy the parents' reference levels and thus remove the pressure. They learn to lie, dissemble, conceal, misrepresent, pretend, and otherwise give the impression of compliance while internally isolating themselves from their parents.

They become, in short, alienated from the adult world. Of course, a lot of the children simply buy into the system and save themselves all that trouble.

Social influence is not social control. But it's hard to learn how to influence (act on) other people while accepting completely that they will not be influenced (be changed) if that is their choice. When we exert influences on other people, hoping for some change in their behavior that's to our own liking, it often happens that there's no visible result. What do we do then? If we just try harder, we're falling into controlling another person, or trying to. If we give up, we haven't achieved what we want. It's hard to find the middle ground, where we give it a good try but on detecting serious resistance give a higher priority to respecting the other's will as much as our own.

I'm not saying that one should never try to control other people. If a kid runs out in the middle of the street, we whisk the kid to safety by whatever physical means is required. If we're being mugged, we do whatever is required to protect ourselves or those we care about. Not everyone goes around respecting other people's wills. We can't just pretend that everyone in the world subscribes to the same system concepts. Well, we can, but it's not always wise.

What really counts is our understanding of human nature. If we understand that all people are basically as autonomous as we are, then we wouldn't want to encourage a system in which autonomy is ignored or overridden by force as a matter of policy (the present most popular system). With that understanding, we try to deal with others in a way that encourages them to understand things the same way, and to realize that if they want to continue being autonomous, they have to support a system in which autonomy is generally accepted as a fact. Once you see that basic concept, you understand the problem we're trying to solve in our social interactions. There's always a conflict between what we want other people to do and what they want to do. If we begin by respecting the will of others as much as our own, there are certain kinds of resolutions of the conflict we will avoid using as long as possible. We will spend more time trying to find clever ways to satisfy all of us, and less time plotting how to get our own way regardless. It seems to me that that would be a pretty nice world to live in. I'd like to persuade others that it's worth a try. But of course I can't control them into doing so.

Rick Marken: Dag says: "As I have talked to you and read your posts for a long time, I get the impression that you think that 'people will do what they will do' regardless, and that as a fellow human being you have no business influencing them. You did admit to me once that you just might have influenced your kids along the way. How? Did you perchance, teach them where they might profitably set their reference perceptions, so that they might function better?" I want people to be able to control their own perceptual variables as skillfully as they can without interfering with the ability of other people to control their own variables. To the extent that one can help another person (or child) to control more skillfully, that is great. I don't care what people want to control (as long as in doesn't interfere with what I want to control), I only want them to be able to control it. My motto is: a control system in control is a control system that's a pleasure to live with (unless that control system is trying to control you or the things you want to control, relative to different reference levels).

If we take the hierarchical control model seriously, then I don't see how anyone could possibly know how to tell another control system "where they might profitably set their own references for their perceptions." This doesn't mean that I would not suggest a reference (or force the results that would be produced by having that reference) under certain circumstances. The classic example is "wouldn't you tell your kid not to run out in the street?" You bet your sweet bippy I would (and did), and I would physically haul them back out of the street if they were in it—different references or not. But I certainly wouldn't say that what I am doing is suggesting a profitable reference setting for the kid. I'm suggesting ways that the kid might want to control the perception of getting hit by a car. If I could (which the model says I can't anyway) get the kid to have as a reference "don't run into the street," then what happens when the street is empty and is the only refuge from a group of bike riders barreling down the sidewalk? Sometimes the "running into the street" reference is good to have set at "yes."

And that's my point; the HPCT model says that there just cannot be a right or profitable setting for a reference signal; reference signals must be able to vary due to disturbances from the environment or the actions of other control systems. What is a good reference setting in situation A will be a bad one in situation B.

What is important in the HPCT model is not the particular setting of any reference (even the higher-order ones that you call standards), but the fact that references vary as part of a closed loop that produces *control* of *perceptions. Of* course, the HPCT model could be wrong, and there could be a *right* set of references at some or all levels. But I'd need some evidence before I reject a model that seems to work so well at making detailed, quantitative predictions of behavior. As it sits, the HPCT model rules out the possibility of "correct" references-except where "correct" is defined as that setting of the references that leads to actions which, when combined with prevailing disturbances, produces *control.* And this just means that "correct" is going to change all the time (sometimes you *must* run into the street, sometimes you *must not* —if you want to *control* other variables).

Dag also says: "I think that just about the only thing that separates humans from other animals is the ability to suggest reference perceptions which the young can adopt because they choose to." What humans (and other animals) teach their offspring is how to *control*, not what level to keep a particular perception, no matter what. I suppose part of teaching control is suggesting references for a perception ("try to bring your arm farther back on the backswing"), but I think the learner is just exploring the ability to vary that perception as a means of controlling others. What a good teacher teaches is *how* to control—not *what* to control, no matter what.

And Dag says: "Influence is for real, and it is important. The world is not populated only by well-behaved, adult, PCT-academics, who object to being 'controlled' by others. To pretend that positive influence through teaching 'standards' or 'principles' is (1) impossible or (2) bad is a cop out. Parenting, management, teaching, leadership, and counseling are about that." I'm not saying that teaching "standards" is impossible. I'm saying that if people actually adopted fixed standards, they'd be dead in the water; they would not be able to control higher-level variables.

I can't help thinking that I am "well-behaved" because I have pretty good control of the perceptual variables that I need and want to control. I have to believe that *most* of those who misbehave are doing so not because they haven't learned about "right" reference levels for certain perceptions, but because they can't control much at all—let alone what you might suggest as the profitable things to control. Society has been trying to make people "well-behaved" by teaching them values, good "standards," etc. for *centuries.* But there are still plenty of misbehaving people, especially in places where people have the least ability to control their own perceptions (due to lack of education, money, skills, resources, etc., etc.). (I have noticed very little serious misbehavior in Beverly Hills; and I hear that Dag's town, Valencia, is a very safe place. Is this because the people in these people have learned the correct "standards"? I think it's because they have excellent control of what they need to control-and not such hot standards sometimes). I believe it is lack of *control* that you perceive as misbehavior, not lack of "good standards;" and I find it mean-spirited and coercive when people claim that the solution to "misbehavior" is getting people to learn better values (standards). How condescending; where is Charles Dickens when we need him? I think "teaching standards" is just that of tire religion again; it's certainly not HPCT.

As for influencing my own kids—of course I want to influence them. But what I really want is for them to be skilled controllers, able to deal with a world filled with unpredictable disturbances that do not allow for inflexibility and simple solutions. I want them educated and loved (so that they can learn with poise). I don't know how to teach control; but I know it's not by teaching the "right" references. One thing that is involved is a respect for the fact that the kid is the only possible system that can know when its references are set properly; it's when there is minimal error at all levels of the hierarchy. My kids are (so far) splendid control systems; that's all I ask.

Dag says: "Character education is, I think, a very useful form of 'social control' that is vitally important, no matter where it comes from." And I say: forget character education. To the extent that you are in the position to do so, teach people how to control (and keep a good supply of degrees of freedom available for allowing that control—i.e., prevent overpopulation) and you will end up with a bunch of very nice characters.

It's all control.

Ed Ford: Concerning standards: Each of us has perceptions of how things ought to be, found at the systems concept level. In order to control for these perceptions, we each set for ourselves certain principles or standards that reflect those concepts and will become the basis upon which we make our decisions.

It seems to me that we set standards for ourselves and in cooperation with others (such as in a home, a community, or within an organization) all of the time. We also insist that others live by the community standards we set, or else we try to control those who refuse to voluntarily follow our standards. We teach our children cooperative standard-setting with others as the most sensible way to live in harmony. That is why we have communities filled with all kinds of standards, called laws. We as communities and families have certain values, and we set certain standards within the home or community that reflect those shared and agreed-upon values. We also teach our children how to set their own standards, and, just as important, we ask them to explore the down-the-road consequences of the standards they've set. This I've called teaching responsibility. I define responsibility as the willingness and ability of people to follow standards and rules and, ultimately, to set their own, without infringing on the rights of others.

I've done group therapy with juveniles quite extensively in various types of settings (mostly schools and correctional facilities). The juveniles are there because of their refusal to obey the standards of the community in which they live and also for having violated the rights of others. I think the purpose of group therapy is to teach those skills which lead to satisfying lives, including learning the skills for making and maintaining satisfying relationships, as well as the skills for becoming a self-sufficient, self-supporting, responsible human being.

The real issue for me is what is the most efficient and effective way to teach these skills at home or in various social settings. Since I am really only an influence, I have found that the best way for me to work with others is to first find out if the living control systems with whom I am dealing (1) want to deal with me, and (2) have reference signals having to do with improving their lives.

I have a close friend who has a 17-1 / 2-year-old son who lives at home, doesn't work, gets up at 6 p.m. and goes out until 5 a.m., is involved in stealing, etc. My friend is running a real conflict, where one reference signal is pulling toward throwing the (sometimes violent) kid out of the house, but there are two other reference signals: one that wants to avoid physical and possibly violent confrontation, and another signal demonstrating a great deal of love for this child. He also has several other reference signals, which include harmony with his wife (who is all for throwing the kid out of the house) and another that involves maintaining the standards for harmonious living within this home. That child is not willing to change his life style and is unwilling to deal with either his father or myself. It's very nice to control for what you want, but when it runs against the prevailing standards of where you live or work, then you have to live with the eventual consequences of your decisions.

The way I teach others how to use their control systems is through asking questions. I find little difficulty with the people I work with (including corrections) to get people to move up one or more levels. In fact, as soon as I get my clients to list their areas of importance to me (systems concepts level) and have them prioritize those areas in terms of importance, that's when I find therapy really gets going. As they begin to identify the areas of conflict (conflict between two reference signals at the highest order), evaluate what area they want to work on or where they would most likely succeed, and set the kind of standards for the area where if they were able to accomplish their goals it would them bring satisfaction, that's when they seem to find some relief. When they say they feel better, what there really saying is "I think I can now figure out a way for making things better in my life." Obviously, the real proof is when they begin to succeed.

With regard to standards within social organizations: Most, if not all, organizations and communities have set standards, and you have to be willing to live with those standards, or you leave (or don't join). There are many belief systems that say that if you want the perceived benefits of being a part of us, and you want identify yourself with us, then you've got to accept our beliefs and abide by the standards we've set that reflect those beliefs. And I think that's fair. I've joined several organizations whose standards were such that I left the organizations. Others I have remained with, the Control Systems Group being one.

Sometimes we set standards that just involve ourselves. I'm a very strict vegetarian, yet I've never tried to impose these standards on my wife or children. I have certain standards in other areas of my life. I figure the way I live my life is the best influence I can provide to others. One thing for sure, I've learned not to try to impose my beliefs and subsequent standards on others.

It's all perception, but we're responsible for our own.

Dag Forssell: Halfway up the HPCT perception ladder, a person might agree with Ed Ford that a husband and wife will have *different* concepts of "wife," but human nature being what it is, there will be an intuitive tendency to say: One to one against does not count. I *know* my *Boss Reality.* I know what a wife is (sort of), and I will continue to use that information. After all, mine is the only percept I have access to.

When we come to a miracle, the natural tendency, given a *long* history of perceiving in a certain way, is to say: It might look like a billion to one against to me, but I *know* my *Boss (Reality),* and nobody is going to take it away from me.

It goes against all intuition and apparent dependability of our basic senses to say "it's all perception," but it is the only conclusion I can defend, given my perceptual constructs. I think that when a person recognizes and acknowledges this, the person is more free to reorganize (without internal conflict), to respect his / her fellow humans (complete with individual perceptual constructs), and to promote a better social order with more degrees of freedom for all.

There is a *Boss Reality.* Our challenge is to perceive it as effectively and accurately as we can, while recognizing that this is *all* we *can* do. The *Boss Reality* does place constraints on our degrees of freedom. I perceive that HPCT provides an effective (and as accurate as can be had at present) perception of the *Boss Reality* of our minds.

The question of how to control well with maximum degrees of freedom for all will quickly demand attention to issues of influence, "social control" if you will, the principles or "standards" we live by, and the quality of information in all corners of our Hierarchical Perceptual Control Systems.

Rick says: "Reference settings depend on the goals of higher-level systems *and* on disturbances to the variables controlled by those systems-there is no one 'reasonable' setting for references at *any* level of the hierarchy. There can't be-and imagining that this is so can lead to internal conflict, interpersonal conflict, or self-destruction..." That there is no one "reasonable" setting for anything at any level might be quite valid. Is that a reason to never discuss any suggested settings at the principle level? I believe a lot of people abstain because of the uncertainties. Your reading of Joan of Arc differs from mine. I read her as

saying that she was willing to die for *her* references, not that they were *absolute.* Self-destructive? Sure! But in the long run, we all live and ultimately die for what we believe in-hopefully, of old age. What do we believe in? HPCT!

Rick also says: "Control systems that think that there is just one 'right' reference value for a perceptual variable are the control systems that really need to learn PCT!" The name is Dag Forssell. I am indeed trying to learn!

Bill says: "Influence is not control unless you (1) insist that your influence have a particular effect on the other person, and if it does not, (2) apply whatever means is necessary to make sure it does have that effect." I appreciate this help at sorting out definitions. I find it a difficult subject. But important to any practical use of PCT.

Violence and social control are bad. Influence might be OK, but we don't much like it either, because it smacks of control. The lines of demarkation get fuzzy. If a wife is unable to influence her husband, eventually she can exercise social control in the form of divorce. If an employer is unable to influence an employee to be productive in the line of business the company is in, then he will have to influence the employee to seek other employment. Some will call it (mistakenly?) social control or even violence. Personally, I have been laid off and have quit. It is a natural consequence of my own and my employers' requirements for degrees of freedom. But there surely is a lot of unnecessary waste, violence, and social control in business. Neither employers nor employees are effective in their control. You find conflict every place you look.

It seems to me that the absence of appropriate influence leads a person to fail to develop the good information content required for good, effective, satisfying control. I continue to be interested in influence as a constructive activity. It is difficult to deal with.

We must show how to apply HPCT for the satisfaction of all. Information offered must tie into what people already (think they) know. It must offer something of immediate interest, address some dissatisfaction or error signal people have, or it is of no interest.

Rick Marken: I have my references for "standards," just like everyone else. If asked, I would say I like people (including myself) with high levels for what I perceive as honesty, integrity, responsibility, and so on. It's difficult to talk about "standards" without having an idea of what constitute "good ones." So I suggest that we move this discussion from a discussion of standards (just one type of perceptual variable) to the model that supposedly informs our understanding of human nature. 1 might prefer particular system concepts, standards (principles), programs, etc.—i.e., I might have a collection of references which can, over time and variations, be perceived as a particular "ideology."

That's just me—my system that grew up over the last 46 years. I'm not interested in pushing my whole ideology, just one component of that ideology: a model of human nature called HPCT.

So, what I believe in (as far as this audience is concerned) is a spreadsheet! The three-level spreadsheet hierarchy described in one of the papers in my book, *Mind Readings*, captures what I believe is the basic functional organization of a human being. Some things are missing —like the reorganization system. But this model gives a good picture of my image of an organized (grown-up) purposeful adult. One nice thing about this model is that it is *all numbers. The* numbers that are perceptions are functions of other numbers—the functions define what is perceived. Numbers are nice, because people don't care that much about them. The perception numbers at level 2 of the model, for example, could be representations of the degree to which some standard (like "honesty") is being perceived (in the spreadsheet level two perceptions are actually functions of linear combinations of intensity perceptions). The spreadsheet has four control systems at three levels; the reference for the highest-level systems are fixed (they are numerical constants), but they could be changed randomly by a reorganizing system. The model acts to keep all of its perceptions matching all of its references. So the level 3 systems adjust the references to *all* level 2 systems (changing the reference *numbers) so* that the level 3 perceptions are maintained at the reference levels. The spreadsheet does this even when you change the environmental variables (also numbers) on which the controlled perceptions are based-that is, it controls a hierarchy of perceptions in the context of changing environmental disturbances and in the context of the changing control actions of all the individual control systems.

If you give names to the numbers at each level of the hierarchy, then things can get personal. For example, if you think of system 1, level 2 as controlling a perceptual "standard" called honesty (as one means of controlling the higher-order perceptions, which might be called "system concepts"), then you have to say that the system is varying its reference for honesty to control whatever perceptions are being controlled by the higher-order variable. This is why I say I don't think that there can be fixed references for *any* perception-it's not because I'm pushing moral relativism or personal autonomy or libertarianism. The only thing I am pushing is the PCT *model* (and I can envision it best and see it working best in the spreadsheet implementation, because I know what the numbers mean; I know this is not the easiest way for many people to visualize the model, but it does have that one nice feature: it doesn't hit any emotional buttons).

So I suggest that when we discuss these big philosophical issues, we try (to the extent we can) to relate them to what we actually know—the

HCPT model. HPCT is a real, working model, and many aspects of it have been tested and passed with flying colors. There are many aspects of the model that we don't understand (like how it could perceive something like "honesty"), and many things that will surely need to be added or changed as a result of research (Greg's suggestion that higher-order outputs might influence lower-level parameters besides references inputs, for example).

I think if we talk about functional organization more, and specific perceptual variables and their references less, we might get a better idea of what HPCT is about. The words (and the fact that people are themselves control systems with their own references for standards and whatever) can really get in the way. HPCT is HPCT—it's not liberalism, radicalism, libertarianism, judaism, mormonism, monotheism, etc., etc. It's a functional model that explains (purportedly) why people behave according to any of these principles. The model is a bunch of numbers that are functionally related to other numbers. It doesn't say what it is "best" for those numbers to represent.

If there is any "value system" implied by the HPCT model, it is just that the model should *work*—i.e., it should be able to keep *all* of its perceptual numbers equal to all of its reference numbers. Anything that prevents the model from doing that is something that should be "fixed."

When Dag says that there might be "right" levels for certain perceptual variables ("standards"), what I hear is the claim that "I can set one of those level 2 reference numbers to a *constant* in the spreadsheet hierarchy, and everything will work even better—the only thing that I have to do is find the *right* number." Well, I know that that is not true—quantitatively: it is not true of numbers in a control hierarchy. If you believe that those numbers are a representation of perceptual variables and that things like honesty are perceptual variables, then I leave the conclusion to you.

But I am open to any model-based (and research-confirmed) evidence that there are *right* constant values for variables in the HPCT model. I mean, HPCT is my ideology, but it is open to test (that is one *nice* thing about numbers).

Greg Williams: I find much appeal in the recent posts by both Dag Forssell and Rick on standards and PCT. It seems to me that Rick's viewpoint, with PCT ("all perception / all control / all numbers") in the foreground, addresses the issues in a general manner, while Dag addresses some particulars. I can see the validity of both in their special provinces—but I think everyone must beware of being overly provincial.

From an examination of the histories of the diverse ethical systems which have flourished around the world at various times, ethical *contextualism* (rather than either relativism or absolutism) might be the best model. The nasty connotations of ethical relativism are apparent when the issue is framed as: In *this* society, you're trying to tell me that anything goes (makes sense, fits in, works)? Obviously ("obviously!") the way we upstanding citizens do it is what is best!!! The nasty connotations of ethical absolutism are apparent when the issue is framed as: Why don't those people in that *other* society do it the way *we* do? Obviously ("obviously!") our way works, and so it should work for them, too!!!

An ethical contextualism offers a middle road, recognizing that, within a particular context (sometimes quite broadly defined—i.e., we're all human), there are certain standards which *do* "work," but also recognizing that if the context is different, those standards might cease to "work." Most tribes, I have read, refer to themselves as The People, which emphasizes their distinctness from others who *aren't* The People. That insularity, rooted in ongoing personal confrontations with a particular context, makes great sense up to a point. Then the conquerors come along, of course, and try to impose a new ethics (no more infanticide, etc.) *and a new context.* If the new ethics precedes the new context (and probably even if it doesn't), there is a great likelihood of pain.

Personally, I would like to see more recognition that individuals' contexts are much more variable *within* our own society than many people like to admit, and so there are grounds for ethical pluralism (e.g., in attitudes toward abortion as influenced by economic status). Yet I understand that there is a perceived need to restrain such pluralism in hopes of keeping "us" (e.g., U.S.) "together" in the face of "challenges" (mainly "foreign competition," it seems, these days) from "outside."

So I can see the cases for local (sometimes *very* local—and possibly *very* ephemeral, too) "absolute" standards *and* for the contextuality of *all* of those standards, seen more globally.

It's all contextual.

Rick Marken: Dag Forssell says: "That there is no one 'reasonable' setting for anything at any level might be quite valid. Is that a reason to never discuss any suggested settings at the principle level?" The change in the height of a column of water depends on the volume, not the mass, of an object that is placed in the water. Is that a reason never to discuss ways to bring the water level to a particular height by suggesting settings for the mass of the object to be added? I think the answer to your question is another question: What do you consider to be a waste of time?

Ed Ford says: "Each of us has perceptions of how things ought to be,

found at the systems concept level." These are called references; they define what we ought to be perceiving. These exist at *all* levels in the model, not just at the systems concept level. We have references for how much pressure to feel on our fingers, and how much like a fist our hand configuration should be in, and how rapidly our hand configuration should be changing.

Ed also says: "It seems to me we set standards for ourselves and in cooperation with others (such as in a home, a community, or within an organization) all of the time." This is the crux, I think. We care about "standards" because they often determine lower-level actions that might influence the variables controlled by other people. I think Greg picked up on this in his suggestion that "ethical *contextualism* (rather than relativism or absolutism) might be the best model .... An ethical contextualism offers a middle road, recognizing that within a particular context (sometimes quite broadly defined—i.e., we're all human), there are certain standards which *do* 'work; but also recognizing that if the context is different, those standards might cease to 'work."' Yes; and the important context is other control systems. My spreadsheet model has to be expanded to two (or more) hierarchical systems simultaneously working in the same environment of numbers (degrees of freedom). I think you would find that these models would quickly run into conflict if their *higher-level* (level 3) systems were controlling for the same variables relative to different reference levels. There would always be less conflict at the lower levels because the references for those levels can be changed by the higher-level systems that see that there are lower-order errors.

Actually, I think I will do this modeling effort; but my intuition is that the only way to solve the problem of multiple interacting control systems, operating in the same environment, is to align the references for the highest-order systems that are controlling the same perceptual variables. I wonder if the solution would be found automatically (through reorganization), or whether there needs to be a system that actually perceives that there is conflict and looks for a cooperative solution. I think the former might work.

So I think it's possible that alignment of higher-order references might be a natural consequence of being reorganizable hierarchical control systems. Of course, the values at which these systems get aligned are not necessarily determined—just as long as they are aligned. I think this is why we see such remarkable differences in cultures. There are remarkable differences between cultures in terms of system concepts like marriage (polyandry, monogamy, polygamy, etc.), and they all work; apparently because everyone buys into that reference. Of course, once pressures lead individuals to shift references (our society seems tacitly moving from monogamy to serial monogamy—largely as a result of an unpredictable disturbance; people are living longer), conflicts between control systems increase, as would be expected until the group is able to "realign."

It's highly unlikely that any society will align on a system concept that demands really "bad" standards like murder. There are standards that are self-correcting (the people who aligned on the system concept that demands murder would be quickly eliminated from the pool of control systems). Note, by the way, that most societies have aligned on system concepts that make it perfectly OK to murder the members of other societies. But that's getting into more substance than I think is appropriate. Back to models.

Greg says: "It's all contextual." OK, I'll buy it. How about another: It's all interacting control systems.

Bill Powers: Rick pointed out that "desirable standards" are not reference levels, but variables. It's easy to show that they are variables just by finding words to indicate other states than the states one automatically assumes for them (the states one likes best):

"courage": bravado, foolhardiness
"conviction": stubbornness, prejudice
"generosity": profligacy, gullibility
"kindness": bleedinghearted sentimentality
"helpfulness": nosy do-goodism
"honesty": bluntness, cruel candor
"honor": hubris, egotism, bushido
"justice": revenge, brutality, litigiousness
"tolerance": naivete, permissiveness
"sound use of time and talents": working for someone else
"freedom of choice": abortion as belated contraception
"good citizenship": supporting the war
"the right to be an individual": offending everyone
"the right of equal opportunity": the right to sleep under a bridge

The problem with lists like these is that they define only dimensions of perception (variables), but, by implication, they specify some particular state of the variables that is "best." The right level for one person is too much for a second and not enough for a third. The right level for today and this person is the wrong level for tomorrow and someone else.

Even the perceptions that go with the words are different for different people. When a manufacturer supports the "Right to Work" act, a labor union opposes it, because the words mean one thing to the manufacturer and another to the union. When an inhabitant of South Los Angeles asks for the right of equal opportunity to work at rebuilding the wreckage, a white construction worker objects because it will deny him or her equal opportunity to make a buck doing the job at a higher wage. Freedom of choice is an empty promise for a person without the means of implementing any choices; for others, it's an excuse for maintaining segregation and shielding themselves from contamination by the rabble. A "sound" use of time and talent means, to some people, not wasting your time on fripperies like music and art and theorizing, but devoting your efforts to maximizing (somebody else's) profits. To a lot of people, honesty means that it's OK to cheat the IRS or a business rival, but not to cheat me.

The names of standards refer to things that are not words, but are shifty attitudes that vary with circumstances. All that makes sense of any kind of standard set to any momentary level is the system concept under which it is adopted. I thought that Ed Ford's recent discussion of standards hit a lot of nails on the head.

I also thought that Rick's statement hit a nail on the head: you can't set a reference signal to a constant value and expect the higher systems to go on working properly. They work by *varying* lower reference signals, not by picking one setting and sticking to it. This isn't "moral relativism"; it's simply recognition that the system concepts that organize and use principles are more important than any particular principle, or any particular state in which to maintain a given kind of principle. Moral rules followed blindly and implacably can generate the cruelest of all human aberrations.

The only reference signals (and perceptions) that can't be changed freely as required by higher levels are system concepts. And the only reason we can't vary our reference signals and perceptions at that level with complete freedom is that there seems to be no place to stand except another system concept-if there is a higher viewpoint, it's impossible to put into words or systematize. If there's free will, the only place it can work is at the top, because everything else is dependent and interconnected. And even at the top, we're free only to be human.

Rick Marken: A public reply to a personal note from Ed Ford: I know no one whose standards I admire more than yours. I think you and I have similar ideas about what constitutes an admirable individual. I think we run into a problem with these damn words. That's why I like models, I guess. They let us back away and just look and see how they work. I think you are able to do this with real people better than nyone 1 know.

Chuck Tucker: Rick, I just could not resist this: By what standards do you determine that you "like" Ed (or anyone, for that matter)?

I think this is a problem too: There is no "good" reason that anyone should like anyone else, but some of us do.

I think what we need is to submit these notions that we (or you) are putting in this model to a test. Devise an experiment that might be able to determine that a person is using a standard with regard to another within the confines of a particular act (this, by the way, was an assignment to my class; no one was able to do it even after reading HPCT all semester).

I think we (really I) must remember that this model (HPCT or PCT) is one which is developed out of engineering and seems to apply quite well to artificial systems that can be built, and to living control systems *up to level three.* There is some information from a variety of studies and other experiences of we who are using this model that it is useful for explaining or understanding or comprehending behaviors. Don't misunderstand me—I firmly believe that this model will work better than any one that I know of—but I realize that it still needs extensive testing. I have not done the work that is required, and I have not seen it done by anyone else, *but* I believe it can and will be done.

Don't you agree?

Rick Marken: Chuck Tucker says: 'Rick, I just could not resist this: By what standards do you determine that you like' Ed (or anyone, for that matter)?" I don't know. But the fact that we like and dislike anything suggests that control is going on. I think it is very difficult to verbally describe all of the perceptual variables that are involved in 'liking a person;' let alone the reference levels (standards) for those variables. If you mean what principle perceptions do I have relative to Ed that I feel are close to my currently prevailing references for those principles (as I sit here at the keyboard and try to describe them), then I think of things like "family"—I like the principles I perceive as exemplified in Ed's relationship to his family. I like the principles I perceive in Ed's interest in and understanding of PCT. Again, these are just words; you would get a better sense of what I'm talking about if you could have my perceptions and my references for those perceptions. Short of that, you could do an informal "test" to see what level of the "family" principle I like to perceive, e.g., by describing people who exemplify different levels of that principle. I admit, for example, that my "liking" for JFK went way down when I heard that he was regularly unfaithful to his wife. You would have to do a lot of testing to figure out what principles were violated for me-for example, my liking for JFK would go right back up if I found that this behavior was done with his wife's consent. My own impression is that what is violated (for me) by JFK's infidelity is a reference level for a particular perception of "respect for other people," not a reference level for "who to sleep with when you're married." But there was some reference for a "standard" (can't we call it a principle as it was originally called? This use of "standard" really confuses me, because it sounds like a reference level) that was violated.

Chuck also says: "I think this is a problem too: There is no 'good' reason that anyone should like anyone else, but some of us do." I think the only reason we do anything is to keep all of our perceptions matching their references. Whether that is a "good" reason or not, I don't know.

And he says: "I think what we need is to submit these notions that we (or you) are putting in this model to a test." Agreed. But it is very difficult to test for control of these higher-order variables; we can barely describe them.

And: "Devise an experiment that might be able to determine that a person is using a standard with regard to another within the confines of a particular act (this, by the way, was the assignment to my class; no one was able to do it even after reading HPCT all semester)." This is very difficult—especially if by "standard" you mean "principle." But it's pretty easy for many other variables. Try getting really close to a stranger; talking really loud during a conversation; using a lot of profanity (if you can—some of these disturbances are hard to produce because they require the "disturber" to set his / her own references to unacceptable levels for his / her own hierarchy). It's really not hard to see variables being controlled-any time someone acts like something is *wrong*, there is a perception that is deviating from a reference—but it's not always easy to name the variable.

I've found (for now) that the discovery of controlled variables is like a Zen exercise; don't try to name stuff; get those words out of your head for a while. Try to just look at the world as variable perceptions; arrangements of objects, relationships between them, etc. Watch how people seem to like certain states of these variables, rather than others. Note that sometimes they seem content with things, and sometimes they protest and complain; the protesting and complaining and the "fixing" and the doing are all evidences that something is not as it should be for a person. You need sharp clinicians to figure out what those controlled perceptions might be.

One of the problems is that most of what people control is too obvious and too "trivial," so it goes unnoticed. People are not generally running around trying to control the "meaning of life." They are moving things from here to there (moving themselves from here to there); carrying out programs, categorizing (it's an "X"—no, it's a "Y"). If you know someone who makes music (well), you might see if they can imitate the "style" of some well-known artist; that's a pretty complex variable (I do a mean Bob Dylan). Control is all around—maybe the problem is that it's too much around—we take it for granted. Bill Powers once said that feedback is like the air we breathe; I think this is true of control, too. Because it is everywhere, it is invisible, unless you know what to look for (like the answer that's "a blowin' in the wind").

Dag Forssell: Rick, You cannot derive any values at all from the HPCT model, especially when it is viewed as a mathematical spreadsheet. I am not claiming that a certain level of these variables or references are *right*. I have meant to offer the observation that perhaps a lot of people get along quite well in spite of holding religious systems concepts that are totally incompatible, because they tend to set references at the principle (what Ed calls standards) level *similarly* anyway. (I have used the word "reasonable," meaning "well thought out," but never in my mind have suggested *absolute* or *constant;* that is your interpretation and contribution—it does make for feisty argument.) Perhaps that shows that more "down-to-earth" systems concepts / understanding based on experience, instead of "intellectual / religious" constructs, .are what really influence the principles most people go by.

I do understand that there is not just one "right" reference value for a perceptual variable anywhere in the hierarchical structure. I do not understand your emphasis on *variable* to describe the list, as if to disqualify reference. As I understand it, precisely the same perception that we call the reference is what "behaves" to create the specified perception of what we call the variable. The words used to describe the reference and the variable perceptions are identical, since the perceptions are identical. You have to specify that you are referring to one or the other. Neither is fixed, since the reference is set at the moment as part of the entire, interacting hierarchy.

Personally, I believe it comes naturally to want to find some meaning in your own life. I think meaning can be found in secular systems concepts just as well as in religious systems concepts.

Rick says: "... I think I will do this modeling effort; but my intuition is that the only way to solve the problem of multiple interacting control systems, operating in the same environment, is to align the references for the highest-order systems that are controlling the same perceptual variables." Great! Maybe we will be able to illustrate more about how control systems disturb one another. You can get part of the way there with rubber bands, but only on one level, of course. I share your expectation about the requirement. This means that we have to talk until we have the same systems concepts, after all. It will not be enough to say that you subscribe to the same principles. This is what Greg observes, as applied to each tribe or subgroup in its context.

This entire exchange has caused me to reflect on my own assumptions and understandings. My ideas relating to character education go

back to 1980-83. I have not scrutinized these particular systems concepts in the light of PCT until now. I have already reorganized some, but I have not settled down yet. I find merit in Rick's observation that it's all control.

HPCT as a model has much to offer. My interests focus on how to teach and apply it. Since we live in a real world with finite degrees of freedom, and a Boss Reality to study, it becomes important to reflect individually on the specific perceptions you fill your own hierarchy with at all levels, so that you can control well. Numbers are not enough. As a parent, manager, teacher, or counselor, it is my challenge to assist those who want to be assisted to fill themselves with good information. Good information will include an understanding of PCT.

There is no such thing as character. There is only effective control. Internalizing the systems concept of yourself as an autonomous control system and adopting the same systems concept on behalf of others (a value judgement) might lead you to principles of similar appearance as those that have been labeled character. As I said, I have not settled down yet. Additionally, perhaps there is no such thing as violence, coercion, social control, influence, etc., which it is why it becomes difficult to distinguish one from the other. What there is is control by one system which creates disturbances for another control system. If two systems control the same variable, you have to look at the coupling of each to the variable, loose or tight, and the resources (or amplification or force) available to each. In arm wrestling, two control systems control the same variable with tight coupling. The control system with the most force minimizes its error signal. The other system gets a large error signal.

Bill says: "Influences should be thought of as disturbances. That is, you can perform an act that by itself would alter the other's perceptual world if it were the only influence." It makes sense to me to see influences as disturbances. Can you see information as disturbances, also? In one book on listening I read long ago, the author suggested that in active listening, you choose to anticipate what the speaker will say next, see what they do say, and compare the two. When you guessed right, you confirm with satisfaction. When the speaker says something else, you think about it intensely. Either way, you are alert and hear well. Of course, you might control so you hear what you want to hear instead. You put the words into your own context.

With this in mind, I can think of reading a post as a lot of small and some not-so-small disturbances. I have to recognize that I am disturbing just the same when I post. Some of the information flies by with minimal disturbance, some is unsettling.

As a parent, I create disturbances for my child in many ways, which the child has to deal with. Thus the child fills with experiences/ perceptions / understanding throughout the hierarchical structure. If I plan the disturbances well, the child learns to control well. I could say that I deliberately create error signals in my child. This thinking agrees with Chuck Tucker's post.

Bill also says: "Parents influence their children by (for example) advice, commands, example, demonstration, and story-telling." I am now beginning to think of all of these forms of influence as made up of disturbances. Does that make sense?

I am controlling and perceiving as best I can.

Mary Powers: Rick and Chuck: Why do we like people, indeed? The real issue to me is why we dislike people. This winter I read Eduardo Galeano's *Memory of Fire*—a three-volume history of the Americas, mainly South and Central (a must read in this half-millenium year of 1992). Pretty brutal. The point here is that it seems "natural" for humans to dislike, fear, and consider subhuman people who are strangers, or different. Often people will like an individual they "get to know" (share reference levels with) and yet continue to dislike other people of the category (Black, Jewish, whatever). We do need to like and be liked, but what about this other reference level?

Ed Ford: The other night I had my meeting of local control theorists who are trying to implement these ideas in their jobs. One of us, Alan Wright, was recently appointed superintendent of schools for the Arizona Department of Juvenile Corrections. He has been working night and day at the two major lockup institutions where the toughest juveniles in the state are sent, trying to implement a new program using PCT as the basis. In the past, diagnostic teams have decided what juveniles would be doing and establishing their plans. Juveniles were staffed monthly and told what they were doing wrong. Juveniles never sensed any control as to when they'd get out. It was always kind of vague. All they learned from the staff was how to be criticized. Now, things have changed, thanks to Alan. First, there are those who know they are getting out (like at age 18) or are going to be transferred to an adult prison. These could care less and continually cause trouble. They have been separated from the rest of the population and are in highly restricted and supervised units. But for the others, things have changed.

Alan and I have been working on the practical applications of PCT to this kind of setting for several years. At Adobe Mountain (the toughest), the juveniles had taken over the place. Alan really tightened the place down. Then each juvenile was asked as he entered the institution what he wanted. The universal answer was to get out. Alan then would ask them what they had to do to get out. He'd explain to those who

didn't know. To the acting-out juveniles, he'd say, "Is what you're doing getting you what you want, which is to get out?"

Alan has the juveniles working in small teams of 36 with three teachers, each teacher directly responsible for 12 juveniles. The job of each teacher is to help the juveniles work toward getting out, which translates into getting certain tasks done in school and following the standards and rules in the classroom. The old idea of being in so long (like six months) and then being released has been replaced by the requirement to get the signoff (approval) of each of the juvenile's direct supervisors in education at the school, and the line officer and case manager where he lives, and the person in charge of activities (work or recreation). Everyone has to sign off, saying that the juvenile is following the rules and working to his best ability and accomplishing his tasks. Any time the juveniles act out and are sent to lockup or to an intensive treatment unit, that time doesn't count against their credit for getting out of the facility. Time is no longer important-only achieving tasks that reflect increased responsibility will get them out of there. The juvenile is given total control over when he gets out of the facility. He has to accomplish certain goals, but he alone has control over how quickly he can get released. Obviously, the more violent the offender's crime, the more responsibility has to be shown over a greater period of time.

It's amazing how the place has settled down. And it's amazing how quickly acting-out juveniles settle down, once they learn they have control over when they get out. To you freedom-loving control systems on the net, this might not sound like PCT, but within the reality of the juvenile correctional system, asking the juveniles what they want and giving them control as to how long they are in a treatment center they don't like has given them a sense of control over their destiny they've never had in the past. There seems to be less violent and more thoughtful reorganization going on. When they do act out, the supervisors just ask, "Is what you're doing getting you out of here?" or "Do you still want to work at getting out of here?"

Dag Forssell: Ed, what a marvelous post! Most encouraging. You show clearly the power of starting with a focus on what people want, instead of a focus on how people behave. You apply disturbances to encourage reorganization over time. You give the delinquents a measure of control they have been denied before.

While your situation is rather extreme, it is not different in kind from many other interactions between parents and their children, business owners and their employees, counselors and their patients.

Again, your results are exciting! Your years of preparation are paying off in a significant way. Congratulations!

Rick Marken: It looks like my distinguished peer, Danny Quayle, has made standards more appropriate than ever. Now, it seems, the official government position is that those nasty social problems would be solved if everyone would just adopt the right values (standards)—and guess who's standards those are? The Trobriand Islanders' values? The Nepalese values? The Danish values (my personal favorite)? Nope—Republican values. Thanks for clearing that up, Dan. It makes me proud to be a member of the '60s generation. If only he could claim that it was the result of smoking too much dope.

I don't like the word "standard"; it can be a synonym for "reference;" so it is very confusing to me when people talk about the *importance* of "standards." It sounds wrong when I think of standards as references, because *all* reference signals are important in the model. And it sounds wrong, for the same reason, when I think of standards as principles. It is no more important to control principles than it is to control intensities.

The PCT model says that we are controlling many levels of perceptual experience simultaneously. Lower-level perceptions are controlled in order to control higher-order perceptions. The higher-level perceptions are in no sense more or less important than lower-level perceptions; all perceptions must be brought to their reference levels in order for there to be control at all. So it is just as important to be able to control the position of your torso as it is to be able to control your position in a perceived relationship as it is to control the principle that is satisfied by being in that relationship, etc.

Principles often have to do with other people (they involve setting references for relationships between you and other people, for example). I think these perceptions seem special only because most of our control problems involve attempts to control variables that involve other people (as one would expect, since people, being control systems, cannot be controlled, and so there will often be large, chronic errors in these systems). It will be very hard to control relationships, programs, principles, etc. that involve other people. Since control is generally poor for variables involving other people, our attention (consciousness) will tend to be examining the control systems at this level (it is a kind of postulate of PCT that consciousness tends to move to the level where reorganization is required-no tests of this that I know of so far; hence, I am talking through my hat).

I am hypothesizing that consciousness (attention) tends to be directed toward control systems involved in the control of variables which involve other people (due to the chronic error that tends to exist in these systems). Better yet, I think we attend to systems involved in the control of variables which involve *at least* the *relationship* between people—most importantly, relationships between *ourselves* and other people. So my hypothesis is that we attend mostly to systems controlling perceptual variables at the relationship level (level 6 and up). We rarely attend to our control of intensities, sensations, configurations, transitions, events, etc. We do attend to relationships (with the boyfriend / girlfriend), programs (soap opera stories), categories ("he was a ..."), principles ("he done her wrong"), and system concepts ("that was no way for a Christian to behave").

I think it does something of a disservice to the PCT model to try to emphasize the importance of one type of perception relative to another. They are *all* important.

If the feeling is that the higher-level systems are more important because they determine the goals of the lower-level perceptions, then this feeling is incorrect (in terms of the model, anyway). The particular reference level that is selected for a lower-level perception depends on the goals of the higher-level perception *and* on prevailing disturbances which are independent of the goals of the system. So, setting my reference for a principle, like "get control of the renter," will result in very different chess moves (relationships) on different occasions; some of those moves might not actually be "good" in terms of other goals (like winning the game) if I just blindly follow the principle.

But you all know that. We just tend to forget it when we are dealing with really "important" principles (the kind that we have been calling standards): principles like "be *kind* to your neighbor" (even when your neighbor is a Nazi who is trying to *kill* you?). The desire to find the "right" references for our perceptions of principles, etc. (i.e., interpersonal perceptions) is strong; and I think it's because consciousness *does* tend to focus on these levels. Consciousness is involved in learning, and the goal of learning is to try to find the "right" reference settings for perceptions involved in what you are trying to learn to do (to control). If we had more difficulty with the lower rather than the higher levels of perceptual control (so that consciousness was always hanging around those levels), we would probably spend all of our time trying to figure out the right configurations, sensations, transitions, and intensities to experience. Sometimes we do try to figure out the "right" settings for these variables—like when we are learning a sport or a musical instrument. Of course, even in this case there are no right settings; just the right variables to vary (by changing references) in order to control the higher-level variable.

I suggest that unless an individual is conflicted at the principle level, there is no reason to try to direct their consciousness to that level in particular (indeed, if they are conflicted at the principle level then you should try to get them—their consciousness—up to the system concept level). I would say that, as a control theorist, I would try to get a person's consciousness *away* from their principle level if there is *no* conflict there.

Putting consciousness where it does not belong can be quite a problem (at least in theory), because it can start a reorganization that is not necessary. If you doubt this, just try moving your consciousness to the perceptual levels that are ordinarily unconsciously controlled in a well-learned skill; I tried this after I had learned to play a two-part invention by heart. I tried to become conscious of what my fingers were doing (the sequence level and some transition and configuration stuff too). The two-part invention turned quickly into an n-part cacophony,

I think Zen people know the potential problems of consciousness. My suggestion to people who are doing therapy (on themselves or on others) is to lay off the levels that are not conflicted. And don't assume that a level is conflicted just because it seems like it is important; I bet very few people have any real problems at the principle or system concept levels. I bet most people just can't control relationships, programs, sequences, stuff like that. I would not *assume* that the problem is always principles (it might be intensities—maybe the person has a boil, not an "attitude").

Another motto: if it works, don't be conscious of it while it's working.

Bruce Nevin: I guess Rick means that one should lay off the levels that don't provide a higher vantage point on levels that are conflicted.

It seems appropriate to attend to principles (standards) if people are conflicted at the program, sequence, or category level, such that taking a point of view through principle perceptions discloses the terms of the conflict. However, I think there are reasons for attending to perceptions at or below the level of conflict.

I think a common (band-aid) resolution of conflict is to ignore lower-level perceptions that cause error at the level of conflict. One way to ignore a perception might be to substitute a copy of the reference signal by the imagination loop. Another way seems to have the effect of making areas of the body blank, dark, numb, foggy, armored—people use different metaphors. The character of the perceptions that are being blanked out might provide clues about the error being ignored, and thence clues about the conflict being band-aid-resolved.

In vipassana practice, after an initial period of attending only to perceptions of the movement of the breath at the nostrils to develop the ability to focus and maintain attention—the first 3 days of a l0-lay course, the first month of a three-month course, etc.—one begins to move attention systematically through the body, from one end to the other, area by area. It is a very common experience for a given area to seem "dark" or devoid of sensory signals, sometimes for extended periods, yet subsequently a great deal seems to be going on there. You're just sitting still, breathing, and moving your attention from place to place, so there's

no evident physical stimulation. In the interim, however, perhaps some emotion-laden imagery or memory has come up to distract you from attending to physical perceptions in the body. Like starlings, if you don't feed them, they go away. Attending to physical perceptions in this way is a way of not feeding them. Their going away unfed seems to be associated with the "waking up" of areas of the body that had been blanked out. Ignoring perceptions seems to have the cost of turning off sensory inputs. People who do body work (massage, polarity, etc.) are familiar with this.

Some forms of therapy dwell on the content of the emotion-laden imagery and memories. Perhaps this can be useful. I suspect it is useful only when people get in touch with their feelings, not in the sense of their emotional reactiveness, but rather in the sense of awareness of physical perceptions in the body.

Rick Marken: Bruce says: "I guess Rick means that one should lay off the levels that don't provide a higher vantage point on levels that are conflicted." Yes. I also mean that, if there is *no* conflict, don't try to become conscious of the non-conflicted systems. It's OK to go up a level from a non-conflict. It's like the piano example—it's OK to be conscious of the fact that I'm playing a two-part invention; it's just a bad idea to focus on the systems that are successfully producing the perceptions that are accomplishing this higher-level goal.

Bruce says: "It seems appropriate to attend to principles (standards) if people are conflicted at the program, sequence, or category level, such that taking a point of view through principle perceptions discloses the terms of the conflict." Absolutely!! And it is OK to go up a level even when there is no conflict; consciousness is just a problem when it is focused on the systems that are currently successfully achieving a higher-level goal; like when you think about *how you* manage to keep the car on the road *while* you are driving (it's perfectly OK, in terms of control ability, to think about how you drive when you are not currently controlling the car).

Assuming that the HPCT model is right and that we really do control perceptions of principles in order to control system concepts, then I am suggesting that, if you direct someone's consciousness to the principles that they are controlling while they are successfully controlling a system concept (like being a Christian or a Dodger fan) *while* they are controlling that system concept, then their control of that system concept will become less skillful. That's OK if there is a conflict at the principle level that prevents control of the system concept; but it's not such a hot idea otherwise (though I think it can be fun; especially if you don't care for the system concept a person is controlling). I think this is what goes on in skillful political debate; get your opponent to look at the principles that they are controlling; suddenly, their ability to defend their system concept deteriorates; not because they see anything wrong with the principle, but because they just *see* it. (This could be another nefarious application of PCT; if you get really good at directing a person's consciousness to certain levels, you could screw up their performance on some task. For example, when you are about to play a game of tennis with your buddy, you might ask, "Say, are you still turning your wrist on the backswing?" Encourage them to think about this during the game. If they do, you are a sure winner.)

David Goldstein: The only problem I have with your therapy suggestion, Rick, is that it is not always easy or necessary or a good idea to be thinking in terms of levels when you are working with a person.

One of the ways you know you are confronting a conflict is when the person is not able to go up a level. The Method of Levels is taking you no place. As Bill Powers has suggested to me, when doing the Method of Levels, don't be so concerned with the levels outlined in the formal theory. Be sensitive to the background perceptions based on what the person is saying and the person's own reactions to what is being said.

A related point to what is being said has to do with a difference I've noticed between Ed Ford and myself in applying HPCT. Ed starts at the systems level and works downward. I start at a lower level and work upward. One advantage of the bottom-to-top strategy is that it avoids what Rick is talking about, namely, directing a person's awareness to levels which "are not broken."

Rick Marken: David says: "Me only problem I have with your therapy suggestion, Rick, is that it is not always easy or necessary or a good idea to be thinking in terms of levels when you are working with a person." *I agree!!* I only used specific PCT-level words because there was talk about the importance of "standards;" where "standards" were alleged to be principles (in HPCT terms). In real life, I would not even try for a second to relate a person's conflicts to the proposed PCT levels. All I suggest is that conflicts can occur at any perceptual level (in theory and in practice), so there is no reason to single out standards (principles) as an important place to look. In fact, the more I think about it, the more convinced I become that real conflicts have to do with pretty low-level percepts, whatever you want to call them, and that the resolution to most conflicts just involves seeing that things can be done in sequence, or that X does not need to be categorized as a Y, or whatever. I think it is rarely necessary to change principles or system concepts (or any high-level perceptual references) to solve most personal problems. I think this is consistent with the fact that people who hold transparently idiotic system concepts (from my perspective) can still get along just great in the world. One exquisite example of this is the fellow who wrote my two-part inventions; J. S. Bach lived a wonderful life and produced the greatest sequences and configurations of sound ever produced—and he did it all for the God of Martin Luther. Silly system concept; great, non-conflicted control system.

I think that when we look at social behavior, we see all kinds of interesting things happening, but it is very difficult to see the possibly very simple sensory variables that are being controlled. PCT is an attempt to help us see beyond our interpretations of behavior—to what behavior is really about: control of perception. And this requires a special kind of looking (based on hypotheses about what variables might be controlled) and testing (to see if disturbances to the variable are resisted). We might be reading a lot into human behavior that is not relevant to what a person is controlling—for example, we say a person has "bad manners" or "poor standards" when they eat with their mouth open —when, in fact, they might just be controlling the amount of pain they feel because they have a toothache.

It is hard to get past our inclination to see behavior as "output." We assume that what we see is what the person is "doing." PCT suggests that we must by to get over that inclination (if we want to understand behavior) and take seriously the proposition that what we are seeing (when we see people "behave") are the means by which people are keeping their *own* perceptions matching their own references for those perceptions. We, as observers, *cannot* see what another person is perceiving or trying to perceive. We can only try to get an idea of what a person might be trying to perceive by doing the Test for the Controlled Variable.

What seems to an observer as control of a complex principle (at the "wrong" level with respect to the observer) might, in fact, be nothing more than efforts to get from point A to point B in the context of variable disturbances.

I think there has been some confusion about what control theory says about control, and what it says about the processes that might influence control—i.e., consciousness. This confusion becomes particularly acute in the discussion of standards (principles) where there is talk about "setting appropriate standards" and such. When we talk like this, who do we imagine to be "setting the appropriate standards"? The hierarchical control model says it is the "higher-order systems." The references for principles are automatically set by the systems controlling systems concepts. The point is that all this varying of lower-level references to control higher-level perceptions is carried out smoothly and automatically by the control hierarchy. I hate to point this out again, but this process is nicely illustrated by my hierarchical spreadsheet and (in a less abstract manner) by Bill's "Little Man." So ordinarily, there is no external "agent" (other than the control hierarchy itself) that sets references —and varies them—and this varying happens automatically. This means that principles are varied *automatically* to control systems concepts; if the systems concept control systems need to vary the honesty principle reference to preserve the perception of the system concept, they *do* it. That's *you* doing it, but there is no choice going on-no conscious decision to be a little more or less honest in this particular situation. It just happens (just as your muscles tense automatically to control the position of your limb).

When we talk about a person "setting appropriate standards," I think we are talking about a phenomenon that is external to the PCT hierarchy. Subjectively, when I talk to myself and say, "maybe I should try X," it is my consciousness that is doing this. Consciousness is like something that hovers over the hierarchy and tinkers with it occasionally; at least, consciousness is what *can* tinker with the hierarchy.

I think we know a hell of a lot less about how consciousness works (in terms of phenomena and models) then we know about how hierarchical control works. But I do think that consciousness (the feeling of having to choose: "should I do X or Y?") only comes up when there is some degree of internal conflict or lack of "output functions" that can be used to control the required perceptions. When you are in control, you are rarely conscious of it, unless you make some effort to notice how well your hierarchy is working. When there is a failure of control (due to conflict, lack of skill, or insuperable disturbance), then consciousness is there. As I said in an earlier post, moving consciousness to systems that don't need attending to can create more problems than it might solve.

Ed Ford: I see counseling as similar to creative writing. When writing, I watch ideas pop out of my mind, as if I had little to do with creating them. I just think about the area where I'm curious or trying to work out a thought, and out something comes. It just pops up, and there it is. The bottom line is that I take advantage of my reorganization system and let it work for me, like creative people do.

Counseling involves using the reorganization system, the same creative process. I don't start at systems concepts and work down any more than I start at a lower level and work up. I begin my session by talking with my clients (what else is there to do?) about what they want, where they see their problems, a little bit about their lives. I have in mind the major areas of importance in PCT that are applicable, such as priorities, values and beliefs, standards, decisions, various areas of perception, our actions, wants and goals, and other stuff. Then I watch myself take certain directions, primarily areas in the clients' lives where both harmony and conflicts exist.

I tend to just watch where I go, letting ideas come out of my mind, not constricting my mind but letting it creatively seek various paths to take. When I occasionally find myself uncomfortable with where I am or what I'm saying, then another idea pops into my mind, and if it makes sense and is compatible with what I want, I go in that direction. That isn't to say that I don't have an overall structure in the way in which I work, or that I don't think about what I'm saying. I'm thinking all the time, but it's within the creative process. PCT has given me a delightful structure, and I've added my own way of understanding and creative process within the boundaries of PCT.

I basically look for where there might be two incompatible goals, or for goals clients have established but over which they have little or no control accomplishing. I also have them look at how they've structured their worlds, and I get them to evaluate the structure they've created.

Typical areas of conflict are a job demanding enormous time, a spouse and / or children needing time, extended members of a family such as sick or lonely parents, physical activities, or intellectual activities. There are all kinds of areas with interrelated and sometimes highly conflicting standards, decisions to be made, and various systems which have been prioritized. It is impossible for an outsider to know all the various areas of importance, their strengths and priorities at any one time, the varying standards, and how they are all interconnected within the total network of a person with whom you are dealing. That's why I think it's best to teach a person how to work out their own internal conflicts —only they know what is really going on. All I know is my own created perceptions of what I think is going on.

I think a big mistake can be made if a person looks at PCT in terms of an individual area of concern and tries to analyze an area in isolation from other areas. Again, there's so much going on. I can't think of a single area of importance to me that isn't tied into lots of other areas of greater and / or lesser importance. Hester, my children, my various jobs, my health, my faith, my friends, CSGnet, things around the house, all kinds of other things as well. These are all very interrelated areas, all with various priorities, depending on the time constraints and other areas of importance. To look for the single or major reason or cause for what people do within their network of reference levels is rather misleading. There seems to me to be too much interrelatedness within our structure of our values and beliefs, how we've prioritized them at any one time, and all the various standards we've set. Added to this is how all of the above can be conflicting with various disturbances when we are attempting to control in various areas.

The most important thing I've learned from control theory is that I'll never understand another living control system, and they'll never understand me. To quote Clint Eastwood (one of my very favorite actors):

“A man has got to know his limitations!” When living control systems come to me seeking help with various conflicts they’re having, I see my goal as a teacher. My job is not to figure out why they do what they do. Rather, it is to help them build confidence in their ability to deal with their internal worlds by teaching them effective and efficient ways of resolving their conflicts and establishing harmony within their worlds. Control theory has given me more help in this area than anything else I’ve learned.

Rick, I think that directing a person’s awareness to levels (or areas) which aren’t “broken” can be very productive. Obviously, if people are doing well in one or more areas, but their belief-in-self systems aren’t, then having them reflect on what they’re doing well can be most helpful in rebuilding confidence. It is best to build from strength, not weakness. Also, sometimes it is best to build more strength in areas of success before attending to weaker areas. Again, as I was saying above, you just have to fuss around and help the client determine which is the best way to go. It isn’t best to set hard and fast rules where you have so much going on.

David Goldstein: Ed, your principle-level generalizations for therapy are, based on your post: be creative and spontaneous. Be a teacher, teach them about HPCT. Be sensitive to signs of conflict and harmony and focus on these areas: Encourage people to believe that they can solve their problems.

At a more specific level, you say: “To look for the single or major reason or cause for what people do within their network of reference levels is rather misleading.” But if HPCT has anything unique to say to therapists, it is: identify controlled variables by means of the Method of Levels and the Test for the Controlled Variable. In a clinical situation, this is much harder than in an experimental situation. If we give up doing this, I am not sure about how HPCT therapy is really any different from other therapies out there.

I know that when you ask people, “What do you want?” in the exploration phase of your counseling, and when you ask people, “is it working?” in your evaluation phase, you are moving in the direction of finding controlled variables. Maybe I simply go further in this direction through the explicit use of the Method of Levels. Asking people questions like you do certainly disturbs them and invites awareness to what is going on inside them.

Ed Ford: David says: “But if HPCT has anything unique to say to therapists, it is: identify controlled variables by means of the Method of Levels and the Test for the Controlled Variable.” I don’t really think PCT actually says anything in particular. The greater the understanding one

has of how the whole system works, the more creative a therapist can be in coming up with all kinds of ways to teach people how to deal with themselves. Certainly the ideas in PCT provide the creative mind various ways to more efficiently help others produce harmony within their own worlds. One of the keys to helping others is, as you suggest above, the whole concept of controlled variables, and how they can be used within the counseling session. However, I believe there are many intriguing ideas flowing from PCT that therapists can use, the idea of controlled variables being one of the more important.

Rick Marken: Ed Ford says: “... I think that directing a person’s awareness to levels (or areas) which aren’t ‘broken’ can be very productive. Obviously, if people are doing well in one or more areas, but their belief-in-self systems aren’t, then having them reflect on what they’re doing well can be most helpful in rebuilding confidence. It is best to build from strength, not weakness.” I guess I didn’t make myself clear. Based on subjective experience (not the PCT model), it seems to me that skill breaks down somewhat when you direct your attention (consciousness) to the means being used to produce a particular result *while you are doing it*. This is easy to demonstrate; while you are typing, think about how you are doing it; how you are moving your fingers, how you are adjusting and coordinating movements of the fingers, etc. You start making mistakes (more of them, anyway) when your awareness moves to these levels of control; control seems to work better when it occurs unconsciously—zen control. But there is no problem when you imagine typing and become conscious, in *imagination*, of how you do things. In fact, certain kinds of conscious imagining are reputed to improve control when you get down to actually controlling. I remember Dwight Stone imagining, over and over again, the details of a high jump event just before executing it. I guess his hope was that once he’d imagined it enough he could just go and do it (control it) unconsciously. The trick is to be able to change easily from conscious imagining to unconscious doing. Sometimes it works; sometimes it doesn’t. Whether imagining itself can actually make things better, I don’t know. Maybe that’s why we dream—but that is usually unconscious (I think).

I have a feeling that consciously focusing on what one does right (in imagination mode, of course) might make one feel better but does not necessarily help a person in other areas. For example, I can’t see how focusing consciousness on, say, one’s ability to throw a football can help with one’s ability to sell cars. I agree that it might help a person control self-confidence a bit (“yeah, you cant sell a car to your mother, but you cam throw the football pretty well”). But that’s just making things better in imagination mode anyway. When the fellow gets back to the car lot, his confidence goes right back to hell. I think it’s better to just get down to the business of helping a person "move up a level" so he / she can see that he / she is creating his / her own problems. Of course, it might help him / her spend the 50 minutes in a session if he / she feels good about you and himself / herself.

Ed also says, in describing a PCT-inspired program for juvenile offenders: "Me juvenile is given total control over when he gets out of the facility. He has to accomplish certain goals, but he alone has control over how quickly he can get released." Well, it's not total control; a fellow (I presume they are all guys) can't get out one second after he gets in—or two, etc. They can control when they get out, but there is a lower (and, I bet, upper) bound to how long they can stay in, no matter how they act.

He also says: "To you freedom-loving control systems on the net, this might not sound like PCT..." I presume that's me. Actually, I only like freedom for people who do not plan to hurt me, my family, or anyone else I like (i.e., everybody). The program sounds just like PCT to me —the juveniles are controlling for getting out; you are controlling for the behavior of the kids, trying to make it look "under control." A program cannot be PCT or not; PCT is just a model of behavior. The program you describe is neither good nor bad—but given that you like it, it is apparently working for you. The kids seem to function in it just fine, too.

I imagine that the delinquents in this facility have set references for and achieved some results that have hurt other people. I am against people hurting other people, and when people do hurt others intentionally (and PCT tells us how to find out if they are doing this intentionally), then I am for preventing that result by any means possible. I am particularly in favor of this kind of intervention if people are doing it because they have become organized in such a way that this kind of hurting is part of the way they are controlling other variables.

So, if these juveniles are organized so that violence to others is just part of their organization, then I don't care what you do to them; just keep them out of my society. If, however, the behavior that got these kids into the facility is the result of reorganizations (because the kids have not been able to get control of their intrinsic variables) or just irrelevant (unintended) side-effects of control efforts that could be eliminated by education or counseling, then I think there might be other ways of dealing with the situation.

One point that might be worth noting: PCT should at least make one sensitive to the possibility that one's efforts to *help* another control system are really an attempt to perceive that system's behavior to be "as we like it." After all, *we* are control systems too, no?

To be continued

The Control Systems Group is a membership organization which supports the understanding of cybernetic control systems in organisms and their environments: *living control systems*. Academicians, clinicians, and other professionals in several disciplines, including biology, psychology, social work, economics, education, engineering, and philosophy, are members of the Group. Annual meetings have been held since 1985. CSG publications include a newsletter and a series of books, as well as this journal. The CSG Business Office is located at 73 Ridge Pl., CR 510, Durango, CO 81301; the phone number is (303)247-7986.

The CSG logo shows the generic structure of cybernetic control systems. A Comparator (C) computes the difference between a reference signal (represented by the arrow coming from above) and the output signal from Sensory (S) computation. The resulting difference signal is the input to the Gain generator (G). Disturbances (represented by the black box) alter the Gain generator output on the way to Sensory computation, where the negative-feedback loop is closed.

Currently available:

Back cover

Closed Loop

Threads from CSGnet

This reproduction of *Closed Loop* was created by Dag Forssell in 2001. Addresses and phone numbers have not been updated. Most are obsolete.

Posted at www.pctresources.com

Proofread as of March 8, 2001

Winter 1993

Volume 3

Number 1

Front cover

Closed Loop

Threads from CSGnet

Winter 1993 Volume 3 Number 1

Edited by Greg Williams, 460 Black Lick Rd., Gravel Switch, KY 40328

CONTENTS

Members of the Control Systems Group receive *Closed Loop* quarterly. For more information, contact Ed Ford, 10209 F. 56th St., Scottsdale, AZ 85253; phone (602) 991-1860.

CSGnet, the electronic mail network for individuals interested in control theory as applied to living systems, is a lively forum for sharing ideas, asking questions, and learning more about the theory, its implications, and its problems. The "threads" in each *Closed Loop*, stitched together from some of the net's many conversations, exemplify the rich interchanges among netters. Some issues of *Closed Loop* also feature research reports by netters, in hopes of initiating new conversations.

There are no sign-up or connect-time charges for participation on CSGnet. The Internet address is CSGL@VMD.CSO.UIUC.EDU while CSGL@UIUCVMD is the Bitnet address. Messages sent to CSGnet via these addresses are automatically forwarded to over 120 participants on five continents, as well as to hundreds of NetNews (Usenet) sites where CSGnet can be found as the newsgroup bit.fistserv.csg-l. CSGnet also can be accessed via CompuServe, AT&T Mail, MCI Mail, or any other computer communication service with a gateway to Internet or Bitnet. For more information about subscribing to CSGnet, contact Gary Cziko, the network manager, at G-CZIKO@UIUC.EDU, phone him at (217) 333-8527, or send a FAX to (217) 244-7620.

Conflict, Belief, Standards: Part II

Martin Taylor: In the posts I've looked at on "standards" and how they should be interpreted in the PCT world, no one has interpreted standards as I would. In interpersonal communication of any kind, including language, you can best achieve control of your percepts if you have some notion of what the other is likely to do that affects your sensory organs. If you don't want to perceive yourself being hit with a 2 x 4, you don't antagonize a Hell's Angel. You model the partner in some way. It seems to me that standards allow you to pre-empt a possibly painful random reorganization by permitting you to set references that are appropriate if the other behaves in a conventionalized way—according to standards. Likewise, if you behave according to standards, your references will be set so that your observable behavior conforms to the expectations of others—they will know what you are controlling for at the relevant level, and they will be able to interpret low-level acts/behaviors as supporting that control.

If there are any "absolute" standards, they will be those that have allowed the social groups using them to survive and prosper. A standard that allows group members to kill one another for fun is not one that is likely to be found in a long-surviving group. Our standards have been evolving since at least the time humanoids diverged from other primates, and there are clearly some sets of standards that work well together but are different from other sets that also work well together. One standard that worked well when relatively isolated tribes wandered around competing for resources involved wariness and intolerance for people not of one's own group. Killing them meant more for one's own group. Racism comes from this. But recently there has come to be only one communicating group in the world, and this long-useful standard seems to be one that will not allow this single group to survive long if it maintains its currency as a model for how to set a reference level.

Standards for grammatical usage seem to have exactly the same theoretical standing as standards for good social behavior. One sets references for using "correct" grammar because it eases the task of communicating partners who use the same standards. If a subgroup uses different standards, there's no problem except that their communication with the main group becomes less effective. If one person decides on a different set of reference levels, they cause communication problems with all of their partners. There's no moral good or bad about it, only a consideration of efficiency.

Bruce Nevin: The previous discussion of "standards" substituted that term for "principles," as in hierarchical perceptual control theory's (HPCT's) level 10. Standards, meaning "norms" or "conventions," can be on any level. Modeling others to facilitate cooperative action with them involves perceptions on many levels. Martin, the convergence of your discussion with the prior one is perhaps this: that people are aware of norms, conventions, and models of others mostly on the principle level, the level at which they attribute motivations and make moral judgments.

Rick Marken: Martin suggests that "standards" should be viewed as conventions that make it easier to cooperate. I agree that there is much to be gained from conventionalized behavior. This is particularly true in the technological world, where it helps enormously to design systems that have a standard response to actions. Thus, we can be pretty confident that a clockwise turn will result in the screw going in or the power going on or increasing. What we tend to conventionalize is the feedback function that relates our outputs to our inputs.

Martin says: "If there are any 'absolute' standards, they will be those that have allowed the social groups using them to survive and prosper." Conventional standards (like the clockwise-turn standards) can be "absolute" to the extent that we can get all objects to abide by this convention. This can be done in principle, though it's difficult (and sometimes not desired) in practice; some people might have a need for a counter-clockwise-in device. But the goal of absolute standards (conventions) is at least feasible for inanimate objects, because these objects have no purposes of their own that might conflict with the convention. Such is not the case with living systems.

The problem is that people are not inanimate objects—and certain individuals in certain circumstances might find that acting according to a particular convention is impossible, not because the person is bad or contrary or immoral, but because he or she is a hierarchical control system that simply cannot act like the knob on a radio. So my argument against "absolute" standards applies as much to standards as social conventions (like grammars) as it does to standards as moral principles. I am all for standards as conventions. The notion of absolute standards—no matter how technologically and socially helpful their existence might be—is inconsistent with human nature (if people are hierarchically organized perceptual control systems). This does not mean that I believe everybody should just go off and do their own thing. I'm just saying that this fact about human nature must be taken into consideration when we think about how people can act cooperatively.

The people who want there to be absolute standards are not "bad" people (from my point of view). The desire for absolutes is quite reasonable. I can understand that desire—especially with respect to people. People should never kill each other or end a sentence with a preposition; people want predictability. All I'm saying is that people are not switches; they cannot abide by such absolute conventions, even if they try. This does not mean that social chaos is inevitable; what I think it means is that we have to find ways to cooperate that take into account the true nature of human control systems. The fact that cooperation is possible in the context of this reality (the inability of control systems to control relative to absolute conventions) is evidenced (I think) by the general spirit of cooperation found despite the diversity (in terms of many conventions) among members of the Control Systems Group itself. It can be done.

Martin Taylor: What I intended was to suggest that a "standard" provides a convenient level at which a reference value can be set, one that has often been found (perhaps by other people over history) to result in a desirable percept. But even with "absolute standards;" there's no compulsion on anyone actually to use them as reference values. As Rick says, such use might conflict with the ability to achieve other reference values. Some day, you might have to try to kill someone if you are to maintain other desired percepts, such as personal survival or freedom.

The existence of absolute standards depends on whether over evolutionary time certain behaviors (in the perceptual-control-theory sense) have benefited the survival and gene-propagation of the people (or others) using those behaviors. If they have, then either by gene transmission or by social transmission, the ordinarily effective behaviors will result in absolute standards. (On social transmission, see F. Boyd and P. J. Richardson, *Culture and the Evolutionary* Process, University of Chicago Press, 1985.)

I find no moral connotation to the idea of "standard," whether absolute or not. The idea of "absolute standard" as "you have to do what I say is right" is, I think, morally and practically repugnant, for many of the reasons adduced by Rick. But "absolute standard" as "that's what people have learned as a usually effective way to behave" is simply a practical concept that improves social interaction.

Rick Marken: I thought Martin was proposing that "standards" be understood as conventions for behavior. For example, there is a convention in the U.S. that we drive on the right. So, when I am on a road, I try to keep my car in a lane to the right of the center line. With regard to perceptual control theory (PCT), this means that I set my reference for the relationship between car and center line at "right of" rather than at some other value, like "left of." I was agreeing that standards of this sort are quite useful for successful social interaction.

Martin says that he "intended... to suggest that a 'standard' provides a convenient level at which a reference value can be set, one that has often been found (perhaps by other people over history) to result in a desirable percept." I agree, except that I think many of these standards (such as which side you drive on) are fairly arbitrary—they work as long as there is agreement among those who need to abide by them in order to avoid interpersonal conflict.

Martin goes on to say that "even with 'absolute standards; there's no compulsion on anyone actually to use them as reference values." Well, there is some social coercion. People can have unpleasant run-ins with the police if they pick the wrong side to drive on. Of course, one is still under no compulsion to set their reference at the conventional level, since he or she is the one setting it.

Then Martin says: "Me existence of absolute standards depends on whether over evolutionary time certain behaviors (in the perceptual-control-theory sense) have benefited the survival and gene-propagation of the people (or others) using those behaviors. If they have, then either by gene transmission or by social transmission, the ordinarily effective behaviors will result in absolute standards." If by "behaviors" you mean "references for certain inputs," then I agree; there might be absolute (fixed, built into the individual, unvarying) references for certain inputs. Such references are almost certainly at the cellular, if not the genetic, level—they are called "intrinsic references" in PCT. If, however, by "behaviors" you mean particular actions, then I don't see how this can be correct; evolution could not possibly select for actions that would have to produce their effects in a disturbance-prone environment. I think a lot of sociobiologists imagine that certain behaviors (in terms of actions) can evolve; for example, they talk about evolution of "aggression." It sounds like they are talking about the evolution of certain visible patterns of outputs. I think the only thing that might be able to evolve is a preference for a certain level of sensory input resulting from these (and / or other) actions.

Finally, Martin says: "But 'absolute standard' as 'that's what people have learned as a usually effective way to behave' is simply a practical concept that improves social interaction." It sounds like you are saying that an "absolute standard" is only *relatively* absolute (it is usually effective at improving social interaction, but not always). If this is what is usually meant by "absolute standard;" then it turns out that I have been advocating a version of this approach to "absolute standards" all along. I've just been saying that some standards are usually effective for lots of people—but not always (they don't work for some of the people some of the time). I just wish some of the others in the discussion of absolute standards would have clarified this point for me. Does this mean that the Ten Commandments are "absolute standards" in your sense of absolute standards—it is usually effective to not steal, but not always? Is that what Judeo-Christians think God meant? What about that first one: thou shaft have no other god before me... usually? Some people got stewed for not obeying that one. Are some standards more absolute than others?

Bill Powers: Just a few ideas to add to the standards discussion. Any given standard, such as "helping the poor," has at least five aspects:

1. The verbal description or name of the standard ("helping the poor").
2. The perceptual meaning of the description or name of the standard: that is, how you can tell when a poor person is being "helped"?
3. The reference level for the standard: that is, what degree of the helping is the desirable degree?
4. The program of actions used to achieve the standard: that is, what actions will help the poor to the desired degree?
5. The system concepts exemplified by the standards: that is, the concept of human nature and of society that defines the goal achieved by helping the poor.

Most discourses on standards focus on the verbal description or name of the standard, under the (incorrect) assumption that it indicates the same principle to everyone. So when old-style Democrats speak of helping the poor, they mean giving them money, advice, and services that the poor people can't obtain for themselves. When Republicans speak of the same thing, they mean doing something that will eliminate the need for giving things to poor people—enabling them to get what they need for themselves, teaching "self-reliance."

The Republicans quite rightly claim that simply giving things to poor people will keep them dependent and poor (they don't learn how to control their own lives). The Democrats quite rightly point out that simply demanding self-reliance ends up punishing people for being poor and creates callousness toward human suffering. Republicans assume that people work in order to maintain a viable economic system that's essential to everyone, and because of financial rewards and incentives. They assume that the healthy society is one in which the members compete for wealth and predominant positions or power. Democrats assume that people work to improve the quality of their lives outside the economic framework, and that the healthy society is one in which nobody has to labor overly long, under unpleasant or dangerous conditions, or in a state of social inferiority. At least that is my view of the "canonical" positions of the two parties. I speak, of course, as a time traveler from a different era.

It's impossible to agree on standards without agreeing on system concepts: the kind of society we live in and our own human natures.

Simply hurling the names of standards back and forth and claiming that they are good gets us nowhere. Even agreeing on the means of achieving standards requires a shared concept of human nature. Those who enjoy power and wealth quite rightly appreciate the advantages of these things; they advocate principles based on the assumption that everyone would be better off with power and wealth, and principles that will help those who already have power and wealth to keep them. Those who value other goals assume correctly that nobody is permanently better off with power and wealth unless everybody has them, and favor principles that spread the wealth even at the expense of those who lose out by accepting the principles.

When we speak of standards as shared principles, we tend to forget how little of this sharing there really is. The story of standards in human societies is a story of conflict, not sharing. This is true in all sizes of groups from the dating couple, through the family, through a whole country. Even when people say, in words, that they agree on a standard, they perceive it differently; even when they perceive it in more or less the same way, they differ on the reference level. We can agree that many poor people need immediate financial aid. Which people? How much? To be spent how? In whose Congressional district?

The other comment I have is more general. We tend to speak of standards in terms of their effects when they are shared, in terms of their roles as characteristics of a society, or in terms of what they do for social interaction. From the theoretical point of view, however, the questions are not just *what* standards are adopted and why they are adopted, but what a standard is, and *how* it can have any effect.

How does a standard influence the behavior of any individual? How does it get communicated? What has to happen inside an individual before the words describing a standard come to have meaning to that person? And what has to happen inside the person in order for any particular interpretation of such a description to attain the force of a reference condition? Without these processes internal to the individual, no standard can have either meaning or effect. We have to understand standards as they exist in and operate in a single person before we can understand how they work in a world populated by many persons.

Finally, we often speak of the advantages or influences that standards have in a society. I think that, very often, these advantages or influences are hypothetical—they're what *should* occur. But I doubt that such things very often *do* occur.

Rick Markers: Bill Powers says: "Finally, we often speak of the advantages or influences that standards have in a society. I think that very often, these advantages or influences are hypothetical—they're what *should* occur. But I doubt that such things very often *do* occur." This has been my point all along—at least in terms of personal (and, to some extent, in terms of interpersonal) control. I see no way in which perceiving certain standards at certain reference levels can *necessarily* lead to successful control of any other perceptual variables—or intrinsic variables. Yet this is an article of faith for many people in society. I imagine that if one of us showed, quantitatively and experimentally, that this faith is not correct, he or she would soon be the victim of a holy war. Now that I think of it, if really good science on control of principles and system concepts were done, it is possible that the results would make the religious / political / scientific establishment take steps that would make the Catholic Church's treatment of Galileo look like a picnic.

One of the things that has particularly irritated me about the current political dialog about values (the one going on in the outside world—not on CSGnet) is that the people who are pushing "family values" most ardently are also the people who have most ardently pushed one of the most fundamental (and, I think, destructive) values of U.S. society—the value of *conflict* (also called *competition).* Every red-blooded American knows that competition is what makes for successful economies. The basic idea (as pink-blooded little me understands it) is that consumers are like judges at a beauty contest (a uniquely American event, itself). Producers (or goods and services) compete to win the patronage of the customers. This competition leads to better and better products from producers (in the sense that they are the products that best meet the customer's needs or wants).

This scenario has one little problem that only Americans with pink-tainted blood might ever even deign to point out; in competition like this, there are generally winners and losers. What happens to the losers? America doesn't like losers, so we ignore them or blame the loss on personal failings (not being a *real man* or a real *woman).* Pinko types like me, however, don't think that losers are just valueless trash; they are worthwhile control systems, with intrinsic reference signals of their own. I worry about the losers because societies with lots of them around tend to be very precarious—and have to take strong measures to make sure that the losers don't try to just take stuff from the winners.

I don't like the "value" attached to competition in this society. I like the "value" of cooperation and community. I think society's emphasis on the importance of "being #1" or "fighting to get to the top" is far worse than the lack of emphasis on "family values." But I doubt that Quayle and Bush will come out in favor of the value of "cooperation" and "community." Do I have bad standards? Is it wrong to dislike competition and to like cooperation?

I will admit that competition (conflict) can accelerate the development of technologies that might help the parties to the conflict "win." Thus, two companies making widgets might progress faster toward the goal of making the "best" widget (the one that satisfies the market best) because they are in conflict (they have to keep improving the widget—the output of each system—or lose the conflict—have their market share of widgets become much lower than their reference).

I think it is this "good" result of competition that has impressed economists. But is this the only way to organize an economy that produces the widgets that we all need to control what we want to control? Must there be winners and losers in order to have an economy that meets the requirements of its members (the winners, anyway)? Can't we organize a society in which everybody is a winner (can control what they need and want to control)—and can't we do it without coercion (the approach that communism used)? It seems to me that the economies of some of the Scandinavian and Western European societies approach a nice compromise between capitalistic individualism and socialistic communalism. Why don't we learn from those economies?

Kent McClelland: Rick, although I agree with and indeed applaud your sentiments favoring cooperation over competition, I wonder whether you're making the choice sound a little too simple. An interesting book by Michael Billig and associates *(Ideological Dilemmas: A Social Psychology of Everyday Thinking,* London: Sage, 1988) has convinced me that such things are not a matter of either / or, at least not in our usual modes of thinking. The book traces the history of Enlightenment thought and shows how contradictory values are built into the public discourse on such issues. Racists, for instance, will typically preface their biased remarks with a disclaimer to the effect that they themselves aren't prejudiced against blacks, but you really can't get away from the fact that, etc., etc. I have no doubt that Bush and Quayle could come up with many heart-warming remarks about the value of community and not see any contradiction at all between that sort of rhetoric and their views on competition.

I think the question is how stable system concepts can come to be constructed from an amalgamation of values or principles that are often contradictory in practice. But maybe such mental and moral flexibility is necessary for us to maintain the perception that the world we observe is consistent with our preferred system concepts. As I believe Rick pointed out, a control system that was stuck with a single reference signal for a principle like honesty (or cooperation!) would be unable to vary its outputs to maintain control of perceptions of the next higher level, just like an arbitrary restriction to a single setting for arm position would cripple your physical control of bodily movements.

Bill Powers: Kent and Rick, cooperation and conflict are outcomes of a social interaction. If people's goals are aligned, there will be cooperation, or at least non-interference. If they are not, there will be conflict and competition.

Competition arises in our society as a consequence of system concepts and principles. One of these concepts has to do with position in a social hierarchy. The idea of the superior person, with others being inferior, sets the stage in some people for a desire to be, or be acknowledged as being, at the top of this social hierarchy. As achievement of this goal requires a relative ranking of people, it is impossible for everyone in the society to achieve it. If even two people wish to be perceived as number one, a conflict must arise, because by definition only one person can be number one (or number anything). The existence of number one creates number two: number n implies number n -1. If one person wants to become a leader, followers must be found, and others who also want to be leader must be fended off, undermined, or otherwise prevented from succeeding. The striving for social position is a pernicious ill in our society, which accounts for a great many of its problems.

I've heard all of the arguments in favor of competition. I don't believe them. I don't think that people with contradictory goals accomplish anything but building up their muscles and cancelling the effects of someone else's muscles, leaving little effort available for real progress. I don't believe there is a "top" in the social hierarchy—I don't even believe there is a social hierarchy. And as long as I don't believe that, there is no social hierarchy for me. This doesn't endear me to people who want such a hierarchy to exist, but that's their problem. There's nothing I want from anyone that would make it worthwhile to play that game. Not even the privilege of living.

And I know for certain that when, in some microsociety, people manage to do without this concept of Number One, everything magically works better: shared goals are accomplished smoothly, easily, and with great pleasure. People get smarter, because they aren't wasting their time and effort trying to counteract what someone else is doing.

I haven't got this system concept worked out in any detail—talking about it too much tends to reduce it to procedures and slogans, anyway. But what I do understand of it, I want to sell. It defines the kind of world that I find worth living in. All I can do to create that world is to persuade others who will persuade others that it's worth a try.

"Leadership," it seems to me, is a role in a social hierarchy. It requires followers. It opens the door to competition and conflict ("I can lead better than he can, so follow me and not him"). The worst result, from my point of view and in my circumstances, is that followers learn from a leader how to follow, not how to explore, teach, and learn.

The attitude of followers toward leaders, in my experience, often tends to be one of admiration, deference, blind loyalty, and even hero worship. It's the attitude of a child toward a favored adult. Many leaders like being on the receiving end of this attitude. It confers power, it allows the leader to indulge in egocentric thinking, it protects the leader from criticism and accountability. The leader can arrive too easily, with the connivance of the followers, at the idea that he or she makes fewer mistakes than ordinary people do. The leader can point to the support of the followers as a way of showing others, outside the group, that there must be something superior about the leader (so they would be better off becoming followers, too). Leaders are corrupted by their followers, and willingly.

Dag Forssell: The idea that a leader is defined as someone who has followers is indeed the predominant interpretation in our society. I believe it is an unfortunate one. It is not the only one available.

I can't conceive of a control system wanting to follow. What a control system wants is good system concepts to inspire good principles, so you can select effective programs, ..., so you can maintain your body chemistry. A control system is *designed to lead itself;* to satisfy its own purposes as it perceives them. "Purposeful Leadership," as I define it, is the development and communication of good information that allows every individual to lead himself / herself in full autonomy. It is a non-manipulative, non-coercive, non-violent approach.

With good information shared and internalized voluntarily, people will be aligned and will automatically cooperate on the mutual concerns.

Bruce Nevin: Dag says: "I can't conceive of a control system wanting to follow." Oh, come on, Dag! You can't mean that, can you? Aren't there many occasions when one control system wants to follow the lead of another control system? And is this in itself pernicious? (Though it can be abused—on both sides of the dyad, be it said! Nor does it end with childhood. Nor is it always childlike, though abuse of childrens' dependency does seem to result in many adolescents and adults coming to abhor and scorn it and fear exposure of it in themselves. One of the sure recipes for childishness.)

Have you ever taken a dance class?

Dag Forssell: OK, Bruce. A difficulty on this net is that anything can be and is taken so damn literally. You have a point, of course. I did learn to lead in dance once upon a time. It is important in ballroom dancing to give clear signals to your partner (follower) with a steady hand. The follower chooses to follow and concentrates on that.

Leadership is often understood to mean that you tell someone *what to do,* then they follow by doing what you tell them to do. This emphasis on doing and instructions fits nicely in a cause-effect world.

I am trying to redefine and sell "Leadership" as the idea that if you want to lead, the *most effective* way is to offer good (a description of that Boss Reality that is as good as you can make it or negotiate it in open discussion) information for your "followers" to evaluate and make part of their own system concepts if they want to (understand, no conflicts with pre-existing concepts, relevant, etc.). Then you step out of the way and let the "followers" control to their hearts' content. You will not need to supervise or "control" their actions, because that is built in.

This form of leadership is inherently non-violent. Teaching it will not work if the top management in a company is coercive, as I perceive most to be. Therefore, the idea *must* be sold at the very top, to the very people who are used to insisting on results or else. I am counting on finding a few who will see it my way, but don't expect many. A few is all I need. Once the process is understood, the leadership / information can come from anyone in the organization.

Bruce Nevin: Following surely cannot mean producing the identical behavioral outputs. We know this because of the variability of behavioral outputs with respect to the reference signal. (Or with respect to the outcome, more or less equivalent depending on success of control.) Nor can it mean assuming the identical reference signals for identical (or equivalent) controlled perceptions. We know this because all the follower has to go by are the behavioral outputs of the leader, among other environmental variables, plus memory and imagination, of course, which are the means for projecting, anthropomorphizing, and so on, which we necessarily do all of the time.

There are two corresponding questions for the other member of the dyad: Can a control system want to lead another control system? Can a control system lead another control system?

From the existence of a large literature and a long history of 'leadership," it seems clear that a control system can want to lead another.

It seems to me clear that A can lead B only to the extent and in the manner that B wants to follow A. This is why virtually all of traditional thinking about leadership boils down to "motivation"—getting others to want to follow you. (Ditto for pedagogy)

Assume that B wants to follow A. The extent and manner depends on B's other goals. B can follow just in terms of proximity. This kind of following ranges from detailed mimicry (mirroring) to very slight correlations, such as B following A with his or her eyes.

Much of what we mean by "follow" is metaphorical, with this literal sense as a basis. We can easily identify the metaphor when we say B is "following A's argument" or "following A's line of thought." The metaphor is not so obvious, perhaps, when we talk of B following A in the sense of coming to A for directions, going off and executing them, and coming back to A for more.

"Following directions" seems to mean to control one's perceptions so that they mimic ("follow") the perceptions that one imagines on hearing or reading the directions.

"Following A's argument" seems to involve imagining the argument for oneself and finding that the imagined line of argument corresponds with what A has said and is saying.

To paraphrase P. T. Barnum, some of the people want to follow all of the time, all of the people want to follow some of the time, but not all of the people want to follow all of the time.

If B is not confident and purposeful in a given situation, B might seek someone to follow until in a situation where B is more confident and purposeful. (Purposeful: has clear goals, is controlling for them without major conflict) We can discuss why this is so.

If I am B in such a situation, I will follow one who appears confident and purposeful rather than one who appears unconfident and irresolute. We can discuss why this is so.

Some people are unconfident and irresolute and conflicted in much of their waking experience. I suspect that many such people came to be so because of childhood experience with adults who emphasized conformity with external authority and arbitrary standards, enforced in punitive ways.

It can happen that such a person feels confident and purposeful in an institutionalized social context with clearly assigned roles and relationships of relative authority, in accord with standards established for those institutions. Such people can become "leaders" within that framework. They know "the system." They become very anxious outside it, and resist contradiction to it. I think that outside the system they fear unexpected punishment; my experience is that outside the system (that is, in circumstances in which they can no longer interpret their perceptions as within the familiar institutional context) they become unconfident and irresolute. They often despise indecision and lack of confidence. (Such people, by the way, are unlikely to be drawn to HPCT at this stage in its history. And this parallels the familiar left / right ideological dichotomy.)

I suggest that charisma depends in part upon the appearance of confidence and purposefulness. As you have suggested, Dag, this connects with sales and marketing, where the pumped-up appearance often outstrips the basis of confidence, and the real purposes are ulterior. But charisma can be genuine. When you're looking for the exit in a crowded waiting room, a person walking quickly in one direction with a suitcase has some charisma.

The ad hoc situational leadership and functional (not authoritarian) hierarchies of anarchism depend upon this, especially in cases where the participants lack detailed knowledge of another's capacities. "You seem to know what you're doing. How do you think we can make this go?"

Now: Can a control system manipulate another control system? Can a control system exploit another control system? I believe these are some of the negative senses of "leadership" and "charisma" that you are resisting. Am I right? I think HPCT does not show that these do not exist. 1t only shows that they cannot work as intended. Social institutions can help people persist in being slow learners about this.

Dag Forssell: Bruce, when I said that "I can't conceive of a control system wanting to follow," I did not mean to be so literal. I meant that it is not the nature of a control system to "follow," whatever that is. I appreciate your post. Your restating my points, paraphrasing rather, is a very good thing. It shows me how my careless wording can be (mis-)interpreted. You are doing a good job of sorting out technical alternatives and aspects of "following."

I am resisting what I perceive to be extremely common stereotype interpretations of leadership and sales, where I sense an interpretation that leadership and sales are indeed "manipulation" and "exploitation." This I read into Bill's original refusal to lead and some comments about sales at past CSG conferences. In turn, this leads to an aversion to consider these major applications of HPCT. Still my perceptions, of course.

If you substitute "manipulate" with "inform" or "guide" or "enlighten" or "teach," and "exploit" with "mutual benefit," the substance of the interaction does not change from an HPCT point of view, but the emotional, stereotype flavor changes dramatically. We are still talking about leadership and sales and mutual economic advantage.

Certainly the members of this net want to sell HPCT to the world. Is this "manipulation" and "exploitation"? I would not label it that. But mention leadership and sales. What comes to mind? Some brutal, selfish "leader" on the one hand, and pusher of overpriced junk nobody wants or needs on the other.

These terms are among the unexamined "human pie slices"—system concepts from pre-HPCT days—that can benefit from some HPCT light. By looking closely at this, perhaps a way to sell HPCT can be found, vastly superior to the frustrating sales efforts in the psychological journals that are discussed on the net, but are not labeled as such. (These journals are a minuscule market compared to the rest of the world, and the one market where we know that HPCT is not welcome).

The way there is to forget about "manipulation" and "exploitation" and instead examine the best interest of and control processes in the other autonomous control system, whether we call it follower or buyer. This done with full visibility to said follower and buyer, of course. There is no need to hide the interest and control processes of the leader or salesperson, either. The exchange of goods or services should benefit both parties. Otherwise we have reverse manipulation and exploitation.

Leadership and sales both can be honorable. HPCT can show how.

Rick Marken: PCT will always have a hard time. People just don't like to believe in autonomy for anyone but themselves. And they will apparently continue to wage war on autonomy even though the consequences of that war are precisely the opposite of what they hope to produce. I am speaking of the "war on drugs"—the greatest and most sustained crime creation program in history. Here is a clear case of trying to do, at a societal level, what we have agreed is useless on an individual level: society is trying to forcibly change the reference level of a controlled variable (drug usage)—trying to force it to zero for everyone. I object to this idiocy, not because I want to take drugs (the usual assumption about those who want to end this drug war idiocy), but because things I care about are endangered.

The only solution is to go up a level ("what do you folks really care about?") or have a police state (a temporary "solution," at best). It looks to me like a solid majority would choose the police state in a second.

If there is a fundamental postulate of PCT, it is that organisms are control systems. A functioning control system is able to make its perceptual experience match its references for that experience; I call this "autonomy"—the normal operation of a control system. Anything that prevents normal operation is the cause of a malfunction. Conflict is an example of a control system malfunction; conflict prevents autonomy—i.e., the ability to control.

The drug war is an example of control systems *in conflict*. *So* the drug war is an example of control systems that are *malfunctioning*. There is no moral judgment here; that would imply that I *like* the goals of one group (the drug warriors) better than I like those of another (the drug takers). In fact, I personally don't care for either of their goals, but that *is not* why I don't like the drug war. I don't like it because there is *conflict* between control systems; this conflict *might* have unpleasant side effects for me (I might get robbed by a druggie who has to pay high prices for highly abundant substances, or have my house broken into by an overzealous SWAT team that's off by a digit on the address of a crack house). But the chances of those side effects are fairly low. I really object because conflict prevents the functional operation of the control systems involved; neither party (drug warrior, drug taker) is able to function as a full-fledged hierarchical perceptual control system.

The conflict would be solved, of course, if the druggies decided to stop taking drugs *or* the warriors decided to stop fighting drug takers (and suppliers). Either approach would end the conflict and people could start functioning again. I favor a solution to this conflict based on the drug warriors changing, because they are the ones who created the conflict by trying to control other control systems. The other control systems (druggies) maintain the conflict by maintaining their references for the perception that the drug warriors want them to change. But somehow (and it's hard for me to articulate it without becoming moralistic) it seems to me that it's a lot easier for the drug warriors to stop controlling for what the druggies are controlling than it is for the druggies to change their own reference for what they are controlling.

The reason the drug warriors are the problem is because they must push against another control system in order to control the variable they want to control. The victim (the druggie) could (and did until the drug warrior came along) control the variable s / he is controlling without creating conflict in another control system at all. So one set of control systems (the warriors) are *creating* conflict by trying to inhibit the autonomy (not consciously, but that is what they are doing) of others. Since the warriors don't understand PCT, they are creating this malfunction out of ignorance. So I still have no moral complaint here. The drug war is just malfunction—producing idiocy (stupidity) that results from a failure to understand the nature of autonomy. So problems like the drug war can be solved, not by trying to articulate better moral principles, but simply by understanding how control systems work. A person who understands control theory simply shakes his / her head in dismay at drug warriors—just as a person who understands plumbing shakes his / her head in dismay at somebody pouring grease down a drain. Both are just watching people create malfunction.

Ed Ford: It seems to me that when two or more living control systems find themselves in the same environment, in order for them to live in harmony and cooperatively, they have to agree on a way things ought to be, a system of concepts, which are best expressed and set forth by agreeing to a set of standards upon which they base their choices as they attempt to find satisfaction while living together. (I see standards as synonymous with rules, criteria, principles, guidelines, etc.) Thus the needed harmony between levels of the hierarchy in social groups. As they live their lives, trying to satisfy their own individual goals, the choices they make, if based on agreed-to standards, will more than likely make it easier for them to live in harmony with each other.

In the order of nature, we first learn to follow standards as children at home and then, ultimately, to set our own. For us to live in harmony, we must always set rules while respecting the rights of others. Whether at home, at school, at work, in an institutional setting, or just buying gasoline, we are constantly surrounded by standards and rules. Thus the need to learn to follow standards as well as to set our own standards while resolving our internal conflicts.

Any time I deal with anyone, whether in private practice or elsewhere, standards and rules are a part of life. Whether it is setting standards for the kind of spouse we want (thus to help us make a choice), or wanting to get along with a parent, or getting through school, or interacting properly on the net, etc., standards are a part of life. However, in order to help living control systems resolve their own internal conflicts and to teach them how to deal with their lives, there has to be a basic understanding of standards and rules and of how consequences and choices are integrated into the standards concept already established in the setting where they are being taught.

My experience over the years has taught me that there are tremendous differences in the understanding of the role of standards, the meaning and place of consequences and choices.

Rick Marken: I just don't get it, Ed. What does "teaching standards" have to do with a PCT-based view of human nature? What I get from PCT is the idea that nothing could be less important—the actual substance of a person's references for relationships, programs, categories, principles, "standards," etc. matters only in terms of how these satisfy higher level goals. The system should just be error-free—and this happens by having working (conflict-free) control systems. Of course, such systems will be setting the "right" references for perceptions like your "standards," but they are right from the perspective of the control systems (they combine appropriately with prevailing circumstances to achieve the higher level goals). What is at any time a "right" setting for a particular standard from the point of view of the control system might very well appear to be a wrong setting from the point of view of someone who "knows the right standards." I know that some of the people you are dealing with have interfered seriously with other people in their efforts to achieve their goals. So, obviously, your goal is to teach them to act without hurting others, i.e., "follow the rules." I think this is great, but you should be clear that this focuses your treatment strategy on getting a person to act in ways that are better *for* you—and, incidentally, for the person him / herself.

A person who wants to perceive him / herself as socially cooperative would be creating a big conflict for him / herself if, for some reason, the reference for a perception with a socially accepted reference (like wearing clothes in public) were changed to a different value. But I don't believe that there are *any* "standards" perceptions which, if controlled at a particular reference level, would be intrinsically internal-conflict producing; conflict depends on what other perceptions a person is controlling and at what level. I think Ed believes that there are certain intrinsically intrapersonal-conflict-producing standards-perception reference settings.

I do believe that there are settings for references for standards perceptions that produce interpersonal conflict—there are *lots* of them. Such conflicts occur because carrying out the purpose tends to produce disturbances to *intrinsic* variables in the other person; there is a biological basis to much (but not all) interpersonal conflict.

Martin Taylor: Rick, when there is conflict, there might be reorganization, and as Bill has often pointed out, that reorganization will tend to drive the conflicting systems into less conflict. If I do not conform to your standards, we both experience conflict if you care enough to try to make my actions conform (you can't see what I am "doing," but you can see my actions), and if your efforts make me unable to satisfy some references. So, point 1, it is not just me who experiences conflict and might reorganize. You might, too.

If a community has developed / evolved a set of standards that results in low levels of conflict when everyone adopts those standards for their actions (again, not for what they are "doing"), the standards will be rather stable. They work, because whatever people are "doing," their actions permit them to control their percepts adequately. That's what is meant by low levels of conflict. If the "standards" don't have this effect, and people find that they experience high levels of error when acting according to the standards, some people will reorganize one way, some another (it's random, after all), and the standards will disintegrate, perhaps to re-form as a new set of standards that provide lower overall error rates. Sets of standards that lead to sustained high error levels in many people are not stable. So point 2 is that if many people adhere to standards, it is because those standards do not conflict with the ideal of low intrinsic error.

I agree that there probably are no standards that we could call "intrinsic," but there are probably some reference levels that cannot be components of stable community standards. These will not be found in the standards of viable communities. But sets of standards probably fit together in clusters that are stable as a group that can be taken into or left out of a total system of standards. Different sets of precepts based on the teachings of long-lasting religions probably form such groups. I would imagine that the number of such sets that could be stable is unlimited, but the societies of the world might have found only a few tens of them.

If an individual lives in a community with stable standards, but does not use them to set the relevant reference levels, that individual will find conflict in many of his / her interactions with other members of the community, whereas the other members will find conflict only in interactions with the deviant. The deviant is more likely to reorganize than are the other members, and if the standard set is truly stable, this reorganization will continue until the deviant acts according to the standards and, at the same time, finds his / her reference levels generally attainable by non-deviant actions.

I suspect that most sets of social standards are not truly stable, and perhaps there are no possible sets of standards that lead always to zero error as a consequence of interactions. In a non-stable, or conflict-retaining, set of standards, all members of the community are liable to reorganization, and the standards themselves will drift in a way directly comparable to linguistic drift—and for the same reason. The result could be the breakoff of heretical groups, or a more or less unconscious shift of mores, or other shifts.

Serious problems arise when individuals who belong to different communities with incompatible standards have to interact. The incompatibility of the standards sets is defined by the existence of conflict when one individual uses one set and the other the other set. One or both must reorganize. When you have large numbers of individuals from each community meeting, then either one community will lose its standards (its "culture") to the other, or both will have to develop supplementary standards to deal with the interactions. That way lies stereotyping of members of "other" communities, but it might be a necessary way to handle the modern possibilities for world-wide interaction.

It's all based on the iterated interactions of individuals, and one-on-one reorganization based on the conflict that occurs.

Rick Marken: Martin says: "If I do not conform to your standards, we both experience conflict..." Not necessarily true. Here you are talking about interpersonal conflict; we "experience" it only in terms of the success (or lack thereof) of our efforts to control variables (in my case, perceiving you as conforming to my standards; in your case, perceiving no loss of control as a result of my efforts—beating you, starving you, locking you up, etc.—to get you to conform to my standards). If one person is a lot stronger than the other, s / he will "experience" no conflict at all in this conflict; s / he will just get the result they want. If both people are about equally strong, they will experience loss of control, i.e., error with respect to some variables they are trying to control. Of course, being people, each will also be able to perceive the cause of his / her lack of control: the other person.

My problem with this whole analysis is just the emphasis on "standards"-setting as a basis for harmonious interactions in groups of

control systems. I think this is almost certainly a crock. As humans, we do happen to be able to perceive at the system level, but that doesn't mean that controlling perceptions at this level is any more important than controlling perceptions at other levels. Herds of animals, for example, work together just fine without agreement on (or ability to experience) standards, system concepts, principles, categories, or whatever. Most everyday conflicts between people are usually over control of perceptions that are at lower levels than "standards"—and people work them out just fine.

I think organisms in groups "get along" when there are a sufficient number of perceptual degrees of freedom (df) to be controlled—and sufficient environmental df to allow all members of the group to control their perceptions. This means that the organisms must be able to perceive the environment in a way that allows simultaneous solution of the perceptual df problem in the constraints of the environmental df. Tom Bourbon's studies of two people controlling the relative distance between lines on a screen contributes more to our understanding of what makes it possible for multiple control systems to "get along" than does all our blathering about standards-setting. Standards are just one thing people have to be able to control—no more or less important than controlling sensations, configurations, transitions, etc. When people can control their perceptions—and when each individual in a group can control his or her own perceptions—then there will be no interpersonal conflict. This is an achievable goal, but to get there, we have to look in the right place; *not* at figuring out what standards people should set, but at figuring out how to provide people—*all people*—with the degrees of freedom necessary to control their own perceptions. We already know how to do this, actually. PCT just shows *why* this is important: The ways to do it are (1) *population control* (to preserve the available df); (2) *education* (to learn about the available df for controlling our own perceptions—and how to control those perceptions more effectively).

People have tried to solve their problems by finding the right standards for centuries (from the beginning of recorded history)—it not only doesn't work, it is the cause of most of our intractable problems (nationalism, religious wars, etc.). I suggest that we approach the problem of interpersonal interaction from a PCT perspective; if people really are input control systems, then PCT should have some scientifically and practically useful things to say about how multiple control systems can get along without conflict. I think the answer is "degrees of freedom," not "standards."

Martin Taylor: Rick, I'm a bit confused. I knew this was foolhardy territory to get into, but I can't see what the discussion of interaction procedures that evolve into conventions known as standards has to do with the system level. You seem to be implying throughout your response that the only place where standards exist is at the system level. I intend the term to apply at all levels relevant to interactions among people, and I think it applies probably more to actions than to behavior (using the PCT distinction that behavior is the control of one's own perceptions, whereas actions are not). Standards include greeting patterns, dress codes, thank-you notes for gifts, and all sorts of things for which the external appearance is what matters.

If I can act according to the standards of my community, and nevertheless control my perceptions with little error, I won't reorganize much, and I will continue to act according to the standards. If I don't act according to the standards and nevertheless am able to control my percepts, I won't reorganize. But in most cases, if the standards matter to many of the people with whom I interact, I will find that not acting according to the standards might impede my control (I might not get the job because I didn't wear a suit to the interview; I might not get a gift from Aunt Mabel because I didn't send a note thanking her for the last one), and I am likely to reorganize. When my reorganization leads me to act in such a way that I maintain control of my percepts, I will no longer reorganize.

Real community standards are those that tend to induce reorganization in people who don't act according to them. As I said before, their stability is determined by the degree of conflict occasioned on average in people who abide by them, because error will lead to reorganization, and if there is a set of standards not very different from the current set but that tend to lead to less error, then the community standards will drift in the direction of that set. The word "community" is diffuse here. It is clearly weighted by the probability of interacting with any particular person, so for most people, I suspect the standards one develops will be closest to those of the parents and older siblings, at least when interacting with them, though other sets of standards might be developed for interaction with others (such as the local gang).

As you can see, I don't think standards are anything that people control (or other herd animals, for that matter). They are the products of reorganization, not percepts. They are the ways that percepts can be controlled when other people are involved in the actions that together form the controlling behavior. All the same, I suppose that people can model desirable organizations, talk about them, and explicitly teach them to the young. But the problem here is how you teach any behavior deliberately. The "standard" you can talk about is a model or a simulation, not the result of a structural reorganization. "Standards" is the result, not the instigator. It is the manifestation of the dynamics of an uncontrolled interaction among control systems, not a prescription

for what should happen. It becomes a prescription by methods fully intelligible within classical PCT.

Rick Marken: Martin, the talk about standards is highly ambiguous—sometimes I think people are talking about reference levels and sometimes about perceptual variables. I thought we had clarified it earlier a bit—my conclusion was that Ed Ford (the main "standards" guy) uses the word "standards" to refer to "higher level perceptual variables"—types of perceptions that might be described by words like system concept, principle, value, belief, etc. He tries to help people set the "right" reference levels for these perceptions. So my reply to your post was really aimed at Ed—I just don't think control of higher order variables is any more important in social interactions than control of other perceptual variables.

You say: "Standards include greeting patterns, dress codes, thank-you notes for gifts, and all sorts of things for which the external appearance is what matters." So what you mean by standards is "perceptual variables that involve another person." Well, now we have another possible meaning for "standard." Why don't we just stick to the PCT-model terminology (and semantics)?

And you say: "As you can see, I don't think standards are anything that people control..." Boy, you've got me. In the quote above, it sounded like standards were social perceptions. Now they are something that can't be controlled. And yet people reorganize when controlling them produces conflict. So it must not be failure to control standards that is leading to reorganization. But the reorganization leads to new, stable standards. So standards are a perceptible (to Martin) side-effect of reorganizing to control perceptions that are not standards? In other words, people control perceptual variables; this can appear to an observer as a process of converging on social standards. Is this it? If so, I completely agree.

Martin Taylor: Rick says: "So what you mean by standards is 'perceptual variables that involve another person.' Well, now we have another possible meaning for 'standard.' Why don't we just stick to the PCT-model terminology (and semantics)?" I've been trying to stick very precisely to the PCT model, but I don't know of any standard terminology to handle what we are talking about.

The problem with any definition of "standard" is that it is something (let's not say what) that one person applies to the observable actions of another. A person might apply standards to himself or herself, but only as an observer, possibly in imagination, of his or her own actions.

Standards have a funny status. I cannot control your behavior, because I have no sensory information that allows me to perceive it. But I can control my perception of your actions, in the same way I can control my perceptions of the inanimate world. To control perceptions of anything, I perform actions (not controlled; the product of all past reorganization), and if my reorganizations have been effective, my perceptual signals come closer to their reference levels. If I hold reference levels for my perceptions of your actions, the same applies. I act, and if the error signal does not decrease, there is a reasonable probability that I will reorganize. You, too. We both reorganize if controlling each other's actions is not proceeding successfully (and if it matters—i.e., if we are working at a reasonably high gain). Our mutual reorganization will probably wind up eventually in a situation where our perceptual errors are not too large. Then, each of us is acting according to the other's standards. This cannot happen if it causes a more-than-compensating increase in errors related to control of percepts outside the interaction. The most likely end-result is that most people in a community use much the same set of standards.

Naturally, the end-result of reorganizing through social interaction and the control of the actions (not the behaviors) of each other will be the existence of perceptual functions in each person that relate to patterns of actions in other people (and perhaps in themselves). Specific reference levels for these perceptions will be associated with the probability of low errors in other perceptual signals, and those reference levels might become the kind of "standards" that Rick was originally talking about.

Ed Ford: Rick says: "I just don't get it, Ed. What does 'teaching standards' have to do with the PCT-based view of human nature?" I am not teaching standards, but the intelligent evaluation and use of the ones people create for themselves. Or, I am trying to help people deal with the standards in the environment in which they find themselves to satisfy their own goals. An example would be helping a person to think through the best way to satisfy the goal of getting released from a lockup facility within the reality of his / her present environment.

When you're down in the trenches, you have to be very practical. People can only achieve their goals by establishing in their own mind criteria (standards, rules, guidelines) upon which they are going to base their decisions. You say the system should be "error-free, " and I agree, and then you go on to say that "this happens by having working (conflict-free) control systems," to which I agree. My question is: how do you help another system get to that point? When you set a reference for driving on the freeway, for establishing a closer relationship with a member of your family, for satisfying an employer or improving your job performance in a working environment, for employing a worker, or just for buying food at a grocery store, you surely do have standards or

criteria based on your references for the choices you're going to make to achieve your goal. And in order for you to function in the environment in which you find yourself in some of the above situations, you are going to have to become aware of what the agreed-to standards or rules are that others in that environment have agreed upon to live by so that you and they can function cooperatively. You can't see a reference for safe driving, but you can see stop signs and speed limits, and you are made aware when you get a driver's license of the various rules or standards for driving.

Obviously, my friend, you have never read my book *Freedom from Stress,* which goes into great detail explaining the relationship of standards, principles, or whatever you want to call them, to the other levels of the hierarchy of control, and how all of that understanding helps people control much more effectively and efficiently for references or goals. You just don't deal exclusively with the highest goal. And, more importantly, the various people with whom I work evidence a need for help in learning how to use their system more efficiently so that they can function more effectively and get what they want.

People are able not only to articulate, prioritize, and evaluate references, but also to set appropriate standards or rules or criteria that will help them reach their goals. Also, these rules or standards will then act as guides for the various choices they have to make *if* they have learned to use their systems properly. What I am trying to say is that you teach people how to use their own systems, to set their own goals, their own standards upon which they can make choices, because PCT teaches me that that is how the system is designed. This hierarchical system is highly interconnected, cross-connected, and interdependent; being able to satisfy goals often demands the awareness and evaluation of all of these various levels. And you know what? It all works.

I never, ever push people to act in a way that would be better for me. That is absolutely wrong. Please explain to me how this focuses my treatment strategy on my goals (except that of helping them to function more effectively and responsibly on their own). Have you ever seen me work with anyone or explain what I do through a role-play demonstration? I suggest that you read the role plays in chapters nine and ten in *Freedom from Stress.* If these people are a part of my life, a necessary part of the environment in which I attempt to live and work cooperatively with others (for example, at work or at home), I have to find out what their goals are, what they are planning to do, how they perceive things, so that I can deal with my life within the reality of the choices these others are making. In my counseling, it is the clients who are asking for help in learning how to deal with their world in such a way that they can satisfy their internal reference signals, including getting along with the people in their lives. It is these living control systems who are asking for help. They are asking to be taught the skills of functioning more efficiently and to learn how to reach their goals with the least hassles. My goal is to help them with what they want. The last thing I want to do is impose my values or beliefs on them. They are going to have to deal with the consequences that are a result of the goals, standards, and choices they make. I teach them how to manipulate themselves, to ride their own bike, to make their own choices, to satisfy their own goals. To manipulate people in such a way as to get them to do what I want is totally against good, sound counseling and teaching, and totally against the PCT design. It is totally repugnant to everything I believe.

Rick Marken: Ed, you say: "People are able not only to articulate, prioritize, and evaluate references, but also to set appropriate standards or rules or criteria that will help them reach their goals." Well, you might be working down there where the rubber meets the road, but you are dealing with some enormously prescient people; apparently, they are able to know what the state of the world (disturbances) will be when they set their standards, rules, and criteria so that these will be appropriate and allow them to reach their goals. How can they do this when the disturbances they will actually encounter are unpredictable and, often, undetectable? I thought that PCT made it clear that the only appropriate settings for *any* references are those that, when the outputs resulting from these reference settings are combined with prevailing disturbances, produce the intended perceptual results. Thus, you *might* be able to direct a person's attention to the perceptual variables that *might* improve his / her ability to *control* other perceptual variables (the ones that he / she came in complaining that he / she could not control), but you cannot *possibly* know in advance the appropriate *settings* for the references for these variables.

And you say: "The last thing I want to do is impose my values or beliefs on them." I never meant to suggest that you did; I know you don't. I am just questioning the idea (at least as you describe it, and as I understand it) that one can help another person control better (which is what I imagine to be the goal of PCT therapy) by suggesting that there might be appropriate settings for one's references for any perceptual variables—rules, standards, principles, whatever. The "appropriate" setting of a reference must vary with circumstance if the intended result is to be produced. So it's not that I think you are trying to impose *your* values—it's that you are suggesting that there are values that are *right* for the client. This is correct, as far as it goes, but the *rightness of* that value is relative; it depends on what they are trying to achieve at a higher level (which I think you clearly understand), *and* it depends on prevailing (and unpredictably changing) circumstances—so that the setting for the value that achieved the higher order goal at one time almost certainly won't do at another time. It is this latter aspect of "setting standards" that I don't hear reflected in your ideas.

Ed Ford: Rick, it is easy to say things in words in the theoretical realm. I wish you would use several examples. It would be much easier for me to understand and to deal with precisely what you are saying. In any event, I will try to respond to what you've said.

Any time we have a goal (reference), and we attempt to achieve this goal, the standards or criteria we set can be set for many reasons, many having to do with other references that interconnect or interrelate to the main reference we have. Whether I am trying to decide on which university to attend, or a young woman to marry, or to drive on a freeway, or to exercise, or to eat "healthier" foods, or where I want to live, or just to call a friend, all of these references are going to involve my making choices which are going to involve other important references. I might set some standards for the kind of woman I want to marry, but in my attempt to satisfy this goal, I might have to adjust my standards if my choices reject me. I might have certain standards for the way a happily married couple should live; obviously, those of us who are (or were) married have found a constant adjusting of standards very necessary to meet the "happy and warm, loving relationship" goal. As a vegetarian, I have very strict standards for what I eat, but I don't try to impose these standards on those with whom I live. Often, when asked to dine at the home of a friend, I am willing adjust some, but not all of my standards (I'll eat some cheese, but never meat or fish). I have been successful at maintaining a no-smoking policy in my house by asking visitors who must smoke to please do it outside.

All of the perceptual variables with which I am trying to deal can be controlled only by satisfying all of the other interconnected references, as well as the one I'm trying to satisfy. Standards can describe in specific terms the kinds of variables you are controlling for; they can also describe the outer limits you are willing to go to to reach or achieve your references, including how much disturbance you are willing to tolerate. Standards can also be tied to other references that are definitely interconnected or interrelated to the present references which you are trying to satisfy.

I am certainly not getting people to articulate "appropriate" settings for references. Rather, they articulate and then evaluate their present settings for their references and see if these particular settings are the most efficient or best settings and the best standards for helping them to reach their goals.

It is the person who has to discover the specifics of his / her conflict and the essential elements within the conflicting area that need to be evaluated, including the references and priorities they've set, the standards and criteria they've established, and the choices they've made—and whether anything in this conflicted area needs to be changed or altered to reduce the conflict. You see, Rick, this is what I've been trying to do. It isn't the counselor who has to discover all of this, it is the person who is having the conflict who has to discover it. All the counselor is doing is helping or teaching the person to better use their system more efficiently and effectively.

You say: "So it's not that I think you are trying to impose *your* values—it's that you are suggesting that there are values that are *right* for the client." No, I am not. I am suggesting that the client find the standards or criteria that work best for him / her in the situation in which he / she finds himself / herself. My job is to teach them how to use their hierarchical systems, as suggested by PCT. When I ask them about their various levels, I am actually teaching them to think "level-wise" and to think about the interconnectedness and the interrelationships involved. From that, they are better able to articulate to themselves (and to me) the specifics of what is going on in their worlds. The more they understand how their living control systems work, the more they are able to use it to their own advantage. Therein lies the beauty of PCT, and especially the levels. When the levels are understood in light of how we function, they become much more useful to us, and our ability to manipulate our own system to our own advantage is enhanced, so that we can satisfy our own internal goals and thus eliminate or reduce conflict to a point where we can live with it.

You say "it depends on what they are trying to achieve at a higher level (which I think you clearly understand), *and* it depends on prevailing (and unpredictably changing) circumstances—so that the setting for the value that achieved the higher order goal at one time almost certainly won't do at another time. It is this latter aspect of 'setting standards' that I don't hear reflected in your ideas." I have nowhere suggested that once someone articulates their individual standards to me, they are locked into those standards. It is the ability of people to recognize and utilize these levels to their advantage to deal with their conflicts that is important. We all change standards all of the time. It is important that they first recognize the existence of the standards, the part they play in how we think, their usefulness in setting and achieving references through the choices they make. I don't care whether they change their standards or not. We all change standards all of the time. It's being able to change within the context of avoiding or reducing conflict that is critical.

When counseling (read teaching) others, it's not what I think, it's what *they* think—my job is to teach them to think by helping them to build confidence in their thinking ability. When they learn PCT and what goes into making up a living control system, they have the road map. My job is to teach them how to use it. They have to learn to use it when I'm not around.

Rick Marken: Ed, your last post on standards cleared up a lot. I know that what you do is teach people to control their own lives more effectively; sometimes, I take issue with the way you describe some of your therapeutic goals. But your last description was excellent and quite compatible with my own sentiments about therapy (and they are just sentiments, since I would never be able to actually *do* therapy as skillfully as you do it).

Ed Ford: I've always believed that there should be no conflict between science and religion. I've recently found evidence of this. My granddaughter, Ruth, age five, from California, was visiting Hester and me. Hester had taken Ruth and her first cousin, Sally Ann, age four, who lives here in Phoenix, to a Christmas tree display, and on the return trip, the two children were in the back of the car, talking. The conversation went as follows:

Sally Ann: My immune system takes care of me.

Ruth: Well, my guardian angel takes care of me all the time.

Sally Ann: All the time?

Ruth: Yes, all the time. She's always with me, everywhere.

Sally Ann: Well, if you just leave the body alone it will take care of itself.

Ruth: Well, my guardian angel takes care of me all the time.

Sally Ann: Well, my immune system takes care of me.

They then went on to another subject.

Greg Williams: Ed, regardless of the potential and (I believe) actual conflicts between science and certain religious ideas, it appears that the major problem is religion vs. religion.

Rick Marken: I think Ed's young relatives were having a religious dispute—no science involved at all. Using scientific terms (like immune system) to describe the cause of perceptions (health) doesn't make it science. "Science" and "religion" are words that refer to lots of different perceptual variables. For me, the best definition of science was given by Bill Powers: "disciplined imagination"; we invent models (imagination) and then test to see if we observe in perception what the model does when "switched on" (discipline). This is a nice definition because it makes it easy to juxtapose it to what I think of as the essence of religion: "faithful imagination." The crux of the difference is the way you ultimately test whether your imaginings are "right." In science, the final arbiter is God—i.e., the cause of one's perceptual experience (we call her Boss Reality). In religion, the final arbiter is People—perceptions are made to fit the faith (too often, violently). (I should note that, by this definition, much that is called "science" is not. Lysenkoism in the USSR is an example of religion—faith in inheritance of acquired characteristics—posing as science.) To my knowledge, there is no religion that would qualify, by this definition, as a science.

Everybody seems to be making up different stories about god(s) and what they say about the meaning of life and how we should behave in it. Seems like what we've got here are variable means to achieve a higher order result—varying across people, anyway. Wouldn't it be marvelous if we could learn to vary these means within one person—ourselves? Then a "Serb" could see that s/he is "Bosnian," too, and vice versa; an Israeli could see that s/he is Palestinian, a Catholic could see that s/he is Lutheran, an Atheist could see that s/he is Muslim, etc. The solution to the problem of religion (like the solution to any conflict resulting from inflexible goals) is not to eliminate the goal but to *rise* above it; PCT can help people get their consciousness to the level that is served by controlling religious perceptions. Once you get up there, you will see that religious goals are arbitrary—but useful for satisfying the needs of that higher level. When you get up there, you see that choosing a religion, ethnicity, nationality, etc., is just as useful (and arbitrary) as choosing a nice book to settle in with on a rainy day; sometimes you want a romance and sometimes only a thriller will do.

Dag Forssell: Rick says: "The crux of the difference is the way you ultimately test whether your imaginings are 'right.' In science, the final arbiter is God—i.e., the cause of one's perceptual experience (we call her Boss Reality). In religion, the final arbiter is People..." What a marvelous, lucid insight. And people can create and defend any system-concept religion they want, teach it, fight for it, and die for it. Witness the sorry spectacle in India. No Boss Reality arbiter there.

Rick also says: "To my knowledge, there is no religion that would qualify, by this definition, as a science." Some years ago, I attended Religious Science, Science of Mind (several times). They would take a text from the Bible, another from the Koran, a third from some Buddhist book. They suggested that there have been many good teachers, but that none is a God any more or less than you and 1. In every affirmation, song, and message, I was able to substitute the word God with "laws of nature." The one thing that was supernatural was "treatment." So I guess they fall down like all of the others.

Religion is more than a system concept, though. It is also a social club. There is where much of the strength and value comes from. And the coercion. If you don't say you believe in what we say we believe in, you can't play in our sandbox. You might get ostracized from your family, friends, and community. Better go to church on Sunday.

Ed Ford: The purpose of my post about Ruth and Sally Ann was to share with friends a delightful and amusing interchange between two innocent children. What I enjoyed most about their conversation was how different their perceptions were and how they tolerated that difference. It was meant to be light and amusing and not a serious comment on or about religion.

Years ago, the foreword to a movie with a religious theme read as follows: "To those who believe, no explanation is necessary; to those who don't, no explanation is possible."

Rick says: "The solution to the problem of religion (like the solution to any conflict resulting from inflexible goals) is not to eliminate the goal but to *rise above* it; PCT can help people get their consciousness to the level that is served by controlling religious perceptions." It depends on whether, for a particular individual, his/her religion presents a conflict. For me, it doesn't. It is highly compatible with everything else that goes on in my head, including and especially PCT. Secondly, I see it at my highest level. It does give me satisfaction at the very highest (in terms of priorities) system concept I have, which, for me, is to be one with my Maker. Obviously, there are standards that flow from that system, and choices I make based on those standards.

And Rick says: "Once you get up there, you will see that religious goals are arbitrary—but useful for satisfying the needs of that higher level." I think it would depend on an individual's own perceptions of when he or she was up there, and how he or she perceived his or her "religious goals" and the other needs at the higher level. It is interesting when someone comes on the net and claims to understand PCT, and then says things that are obviously different from what we all have experienced through our own individual work. Even among those in the CSG, we all understand PCT according to how we have created it in our perceptual systems, from what we have done, read, observed around us, perceived as useful, experienced in creating ideas from it, perceived from building models based on it, etc. Many of us have understandings that others will never have. My wife's understanding of what it is like to have a child is quite different from mine, and, obviously, I'll never understand her experiential knowledge.

Rick also says: "When you get up there, you see that choosing a religion, ethnicity, nationality, etc., is just as useful (and arbitrary) as choosing a nice book to settle in with on a rainy day; sometimes you want a romance and sometimes only a thriller will do." Again, "you see" refers to what *you* perceive, not what *everyone will* perceive "when you get up there." I think that you presume a lot when you state that you understand everyone's knowledge of how they will perceive an experience you've only had in terms of your own individual perception of your own created experience. Your knowledge of religion is limited to what you presently perceive, just like my knowledge of PCT (or anything else) is limited to only what I have built into my own perceptual system.

Dag says: "Some years ago, I attended Religious Science... They suggested that there have been many good teachers, but that none is a God any more or less than you and L" That was their perception, and you accepted that as yours. And I respect that. However, I don't happen to agree with that statement. That's my perception.

And Dag says: 'Religion is more than a system concept, though. It is also a social club." Again, that is your perception. And again, I don't happen to agree with that statement. My own particular religion is based on fact, not fiction. It is also based on 50 years of thought, study, research, and lots of reading.

Rick Marken: Ed quotes: "To those who believe, no explanation is necessary; to those who don't, no explanation is possible." Apparently, that's true. What I want to understand is *why* it is true. I want an explanation of believing, itself, whatever the beliefs themselves might be.

Ed says: "It depends on whether, for a particular individual, his / her religion presents a conflict." I didn't mean that religion is a problem because it creates intrapersonal conflict. I'm sure most devout people are quite unconflicted about their religious beliefs. The problem with religion (and other high-level goals of the same sort that become fixed—ethnicities, nationalities, etc.) is *interpersonal conflict. I* don't know if you've looked at your local newspaper lately, but mine is *filled* with violent, interpersonal conflicts over religions, nationality, ethnicity, etc. People are fighting their brains out to defend perfectly arbitrary goals; I consider this a problem-one that is so unnecessary that it is unbelievable. And the solution, of course, is for each person to be able to see that their own ethnic, religious, or national goals, though important to themselves, are perfectly arbitrary; that it's like arguing over whether cars should be driven on the left or right.

To me, religion is (as I have said before) just something that people do—like being a control theorist. PCT is trying to understand *all* of human behavior, and religion is certainly one of the most important (and troublesome) things that people do. It should be something we in PCT try desperately to understand.

Bill Powers: It will not be possible for science and religion to get together until both realize that neither is Revealed Truth, and that both are human ideas. Of course, that is precisely what both sides have been rejecting since the start of science. One side points the finger at Nature, the other at God. Neither side, apparently, notices whose finger is doing the pointing.

Greg Williams: I'd like to add one additional observation to Bill's post regarding the possibility of the religious and scientific "sides" getting together. In several forms of religion, and some (at least historical) forms of science, accepting authority and having faith have been / are now valued more (sometimes *much* more) than adjusting beliefs in the light of new evidence. Modern science at least gives lip service to the idea that one's *own* finger should be doing the pointing, unencumbered by pleas or threats from others. But that is anathema to some modern religions. One reason that a discussion of "science vs. religion" is appropriate on CSGnet, in my opinion, is that the issue of self- vs. otherdetermination is right at the heart of what control theory has to say about the chances of an individual successfully coping in a disturbance-filled world. On the other hand, high-level reference signals (within a broad spectrum) appear to be very loosely coupled to day-to-day survival (assuming you aren't in a holy war, of course), so I don't feel much missionary zeal for going around begging folks to recant what they accept on authority. And if I did, I wouldn't rail against the beliefs themselves so much as against why they are held. As a general principle (based on PCT ideas), it would appear that breaking correcting loops (e.g., accepting dogma uncritically) is dysfunctional. Yet people do it all the time and seem none the worse for it. Of course, their neighbors might be *much* worse for it!

May your neighbors not be *extremely* dogmatic.

Rick Marken: I partly disagree with Bill's post on religion and science getting together. I think there are many scientists (the good ones) who understand that their models are human ideas, and that "nature"—the cauldron in which these ideas are tested—is just their own perceptions. I think there are also religionists who understand that religious models (myths) are human ideas, and that "spiritual experience"—the cauldron in which these ideas are tested—is just their own perceptions (human experience). I would venture to guess that there are far more scientists like the above than there are religionists. The reason for this is that implicit (or explicit) in most religions is the idea that you *must* believe that these ideas are Revealed Truth or else you, your people, or the human race are in deep trouble. I don't think this latter assumption is part of science—although I agree that many scientists act as *though* such an idea were part of the game; that is where science and religion become one—as Bill says, when their ideas (models, myths) are treated as revealed truth rather than human invention—invented for a *purpose*.

Bill Powers: There are two sides to religion. One of them, the good side, consists of the attempt to adopt and live out principles that make civilization possible. As most people never think about such things except in the context of a religion, one wonders what the world would be like without such formalized social systems of belief.

The bad side shows up because people have different religions. If those living under principles of love and tolerance could actually live up to those principles, all would be well. But aside from the fact that not all religions preach universal brotherhood, it doesn't seem possible for people to live up to their religious principles when those principles disagree with someone else's.

The basic reason, I think, is the assumption of supernatural origin of the religious principles. When you believe that you are in receipt of the word of God, directly or through an authorized dealer, there can be no tolerance for deviations. The word of God is absolute. This means that if a different group claims to have heard a different word, or a different interpretation of words, the other group must simply be wrong. Every religious group must feel this way about every other group, no matter what they say. Very quickly, this comes down to the choice of converting the other group to the true belief ("saving" them), isolating the other group, or eliminating the other group.

Each group, of course, must resist all attempts by the other groups to evangelize, because succumbing would be going against the word of God. The loop gain, with respect to adhering to the word of the Infinite, must be infinite. This means that even minor differences of doctrine can lead to maximum conflict.

All that saves us from continuous violent confrontation between religions is that very few people are actually as religious as they think they are, or claim to be.

The greatest mystery of the human mind, in my view, is the phenomenon of Belief. Nazis are easy to deal with, because their beliefs are threatening to our physical safety, and we can flatly reject them. But what about other belief systems, invented and accepted apparently at random? Is the human mind just naturally susceptible to any belief that comes around, no matter how childish and full of holes? Is there something about our highest levels of organization that demands some belief, any belief, to fill the vacuum?

It seems to me that before we can have anything approaching sanity on our planet, we must begin to understand how belief systems get formed and how to keep them from overpowering people—how to leave a little freedom of belief, so that knowledge about the *whole* world of experience can play a part in forming belief systems. I haven't the slightest idea of how to do that, except by continuing to point out that different people believe different things, a fact that ought to give anyone pause who is convinced that his / her own belief system is the only right one.

Or is this a level at which we are all helpless, including me?

Rick Marken: Bill says: "The greatest mystery of the human mind, in my view, is the phenomenon of Belief." I agree. We should explore this from a PCT perspective. The problem, of course, is that, when it comes to many of one's own beliefs, they are not treated as beliefs, but as knowledge. I think many of our most tenacious intrapersonal and *interpersonal conflicts* are the result of controlling perceptions based more on beliefs (replayed reference signals) than on Boss Reality.

I think it would be worthwhile to say what beliefs are in the context of the PCT model—to describe examples of the everyday beliefs that people are walking around with (from the divine, like religious beliefs, to the profane, like beliefs about the "right" foods to eat); also, it would be nice to discuss the difference (from a PCT perspective) between belief and knowledge. I know this is a difficult discussion to have, precisely because beliefs are so important to people. With Bill, I ask, "*Why* is this so?" Why do people fight to prove that what they do not know is so? There must be a reason that this species has been willing to persecute itself for millennia over fantasies. It must be an aspect of our nature as control systems. What is it? I think that this could be a very satisfying (and even therapeutic) investigation.

Bill asks: "Or is this a level at which we are all helpless, including me?" No. I think people, like you (and me?), who are willing to consider the possibility that *anything* we think could be just a belief and, more importantly, are willing to wonder what a belief is, are not helpless victims of our beliefs (at least, when we are able to keep our awareness "above" the levels that create those beliefs—something that I don't do nearly as often as I would like). I think it requires some effort to defeat some of the insidious consequences of belief, but it can be done, I think.

Ed Ford: It seems to me that belief systems are formed by living control systems as they try to establish harmony within themselves as a result of their attempts to find satisfying experiences from the environment in which they find themselves. The choices we make and the standards we've set ultimately evolve into systems of ideas, or the way we think things ought to be. I think this harmony, this internal peace or internal integrity, is what the living control system is continually striving toward. Obviously, our knowledge of what's available is limited by our perception of the environment in which we find ourselves, plus what becomes available to us through reorganization. What we create out of what we perceive is what ultimately becomes what we are.

I think humans tend to accept the system concepts of those who they perceive love them and whom they love or admire. Obviously, if there are internal peace and harmony where we live, then the prevailing system concepts of our parents / friends are most likely to be perceived as acceptable. Those systems are ultimately tested when children (and later adults) are faced with choices which are in conflict with the prevailing or accepted system concepts. But to me, the ultimate test of a system concept is that first it brings internal harmony or peace to the person.

I don't believe a belief or value system (system-concept level) overpowers a person. I believe many people choose systems and elements of those systems and create their own standards from how they perceive those systems to justify the choices they're making, in their attempts to find that elusive peace and harmony that all living control systems are trying to establish.

When a person harms another living control system, his / her system concept is brought into disrepute. And this shouldn't be. I don't think it's right to blame Christianity for the acts of those who, claiming to be Christians, do harm to others, any more than it's fair to blame any system of ideas on those who claim to be adherents, but who go about harming others.

The second important test of any system concept is the respect shown to those "who don't belong, who don't believe." Therein lies the critical test of any system of beliefs, namely, that everyone is shown respect, as having value as a person. That, to me, is the real test of a valid system of beliefs. If from a system concept I am able to establish standards and make choices that bring me the internal peace and harmony within my system *and* at the same time that system concept leads me to see value in others and respect their right to make choices, then the system has value. In short, when we harm others, we harm ourselves, and in the process the very harmony and peace we are seeking are lost.

When a person is in conflict and uses a system concept to justify actions which bring harm to others, I don't think the system concept is wrong, I think the person is wrong. And I don't think the belief system overpowered them, they merely used the system "to justify their own means." I think people tend to overpower themselves by setting impossible standards or goals, by trying to change things over which they have no control, or by making ineffective choices in a desperate attempt to bring harmony or peace to their system.

Because I'm a living control system by design, my system concepts are very unique to me. No one quite perceives things the way I do. And I think the test for whether our systems of beliefs are valid are our own internal harmony and peace, and the respect and value we assign to others.

Rick Marken: Ed says: "I don't believe a belief or value system (system-concept level) overpowers a person." A belief and a system concept are not the same thing. A belief in PCT (I think) is an imagined perception: this means that beliefs can occur at any level of the hierarchy (except for the lowest); we can believe that the sky is blue (sensation), that it will rain (fluid transitions?), that we're loved (relationship), etc. We can also have beliefs that are system-concept-level perceptions —I can believe that I am a control theorist.

Beliefs (by my definition) can also differ in terms of one's ability to produce or experience them as perceptions (rather than just as imaginations). I believe my car is in the lot, and I can produce that perception; I believe that Mozart was the means by which God spoke to humanity, but I can't produce that perception (I can certainly produce the imagination).

Our ability to "believe" is, I think, one of the things that makes life fun; it makes it possible to be entertained by stories, plays, and such. I think it also makes life a bit more tolerable (as Ed said, it helps us "find that elusive peace and harmony that all living control systems are trying to establish"). It does this by "filling in" the unachieved aspects of the perceptions we are controlling; we believe that we are "loved," for example—and we create a perception that is based mostly on Boss Reality but that is "filled in" a bit by belief (imagination) so that our control seems a bit better than it might actually be.

But you can see that what is good about belief is what could also make it a problem; belief makes stories fun because we treat the imaginations as though they were "real" perceptions; but what happens when we forget that they *are not and never were real perceptions?* We get what we see—people willing to die or kill to control for imagined perceptions.

I think it is interesting that when the "filling in" done by belief gets to be a bigger part of perception than the part constrained by Boss Reality, we call that "insanity." But when the "filling in" is *total*—so that there is no constraint of Boss Reality—just belief based on made-up stories (the Bible, the Koran, etc.), we (some of us) call that "wisdom" I suggest that we call it what it is: "total insanity."

The Blind Men and the Elephant: Three Perspectives on the Phenomenon of Control

Richard S. Marken
10459 Holman Avenue, Los Angeles, CA 90024

Abstract

Behavior has been described as a response to stimulation, an output controlled by reinforcement contingencies, and an observable result of cognitive processes. It seems as if these are descriptions of three different phenomena, but they are actually descriptions of three different aspects of the same phenomenon: control. Control is like the proverbial elephant studied by the three blind men; what one concludes about it, and how one tries to explain it, depends on where one stands. I suggest that the best place to stand is where one has a view of the *whole* phenomenon, be it elephant or control.

Introduction

The behavior of living organisms (and some artifacts) is characterized by the production of consistent results in an unpredictably changing environment, a phenomenon known as control (Marken, 1988). Control can be as simple as maintaining one's balance on uneven terrain, or as complex as maintaining one's self-esteem in a dysfunctional family. Control is a pervasive aspect of all behavior, yet it has gone virtually unnoticed in psychology. What *has* been noticed is that behavior appears to be a response to stimulation, an output controlled by reinforcement contingencies, or an observable result of cognitive processes. Each of these appearances is what would be expected if people were looking at control from different perspectives. The situation is similar to that of the three blind men who were asked to describe an elephant. The one near the tail described the elephant as a snake, the one near the leg described it as a tree trunk, and the one near the side described it as a wall. Each description gives an accurate picture of some aspects of the elephant, but a false picture of the elephant as a whole. If behavior involves control, then psychology has given an accurate picture of some aspects of behavior, but a false picture of behavior as a whole.

The basic requirement for control is that an organism be in a negative-feedback situation with respect to its environment. A negative-feedback situation exists when an organism's response to sensory input reduces the tendency of that input to elicit further responding. Negative feedback implies a closed-loop relationship between organism and environment; sensory input causes responding that influences the sensory cause of that responding, as shown in Fig. 1. It is hard to imagine an organism that does not exist in such a closed-loop situation, because all organisms are built in such a way that what they do affects what they sense. Eyes, for example, are located on a head that moves, so that what the eyes see depends on what the head does. To the extent that what the head does depends on what the eyes see (such as when the head turns in response to an attractive passerby), there is a closed loop; sensory input causes responding (head movement), which affects the cause of responding (sensory input). The feedback in this loop must be negative, because behavior is typically stable (organisms do not normally exhibit the "runaway" behavior that characterizes positive-feedback loops, such as the feedback from a microphone that amplifies its own output).

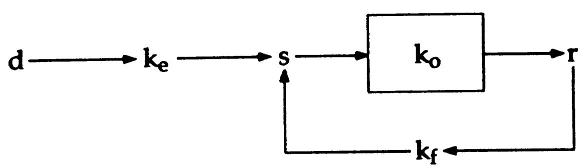

Figure 1. Closed-loop feedback relationship between an organism, represented by the rectangle, and its environment, represented by the arrows outside of the rectangle. A sensory variable, s, influences responding, r, via the organism function, k_o. Responding influences the sensory variable via the feedback function, k_f. The sensory variable is also influenced by an environmental variable, d, via the environmental function, k_e.

The fact that organisms exist in closed negative-feedback loops means that two simultaneous equations are needed to describe their relationship to the environment. These are given as equations (1) and (2) below. The terms in these equations are summarized here for reference in the discussion that follows:

s = sensory variable;
r = response variable;
s* = reference value for sensory variable, such that r = 0 when s = s*
d = environmental variable;
k_o = organism function relating sensory variable to response variable;
k_e = environmental function relating environmental variable to sensory variable; and
k_f = feedback function relating response variable to sensory variable.

For simplicity, I assume that all functions are linear, and that all variables are measured in the same units.

Equation (1) describes the effect of sensory input on responding, so that

(1) $r = k_o(s^* - s)$.

This equation says that responding, r, is a linear function of sensory input, s. The sensory input is expressed as a deviation from the value of input, s*, that produces no responding; s* defines the zero point of the sensory input. Equation (2) describes the effect of responding on sensory input. For simplicity, I assume that responding, r, adds to the effect of the environment, d, so that:

(2) $s = k_f r + k_e d$.

The variables r and d have independent (additive) effects on the sensory input, s. The nature of the environmental effect on sensory input is determined by the environmental function, k_e. The feedback effect of responding on the sensory cause of that responding is determined by the feedback function, k_f.

Equations (1) and (2) must be solved as a simultaneous pair in order to determine the relationship between stimulus and response variables in the closed loop (see Appendix, below). The result is

(3) $r = (1 / ((1/k_o) + k_f))s^* - (k_e / ((1/k_o) + k_f))d$.

Equation (3) can be simplified by noting that the organism function, k_o, transforms a small amount of sensory energy into a large amount of response energy (such as when a pattern of light on the retina is transformed into the forces that move the head). In control engineering, k_o is called "system amplification" or "gain," which can be quite a large number.

With sufficient amplification (such that k_o approaches infinity), the $1/k_o$ terms in equation (3) approach zero, so equation (3) reduces to

$(4)\ r = s^* / k_f - (k_e / k_f)d.$

Equation (4) is an input-output equation that describes the relationship between environmental (stimulus) and response variables when an organism is in a closed-loop, negative-feedback situation with respect to its environment. The result of being in such a situation is that the organism acts to keep its sensory input equal to s*, which is called the reference value of the input. Equation (4) shows that the organism does this by varying responses, r, to compensate for variations in the environment, d, that would tend to move sensory input away from the reference value; this process is called control.

Three Views of Control

All variables in equation (4), with the possible exception of s*, are readily observable when an organism is engaged in the process of control. The environmental variable, d, is seen as a stimulus, such as a light or sound. The response variable, r, is some measurable result of an organism's actions, such as bar pressing or speaking. The reference value for sensory input, s*, is difficult to detect because an observer cannot see what an organism is sensing. But s* is the central feature of control, since everything an organism does is aimed at keeping its sensory inputs at reference values. The value of s* can be constant or variable, its value at any instant being determined by properties of the organism itself.

Because reference values are difficult to detect, it will not be obvious to an observer that an organism is engaged in the process of control. What will be obvious is that certain variables, particularly the environmental and response variables and the relationship between them, will behave as described by equation (4). Thus, equation (4) can be used to show how control appears to someone who does not know that it is occurring. It turns out that there are three dearly different ways of looking at control, depending on which aspect of the behavior described by equation (4) one attends to.

The stimulus-response view. This view of control sees behavior as a direct or indirect result of input stimulation. An example of stimulus-response behavior is the so-called "pupillary reflex," where changes in a stimulus variable (illumination level) lead to changes in a response variable (pupil size). The stimulus-response view is the basis of several current approaches to understanding behavior, such as the "synergistic" or "coordinative structure" theory of motor coordination. Warren, Young, and Lee (1986), for example, describe a synergistic model of running in which "vertical impulse is directly modulated by the optical variable..." (p. 264). The behavior of running is seen in stimulus-response terms; an optical stimulus variable determines ("modulates") the value of a response variable (vertical impulse). The stimulus-response view is also the basis of a recent theory of attention (Cohen, Dunbar, and McClelland, 1991) in which connections between printed-word stimuli and verbal responses in the Stroop effect are modulated by connections in a neural network.

Equation (4) shows that behavior will look like a stimulus-response process when the reference value for sensory input, s*, is a constant. If s* is zero, then responding is related to environmental stimuli as follows:

$(5)\ r = -(k_e / k_f)d.$

Equation (5) shows that, when there is a fixed reference level for sensory input, it will look to an observer of behavior as though variations in an environmental stimulus, d, cause variations in a response, r. This is what one sees in the pupillary reflex, where pupil size, r, is proportional to illumination level, d. Of course, this relationship between pupil size and illumination level is precisely what is required to keep a sensory variable (sensed illumination) at a fixed reference value (s* = constant).

When looking at an apparent relationship between stimulus and response, one's inclination is to assume that the nature of that relationship depends on characteristics of the organism. Equation (5) shows, however, that when an organism is engaged in control, this relationship depends only on characteristics of the environment (the functions k_e and k_f); the organism function relating sensory input to response output, k_o, is rendered completely invisible by the negative-feedback loop. This characteristic of the process of control has been called the "behavioral illusion" (Powers, 1978).

The reinforcement view. This view of control sees behavior as an output that is shaped by contingencies of reinforcement. A reinforcement contingency is a rule that relates outputs (like bar presses) to inputs (reinforcements); in equation (4), this contingency is represented by the feedback function that relates responses to sensory inputs, k_f. The reinforcement view is the basis of at least one influential theory of generalization and discrimination (Shepard, 1987). In a connectionist implementation of the theory, a reinforcement contingency is used to shape the formation of generalization gradients (Shepard, 1990). The reinforcement view is also the basis of modern theories of operant behavior. According to Domjan (1987), the contemporary perspective on operant behavior focuses on how contingencies "restrict freedom of action and... create redistributions of various types of activities" (p. 562). In other words, contingencies shape (redistribute) responses (activities).

Equation (4) shows that it will look as though contingencies (the feedback function) control responses when s*, d, and k_e are constants, as they are in typical operant conditioning experiments. In these experiments, s* is the organism's reference value for the sensory effects of the reinforcement. The environmental variable, d, is the reinforcement, which, if it is food, is typically a constant size and weight. The sensory effect of a reinforcement can be assumed to be directly proportional to its size and weight, making k_e -1. So, equation (4) can be re-written as

(6) $r = S^* / k_f - D / k_f$,

where S* is the constant reference value for sensed reinforcement, and D is the constant value of the reinforcement itself.

The only variable in equation (6) is the feedback function, k_f, which defines the contingencies of reinforcement. One simple contingency is called the "ratio schedule," in which the organism receives a reinforcement only after a certain number of responses. The ratio corresponds to the function k_f in equation (6). When the ratio is not too demanding, it is found that increases in the ratio lead to increased responding. More demanding ratios produce the opposite result; increases in the ratio lead to decreased responding (Staddon, 1979). Either of these results can be produced by manipulating the relative values of S* and D in equation (6). The important point, however, is that the apparent dependence of responding on the feedback function, k_f, is predicted by equation (6). To an observer, it will look like behavior (responding) is controlled by contingencies of reinforcement. In fact, the relationship between behavior and reinforcement contingencies exists because the organism is controlling sensed reinforcement; responding varies appropriately to compensate for changes in the reinforcement contingency, so that sensed reinforcement is kept at a constant reference value, S*.

The cognitive view. This view of control sees behavior as a reflection or result of mental plans or programs. This kind of behavior is seen when people produce complex responses (such as spoken sentences, clever chess moves, or canny investment decisions) apparently spontaneously; there is often no visible stimulus or reinforcement contingency that can be seen as the cause of this behavior. The cognitive view is the basis of numerous psychological theories that propose mental algorithms to explain the appearance of cognitive behavior. Examples of such theories are the ACT (Anderson, 1983) and SOAR (Newell, 1990) models of cognition, and hierarchical models of the generation of movement sequences (Rosenbaum, Kerry, & Derr, 1983).

Cognitive behavior is most obvious when environmental factors (such as stimulus variables and environmental and feedback functions) are

held constant. When this is the case, equation (4) becomes

(7) $r = s^* / F + K$,

where F is the constant feedback function, and $K = (k_e / k_f)(d)$, a constant. Since s* is typically invisible, equation (7) shows that there will appear to be no obvious environmental correlate of cognitive behavior. An observer is likely to conclude that variations in r are the result of mental processes-and, indeed, they are. But it is actually variations in s*, not r, that are caused by these processes; variations in r are the means used to get sensory inputs equal to s*. Thus, chess moves are made to keep some sensed aspect of the game at its reference value. When the environment is constant, r (the moves) might be a fair reflection of changes in the reference value for sensory input. However, under normal circumstances, r is only indirectly related to s*, variations in r being mainly used to compensate for variations in the environment that would tend to move sensory input from the reference value, s*.

Looking at the Whole Elephant

The blind men never got a chance to see the whole elephant, but if they had, they would have instantly understood why it seemed like a snake to one, a tree trunk to the second, and a wall to the third. Psychologists, however, can take a look at control and see why the appearance of behavior differs, depending on one's perspective. What is common to the three views of behavior discussed in this paper is the reference for the value of sensory input, s*. Organisms behave in order to keep sensory inputs at these reference values (Powers, 1989). They respond to stimulation in order to keep the sensory consequences of this stimulation from moving away from the reference value, so it appears that stimuli cause responses. They adjust to changes in reinforcement contingencies by responding as needed in order to keep the sensory consequences of reinforcement at the reference value, so it appears that contingencies control responding. And they change their responding in order to make sensory input track a changing reference value for that input, so it appears that responding is spontaneous.

What appear to be three very different ways of describing behavior can now be seen as legitimate ways of describing different aspects of one phenomenon—control. Each is just a different way of describing what an organism must do to keep its sensory inputs at their reference values. Indeed, once one understands that the appearances called "behavior" are the visible consequences of an organism's efforts to control its sensory inputs, the problem of explaining behavior changes completely, from an attempt to build models that simulate the appearance of behavior (stimulus-response, reinforcement, or cognitive) to an attempt to build models that control the same sensory inputs as those controlled by real organisms. In order to build the latter type of model, it is necessary to learn what sensory variables are actually being controlled by organisms. This type of investigation cannot be done by simply looking at the appearance of behavior. Methods based on control theory can be used to test which sensory variables an organism might be controlling at any time (Marken, 1992). These methods make it possible to take off the blindfolds and see the whole elephant—the phenomenon of control.

Appendix

Given the two system equations

(1) $r = k_o(s^* - s)$ and (2) $s = k_f r + k_e d$,

we want to solve for r as a function of s. First, substitute equation (2) for s in equation (1) to obtain

(A.1) $r = k_o(s^* - (k_f r + k_e d))$,

which expands to

(A.2) $r = k_o s^* - k_o k_f r - k_o k_e d$.

Move all terms with r to the left side of equation (A.2) to obtain

(A.3) $r + k_o k_f r = k_o s^* - k_o k_e d$.

Factor out r on the left side of equation (A.3) to obtain

(A.4) $r(1 + k_o k_f) = k_o s^* - k_o k_e d$.

Divide both sides of equation (A.4) by $(1 + k_o k_f)$ to obtain

(A.5) $r = (k_o / (1 + k_o k_f))s^* - (k_o k_e / (1 + k_o k_f))d$.

Finally, divide k_o out of the numerators on the right side of (A.5) to get equation (3):

(3) $r = (1 / ((1 / k_o) + k_f))s^* - (k / ((l / k_o) + k_f))d$.

Acknowledgments

I would like to acknowledge the earlier ideas of Wayne A. Hershberger and W. Thomas Bourbon on this topic, which provided inspiration for the integrative analysis presented in this paper.

References

Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.

Cohen, J. D., Dunbar, K., & McClelland, J. L. (1991). On the control of automatic processes: A parallel distributed processing account of the Stroop effect. *Psychological Review*, 97, 332-361.

Domjan, M. (1987). Animal learning comes of age. *American Psychologist*, 42, 556-564.

Marken, R. S. (1988). The nature of behavior: Control as fact and theory. *Behavioral Science*, 33,196-206.

Marken, R. S. (1992). *Mind readings: Experimental studies of purpose*. Gravel Switch, KY: Control Systems Group.

Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.

Powers, W. T. (1978). Quantitative analysis of purposive systems: Some spadework at the foundations of scientific psychology. *Psychological Review*, 85, 417-435.

Powers, W. T. (1989). *Living control systems: Selected papers of William T. Powers*. Gravel Switch, KY: Control Systems Group.

Rosenbaum, D. A., Kerry, S., & Derr, M. A. (1983). Hierarchical control of rapid movement sequences. *Journal of Experimental Psychology: Human Perception and Performance*, 9, 86-102.

Shepard, R N. (1987). Toward a universal law of generalization for psychological science. *Science*, 237, 1317-1324.

Shepard, R. N. (1990). Neural nets for generalization and classification: Comment on Staddon and Reid (1990). *Psychological Review*, 97, 579-580.

Staddon, J. E. R. (1979). Operant behavior as adaptation to constraint. *Journal of Experimental Psychology: General*, 108, 48-67.

Warren, W. H. Jr., Young, D. S., & Lee, D. N. (1986). Visual control of step length during running over irregular terrain. *Journal of Experimental Psychology: Human Perception and Performance, 12*, 259-266.

Models and Their Worlds

W. Thomas Bourbon
(Department of Neurosurgery, The University of Texas Medical School-Houston, 6431 Fannin, Suite 148, Houston, TX 77030)

William T. Powers
(73 Ridge Place, CR 510, Durango, CO 81301)

Abstract

Many seemingly plausible models of behavior demand implausible models of the physical world in which behavior occurs. We used quantitative simulations of a person's performance on a simple task to compare the models of causality and of how the world works in three theories of behavior: stimulus-response, cognitive, and control-theoretic. Our results demonstrate that if organisms in fact functioned like the first two models, they could survive only in implausibly stable worlds; if like the third, they could survive in a changeable world. Organisms inhabit a changeable world that does not satisfy the demands of popular behavioral theories. For the sciences of behavior, the implications are clear: either cling to theories that do not mesh with knowledge of how the world works, or abandon many cherished notions about how and why behavior happens in favor of models that deal adequately with change.

Models and Their Worlds

The question usually addressed by behavioral theorists is "Why do organisms behave the way they do?" One group answers "Because the world outside them is the way it is"; another group answers "Because the minds or brains inside them are the way they are." In either case, behavior is at the end of a linear sequence of cause and effect, a consequence of antecedent stimuli from the environment or antecedent commands from the mind or brain. As an alternative, one can propose that organisms behave to control what happens to them. In the process, their actions affect the world outside of them. 'Why is the world the way it is? Partly because organisms behave the way they do."

"The world" is the part of the surroundings on which an organism can act, and which, in turn, affects the organism. Every statement about the antecedents or consequences of behavior either includes or implies notions about how the world operates. Every theory of behavior is, in part, a theory about the world in which behavior occurs.

In this paper, we reduce three models of behavior to elemental form to identify and test their ideas about causality. Two models represent core assumptions in most popular theories; the third is the model from perceptual control theory (PCT). We require each model to simulate and predict the same behavioral events that occur when a person performs a simple task, but we go a step further. For each model, we determine whether its implications about how the world and behavior affect one another are reasonable and true to what is known about the physical world.

Three Models

For convenience, we call the two popular models the "stimulus-response" (S-R) model and the "cognitive" model. Our simple versions of these models are not intended to represent, in detail, any specific variations on those two themes, but we believe they faithfully represent core assumptions about causality embraced in those themes. Our method of testing requires that each model predict moment-by-moment values of several continuous environmental variables, a challenge to which behavioristic and cognitive models are rarely subjected; hence, simple computational versions of those models are not readily available, and we constructed our own. Anyone who rejects our versions of those theories should identify acceptable versions and then require their models to duplicate the quantitative results we report here.

The stimulus-response model. Our S-R model represents all theories that say external influences determine behavior. Such models sometimes (but by no means always) recognize that motor actions produce environmental consequences, but all insist that action is a dependent variable. A behavioral episode begins with an independent antecedent (stimulus, context, event, occasion, relationship, or treatment), followed (in some theories) by an effect on the organism, then (in all theories) a behavior as a dependent variable, and finally the consequences of that behavior. Environmental consequences of action simply follow from what the environment did to the organism; if any consequences of action modify subsequent influences on the organism, that is merely another change in the independent variable, followed in a lineal causal chain by another action and another consequence.

We expect most behaviorists to say that our S-R model is "reflexological"—a version of behavioristic theory many behaviorists disavowed years ago—and to echo the comment: "There may not be a reflexologist alive" (Shimp, 1989, p. 163). Protests aside, at the core of every behavioristic theory is a claim that the environment controls behavior. From the beginning, behaviorists have asserted, like Donahoe and Palmer, "Although the organism is the locus of environmental action, it is the environment, and not the organism, that is the initiator and shaper of behavior" (1989, p. 410). When Hayes and Brownstein (1986) discussed prediction and control as criteria for evaluating behavioristic analyses of behavior, they said, "One could ask, for example, how do we know that *this is* the relevant stimulus for *this* behavior? The answer is of the general form that when we change *this* stimulus (and not *that* stimulus), we get a change in *this* behavior (and not *that* behavior)" (p. 178, emphases in the original). And Skinner claimed, "The ways in which behavior is brought under control of stimuli can be analyzed without too much trouble..." (1989, p.14).

Here, we merely test results that would ensue were it in fact true that independent environmental stimuli specify instantaneous details of behavior and its consequences.

The "cognitive" model. Our cognitive model stands for all theories that say actions originate not from current external events, but from internal causes, inner traits, tendencies, propensities, sets, plans, attitudes, aspirations, symbol-generating processes, programs, computations, coordinative structures, or some kind of systematic endogenous brain activity. No major theory of this sort proposes that behavior is entirely spontaneous; in one way or another they say the internal causes of present behavior formed and changed slowly, during past experience with the outside world—the recent past in some theories, the geologically distant past in genetic theories of behavior. In cognitive theories, the link between present behavior and influences in the present external world ranges from weak to almost nonexistent. In many texts on cognitive theory, there is no mention of overt action, much less an attempt to explain such actions. When there are explanations, the causal chain runs from input to cognition to command to action to consequence.

Kihlstrom (1987) succinctly identified the linear causal model in cognitive theory: "Cognitive psychology comes in various forms, but all share an abiding interest in describing the mental structures and processes that link environmental stimuli to organismic responses..." (p. 1445). Each step of the assumed chain from stimulus (input) to response (output) is described in detail by various cognitive theorists. For example, Real (1991) describes how inputs from a variable world would be transformed, in three sequential stages, into cognitive "representations":

> ... three stages may be viewed... as three components of a single dynamical system mechanistically tied to the organism's nervous

system. The encoding of information would... correspond to initial inputs, computational rules correspond to transient dynamics, and representations would correspond to the equilibrium configurations resulting from the transient dynamics. The animal reaches a representation of the environment through the operation of specific computational rules applied to a particular pattern of incoming sensory information (p. 980).

In a discussion of computations which they assume cause movement, Bizzi, Mussa-Ivaldi, and Giszter (1991) complete the chain between representations and actions: "... the central nervous system must transform the neural representation of the direction, amplitude, and velocity of the limb, represented by the activity of cortical and subcortical neurons, into signals that activate the muscles that move the limb" (p. 287).

Some theories combine cognitive and S-R models. In their simplest forms, hybrid models say that the mind-brain receives "inputs," then produces direct transformations of coordinates from "perceptual spare" to "action space" that are required to initiate commands to move the body or part of the body to a point specified in the input (as examples, see P. M. Churchman, 1986; P. S. Churchman, 1986). Such models reduce cognition and neurology to a simple table-look-up.

A more complex hybrid S-R/cognitive model was endorsed by the cognitive theorist Allen Newell (1990) in the 1987 William James Lectures. Newell spoke of how "It is possible to step back and treat the mind as one big monster response function from the total environment over the total past of the organism to future actions..." (p. 44). On a more immediate scale, he said, "The world is divided up into microepics which are sufficiently distinct and independent so that the control system (that is, the mind) produces different response functions, one after the other" (p. 44). For strategic purposes, Newell places his theory in the category of cognitive theories that he says do not effectively explain how perception and motor behavior are linked to central cognitive processes. Then he says that such theories "... *will* never cover the complete arc from stimulus to response, which is to say, never to tell the full story about any particular behavior" (p. 160). In his allusion to the reflex arc, Newell remarkably implies the equivalence of the causal models in his cognitive theory and in reflexological theory.

In either their simple or complex forms, hybrid S-R/cognitive models produce results identical to those of S-R models, so we will not discuss them further.

The perceptual control theory model. The PCT model, which we discuss later at some length, is the least familiar of the three models. In brief, it proposes that there is a simultaneous two-way interaction between organism and environment (see Hershberger, 1989; Marken, 1990; and

Powers, 1973, 1989, 1992). In PCT, the basic unit of behavior is not the linear input-output chain, but the negative-feedback loop, which has properties different from the units of the other two models and implies interesting consequences about the way an organism's actions alter the outside world.

"Models"

We use the term "model" in the very narrow sense in which an engineer would use it: a precise quantitative proposal about the way some system operates in relation to its environment. Most behavioral scientists use *descriptive* models, which merely rephrase (usually in words; sometimes in mathematical form) previously observed relationships between organism and environment. There are unlimited ways to restate behavioral data. If each of them passes as a *model* of behavior, then the list of seemingly plausible models is also limitless. The availability of many equally plausible descriptive models is behind the mistaken assumption, common in behavioral science, that models are poor substitutes for real understanding—that if one understood the phenomenon at hand, one would state the facts, not a "mere" theory or model.

But "model" also means, in the present context, a *generative* model, in which the proposed organization is stated in a way that can be used to calculate behavior as a function of moment-by-moment variations in the independent variable. By that usage, a model does not substitute for knowledge. To the contrary, simulation of a well-posed model rigorously tests one's presumed knowledge of the causal principles at work in behavior.

S-R theory as a model. Calculations of the correlation between a dependent and independent variable produce a correlation coefficient, a regression coefficient, and an intercept. In most behavioral research, little attention is paid to the regression coefficient and intercept, one reason being that the typical scatter of the data is large enough to make a linear regression line almost useless for predicting behavior. But, by the logic of the S-R approach, the regression equation constitutes both a generative model and a description. It is a first approximation to a proposed law of behavior: at every moment, the behavioral measure is proportional to the magnitude of the independent variable. If that law is true, one can vary the independent variable and calculate (predict) the dependent one strictly from the previously determined regression equation.

It can be argued that this strict interpretation of a regression equation is inconsistent with the state of the art in behavioral science—all we can hope for now, in most cases, is to establish the presence or absence of a statistically significant relationship. Our reply gives the benefit of the doubt to the theory underlying the S-R concept. If, given as many years as necessary, methodologies improve, sources of variance are eliminated, and better data are obtained, then regression equations will become meaningful. When they do, there will be an obvious test for whether a proposed regression coefficient is a law of behavior. In the regression equation, one can impose a new pattern of the independent variable and calculate the resulting pattern of behavior, the dependent variable. The modeled result can be compared against what happens when the organism encounters the altered independent variable. In more elaborate form, this process of testing a model against actual events is the basic methodology of the physical sciences. Used in this way, the regression equation is a generative model.

We use an alternative to waiting for years for data to improve: we apply this method in an experiment so simple that the regression line is highly meaningful, and random variation is a minor factor. We subject the S-R model to a test under conditions that should make it work as well as it ever will.

Cognitive theory as a model. We give the cognitive model a similar treatment. Cognitive models are more difficult to test and defend than S-R models; there is no simple way to determine whether a given cognitive model is correct, as well as plausible. No matter how well a model proposing a specific organization of the mind-brain predicts behavior, one cannot test the model objectively by, for example, deriving a regression line based entirely on observable variables. There is no way to know whether some other cognitive model would not work as well or better. There is only one regression line that best fits the behavioral data, but there are many seemingly plausible cognitive models.

Kugler and Turvey (1987) aptly described the problem of non-unique computational models for behavioral output:

> Whereas physical events are said to follow uniquely from their causes, internally consistent, logical descriptions of the causal process are multiple .... How does one get from the existence of multiple (logical) descriptions to a unique (causal) description? Dressing up logical formulae in instantiable programs does not resolve the uniqueness problem. Many programs can give rise to the same sequence of machine outputs (p. 28).

To avoid problems of this sort, we give cognitive models the same benefit of the doubt that we give S-R models. Given proper knowledge of the history and properties of the environment, and the correct internal computations, the ideal cognitive model should calculate exactly the motor outputs required to produce a preselected result. Of course, even a perfect cognitive model would require experience with an environment to build up knowledge of its properties: if the environment changed, the model would need new interactions with the altered form before it could again compute the correct action.

We test the cognitive model by assuming that it is perfect: it makes optimal use of information and computes the same required action on successive trials, and the motor systems perfectly obey its commands.

The reasoning behind our approach to the models is simple: in a well-defined experiment, if quantitative predictions by both the S-R and cognitive models, given the benefit of every doubt, are incorrect, and the PCT model predicts correctly in the same experiment, there will be excellent reason to say that the control-theoretic model is right and the other two are wrong, for that experiment. How far one generalizes the result depends on how dear are the parallels with other experiments and the simple one we use: we leave such judgments to the reader.

Perceptual Control Theory as a Model

Perceptual control theory always considers two simultaneous relationships: (a) the observed dependence of stimulus inputs on behavioral outputs and independent events, and (b) a conjectured dependence of behavioral outputs on stimulus inputs.

The environment equation. The first relationship the PCT model describes is how the input to an organism depends on the organism's actions and on disturbances arising simultaneously with behavior but independently of it in the external world. To simplify this part of the model, we restrict all variables in the experiment to change in a single dimension, described later. Consequently, the variable at the organism's input is simply the sum of a physical effect from the organism's output and another physical effect from an independent disturbance. The apparatus (a computer system) records exactly what these relationships are and exactly what disturbance is acting at any moment. This part of the model is completely determined by the experimental setup; it is a statement of fact, not a conjecture, and it is illustrated in detail by Bourbon, Copeland, Dyer, Harman, and Mosely (1990).

The organism equation. Perceptual control theorists assume an organism can be modeled as a system that senses some aspect of the environment that is then represented internally as a one-dimensional perceptual variable. The magnitude of this variable is compared continuously against a reference signal (or reference magnitude) inside the organism or the model of the organism. Any difference between the reference signal and the perception is a non-zero “error signal” which drives action, again in a single dimension of variation.

This part of the model can be treated exactly as a regression equation. The slope of the regression line represents the incremental ratio of output to input, and the intercept represents the setting of the internal reference signal. The slope reflects measured output as a function of measured input; the intercept is the magnitude of input for which the output does not change. Control theorists assume that the value of the input for which the organism produces no change in output is the input that the organism specified in advance.

The system equations. The organism and environment equations form a system of equations; for examples, see Pavloski, Barron, and Hogue (1990, pp. 33-37); Powers (1973, pp. 273-282; 1978, pp. 422-428); and Runkel (1990, pp. 93-99). There are two system variables (the input and output variables) and two equations. The input and output variables appear in both equations, and each must have only one value at a time. Consequently, the system can be solved for each variable as a joint function of any system constants and the values of the two independent variables (the external disturbance and the internal reference signal).

Our experiments use random disturbances that cannot be represented by any reasonable analytic equation. Consequently, in the PCT model, we calculate numerical solutions of the system equations. Numerical solution of system equations, with time as a parameter, is called simulation.

Simulation. Simulation recreates, through computation, a continuous relationship among system variables and independent variables. The experimenter causes a pattern of changes in the independent variables, while the equations for the model continuously compute the states of dependent behavioral variables at the input and output. For a good model, the results of a simulation look very much like a recording of an organism's actions in an experiment where the independent variables change in exactly the same way as during the simulation; for a bad model, the results of the simulation do not resemble those produced by the organism.

Simulation involves at least two stages. The first matches simulated behavior to real behavior, after the fact, by adjusting the parameters in the model. The second stage uses a new pattern of variation in the independent variable, with the model's parameters set as previously determined, and records the behavior of the model. Then the new pattern of variation is applied to the person, whose behavior is recorded and compared with the model's behavior. In the sciences and in engineering, models are often tested in a third stage (as we do here), with both a new pattern of variation for the independent variable and a new kind of environmental disturbance, not used in the original parameter determinations. In this third stage, the model predicts, in simulation, relationships not previously observed.

Reduced to its essentials, the logic of simulation resembles more familiar ways of studying relationships and testing to see if they generalize. It is, however, much more exacting: it compares modeled and actual behaviors instant-by-instant, rather than in terms of static data sets. For the present experiments, the models predict thousands of values for several variables, all of which are compared with the values produced by a participant. The success or failure of a prediction is immediately obvious.

Some people argue that models which work properly in very simple situations might not work when complexities occur. The converse of that hypothesis, also sometimes offered, is that failure of a behavioral theory in a very simple experiment doesn't necessarily mean that it will fail in more realistically complex studies. But engineers, who deal with both simple and complex systems, would not agree. Certainly, a model that works in a simple situation might need considerable revision to work in a more complex situation. But if a model fails to work in the simplest possible circumstances, there is no chance that it will successfully predict more complex phenomena. Complexity can be an excuse for failures of a model in a complex situation, but not in a simple one. If the core assumptions of a model fail in simple experiments like ours, there is no chance the model will work in more complex circumstances.

The Experiment

The Task

Participants in this three-phase experiment move a control handle in one dimension, forward and backward. On a computer screen in front of them is a short horizontal bar, the "cursor," distinct from the background, which moves up as the handle moves forward and down when it moves back. Flanking the path of the cursor are two more bars, the "target," that remain even with one another and move slowly up and down the screen, following a path generated by the computer. The person's task in all phases of the experiment is to keep the cursor exactly between the target lines. (There is nothing special about that relationship between cursor and target; the person could easily select any other.) This task is known as "tracking." When the target is stationary, it is called compensatory tracking; when the target moves, as it does here, it is called pursuit tracking.

We can easily modify the experiment to include perceptual variables other than spatial position. For example, the handle can be set to alter the size or shape of a geometric figure, change the magnitude of a number displayed on the screen, or alter the pitch of a sound. And tracking can occur across stimulus attributes and sensory modalities, as when a person uses the handle to make the pitch of a sound match the magnitude of a number or the vertical position of a target. All relationships observed during a simple tracking experiment are found in these other tasks; any of them can be used to make the points we make here.

The Conditions: Three Phases

Phase 1. In Phase 1, the target moves up at constant speed to a preset limit, then down at a constant speed to another preset limit, and so on, in a triangular wave. Each excursion up or down takes 2.8 seconds. The person practices as long as necessary to keep the cursor between the targets with an error of no more than three per cent of the total movement averaged over one minute. Data from the final minute of practice when this criterion is reached are saved as the data for the experimental run.

The relevant parameters are estimated for each model, and then the models reproduce the person's behavior. In the next two phases, we use the parameters thus determined to create a simulated run before the person runs a single one-minute trial. No model is altered, in any way whatsoever, from this point on.

Phase 2. Conditions in Phase 2 are the same as in Phase 1, except that there is a probability of 2/3 that the target speed will differ from the last speed on any given up or down excursion. The speed of each excursion is selected randomly from 1.4, 2.8, or 5.6 seconds per excursion, with a mean of 2.8 seconds over the one-minute experimental run (the same mean excursion time as in Phase 1). The person must still move the handle to keep the cursor between the target marks. A few minutes prior to the person's run, each model is run with the same randomly generated pattern of variations in target speed that the person will experience. The person gets no practice: the first run under these new conditions is the only run for Phase 2.

Phase 3. Conditions are the same as in Phase 2, except that now a smoothed random disturbance also acts on the cursor. The disturbance is created at the start of the entire experiment by smoothing the output of a random-number computer algorithm and storing the resulting waveform. The same disturbance is used in runs by the models and the person. Cursor position is determined by the sum of handle displacement from renter and the momentary magnitude of the disturbance. Again, the person does a single one-minute run with no practice. A few minutes before the person's run, each model predicts the results, with a new pattern of target excursions and with the disturbance acting on the cursor.

The experimental variables. During each 60-second experiment, each variable is sampled every 1 /30 second, for a total of 1800 values per variable. In the figures illustrating the results, every third value is plotted. There are three measured variables: the positions of the target (T), handle (H), and cursor (C).

Phase 1

The person's data. The person kept the cursor even with the target, as shown in Fig. 1A. The perfectly regular triangular wave in the upper part of the figure is the vertical target position across time. The slightly less-regular wave that closely follows it is the cursor position created by the person. In the lower part is the handle-position record, identical to the cursor-position record because handle position directly determined cursor position. (The handle-position plot is scaled to be the same amplitude as the cursor-position plot; we use this scaling in all figures).

The mean vertical distance between the cursor and target was -0.8 units of screen resolution (S.D. -1.8; total vertical range on the screen = 200 units). The following Pearson correlation coefficients describe the

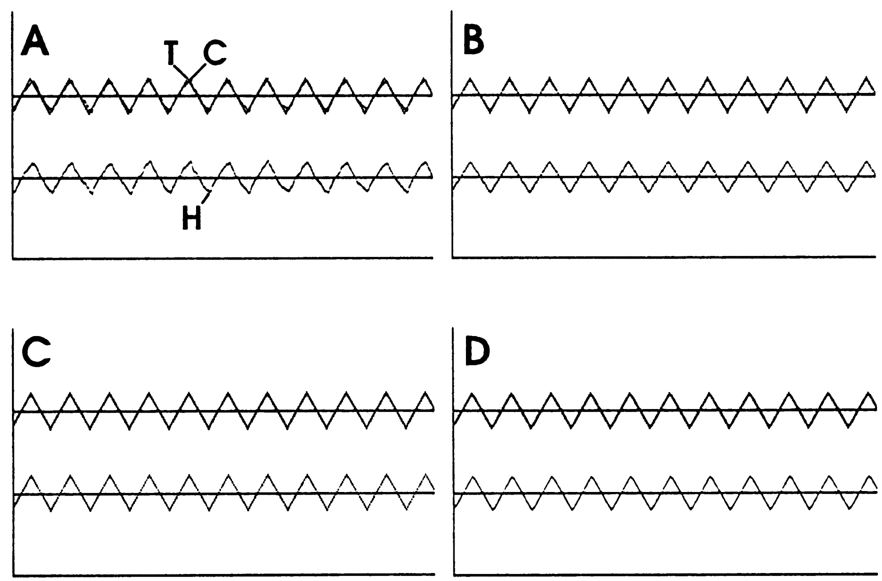

Figure 1. Results of pursuit tracking, Phase 1: data from the person (A); reconstructions of the data by the stimulus-response model (B); by the cognitive model (C); and by the control-system model (D. In A, H = handle, T = target, and C = cursor. For target and cursor, ”up” in the figure is toward the top of the computer monitor; for handle, “up” is away from the person. The duration of each experiment is 60 seconds.

relationships among variables in Fig. 1A: between positions of the cursor and target, .977; handle and target, .977; and handle and cursor, 1.0. In the regression of handle on target, the slope was 0.89 (the person moved the handle the equivalent of 0.89 screen units for every movement of one unit by the target), and the intercept was -0.8, identical to the average difference between positions of the cursor and target.

Testing the models: The rationale. In simulations of the models, computations begin with all variables set to the same initial values from the first moment of the run by the person and are repeated 1799 times, once for every 1/30 second in the run by the person. Each model produces handle positions in its unique way, but a common procedure determines cursor positions.

Establishing the S-R model. We remind readers that we do not compare the relative merits of the many varieties of behavioristic theory, nor do we examine or challenge behaviorists' descriptions of conditions in which learning occurs. We merely examine consequences that would ensue were behavior controlled by an independent antecedent variable —were behavior literally "under environmental stimulus control."

Our simple S-R model is rigorously true to the requirements laid down for laws of behavior by B. F. Skinner (1953):

> The external variables of which behavior is a function provide for what may be called a causal or functional analysis. We undertake to predict and control the behavior of the individual organism. This is our "dependent variable"—the effect for which we are to find the cause. Our "independent variables"—the causes of behavior-are the external conditions of which behavior is a function. Relations between the two—the "cause-and-effect relationships" in behavior—are the laws of a science (p. 35).

In our simple experiment, the only independent variable is the position of the target, determined solely by the computer program. The position of the handle depends on the actions of the person, so it is a pure dependent variable, which we model as a response to target position. In Phase 1, the handle determines the position of the cursor, which is a remote (from the person) consequence of behavior, not a cause.

Cursor movement is also a "stimulus," by any traditional definition, but it is not independent of behavior; it lies at the *conclusion* of the assumed causal chain. At best, it might be a "reinforcing" stimulus. Behavioral theorists claim that reinforcement produces long-term changes in the probability of a general class of actions (an "operant"). For example, some might say that, at an earlier time, cursor movement reinforced handle movement, which explains why the person uses the handle now. But reinforcement theory does not explain or predict how

a person produces moment-by-moment changes in behavior and in its consequences.

We use a regression equation as our S-R model. For the handle and target positions in the person's data, shown in Fig. 1A, the slope (m) of the regression of handle on target is 0.89, and the offset (intercept, b) is -0.8. We represent target position as t, handle position as h, and cursor position as c. Therefore, the S-R model for handle position is of the form

$$h = mt + b,$$

and the position of the cursor is modeled as

$$c = h.$$

Results of running the S-R model. To "run" the S-R model, we start with all variables at their values during the first instant of the run by the person, then we multiply the remaining 1799 target-position values, in sequence, by the slope m and add the intercept b, and obtain the successive predicted positions of the handle and cursor, shown in Fig. 1B.

The positions of handle and cursor created by the model resemble those from the person: the correlation between modeled and actual handle positions is .977; between modeled and actual cursor positions, also .977. Our simple reflexological model accounts for % per cent of the variance (r-squared) in the behavioral data from Fig. 1A; the regression equation is highly meaningful.

Establishing the cognitive model. Our goal with the cognitive model is not to compare the many diverse computational algorithms studied by cognitive and brain scientists. We merely examine the consequences that would ensue, were it possible for a system to reliably compute the same output, no matter how it does the computation. Our cognitive model assumes that, during the practice period, some central process learns and models the amplitude and frequency of target movements and computes commands that cause the muscles to move the handle, and thus the cursor, in a pattern as close as possible to that of the target.

A detailed version of this model would use a program loop simulating a "higher cognitive process" to compute handle positions independently of target movements. It would generate commands for the amplitude, frequency, and shape of the movements. But severe phase errors (mismatches in timing between the positions of the target and the model's handle) would develop unless we gave the model exact information about the frequency of the target and started it at exactly the right moment with exactly the right initial conditions. To assure that there were no errors, we would tell the model exactly how to move the handle to re-create the results of Phase 1. To achieve the same result, without the complex computations, we simply assume that, however the cognitive model works, it works perfectly: it computes handle movements to match the average pattern of previous target movements. For the last minute of practice, it uses information accumulated earlier to command movements that reproduce the movements of the target (of course the model we use here does not actually need any practice).

This makes the cognitive model exceedingly simple: it is of the form

$$h = t.$$

Handle movements perfectly reproduce movements of the target that occurred during the practice run, and the resulting cursor movements also perfectly reproduce the movements of the target.

Results of running the cognitive model. A run of the cognitive model is extremely simple: since h = t and c = h, we simply plot the successive target position values as c and as h. The upper trace in Fig. 1C shows target and cursor positions perfectly superimposed; the lower trace of handle position is identical to the upper traces. The positions *of* handle and cursor created by the model are like those from the person: the correlation between modeled and actual handle positions is .977; between modeled and actual cursor positions, also .977.

Establishing the control-theory model. The environment part of the PCT model is just a description of the external situation: cursor position depends on handle position plus the magnitude of any possible disturbance. The environment equation is

$$c = h + d.$$

In Phase 1, the disturbance magnitude is zero.

The fact that the cursor is also a dependent variable wholly or partly determined by handle position is not a problem, because both the organism equation and the environment equation form a single system of equations. We symbolize the perceived separation of cursor and target, c - t, as p, which we take as the real input variable. This variable p is compared against a reference level p*, which specifies the state of p at which there will be no change in output; it is the value of p that the person intends to experience. Any difference between p and p* is called "error." The output, which is the handle position h, is the time-integral of error and takes the form

$$h = k[\text{int}(p^* - p)].$$

The constant k is the "integration factor." It represents how rapidly the person moved the handle for a given difference between the perceived separation p and the reference separation p*; k is expressed in units of screen resolution the cursor would move per second for a given amount of perceived error.

To fit the model to the subject's behavior, we estimate p* and k, the only adjustable parameters of the model. We set p* equal to the average value of cursor-minus-target during the person's run in Phase 1. (By estimating p* from the data, we avoid claiming that we know the person is trying to keep the separation of target and cursor at zero. The person can maintain any reasonable separation-there is nothing special about p* = 0.) To estimate k, we insert the estimated value of p* into the model, then we insert an arbitrary value of k and "run" the model, a procedure we explain below. During each of several successive runs of the model, we insert a new arbitrary value of k and calculate the root-mean-square (RMS) difference between all of the cursor positions from both the model and the person. The best estimate of k is the one from the run with the smallest RMS difference.

To "run" the model, we start the handle position at the subject's initial handle position during Phase 1, and then do the following computer program steps over and over, changing the value of t on each step to re-create the target movements:

```
l: c: = h + d
2: p: = c-t
3: error: = p* - p
4: h: = h + k • error • dt
```

where dt is the physical duration represented by one iteration of the program steps. In all of the experiments reported here, each iteration represents 1 / 30 second, so dt = 1 / 30 sec. For the various terms in the program steps, k and p* are the system constants: k is the tentative value of the integration factor and p* is the estimated reference signal; t is the momentary target position, c is the cursor position, h is the handle position, and d is the disturbance magnitude-here, 0.

The fourth program step is a crude form of numerical integration; the notation means that the new value of h is computed by adding an amount (k • error • dt) to the old value of h. These are program steps, not algebra: do not cancel the h's! The "colon-equal" sign is the replacement operation, which replaces the previous value of the variable on the left with the new computed value of the argument on the right.

Results of running *the PCT model.* In the person's run during Phase 1, p* was estimated as -1 unit on the screen (-0.8 rounded), which means that, on average, the person kept the cursor slightly below the target.

Following the procedure described above, the estimated best value of the integration constant k was 8.64 in units of resolution per second.

The results of a run of the model with those estimated values of p'“ and k are shown in Fig. 1D. The positions of handle and cursor created by the model resemble those from the person: the correlation between modeled and actual handle positions is .989; modeled and actual cursor positions, also .989.

Summary of Phase 1. The person performed the tracking task reasonably well, and simulations of all three models produced results like those from the person. After this round of simulations, all three models remain defensible as explanations of the person's performance.

Phase 2

Next, we use the three models to predict behavior when one condition changes, then the person does a run under exactly the same conditions as those encountered by the models. The changed condition is that the target now moves up and down at randomly varying speeds. The mean speed is still 2.8 seconds per excursion, but on every successive excursion, there is a 2/3 probability of a change of speed that lasts until the end of the excursion, and then the next speed is selected randomly. The random changes are generated beforehand and recorded, so the same changes are presented to all three models and to the person. We have already established the three models, so our descriptions of the results are brief.

The person's data. Fig. 2A shows data from the person's run, after the models made their predictions. The person made the cursor follow the target about as well as in Phase 1. The mean vertical distance between cursor and target was -1.4 units of vertical screen resolution (S.D. = 2.2). The following Pearson correlation coefficients describe relationships among variables in Fig. 2A: between positions of the cursor and target, .966; handle and target, .966; and handle and cursor, 1.0.

Prediction of the S-R model. The linear regression equation developed after Phase 1 accurately predicts the positions of the cursor and handle despite the changes in target speed, as is shown in Fig. 2B. This is possible because, just as in Phase 1, the required handle movement is simply proportional to target movement at every instant. The positions of handle and cursor created by the model are like those from the person: the correlation between modeled and actual handle positions is .989; between modeled and actual cursor positions, also .989.

Prediction of the cognitive model. The results for the cognitive model, shown in Fig. 2C, reveal the first obvious failure of a model. The positions of handle and cursor created by the model are not like those from the person: the correlation between modeled and actual handle

positions is .230; between modeled and actual cursor positions, also .230.

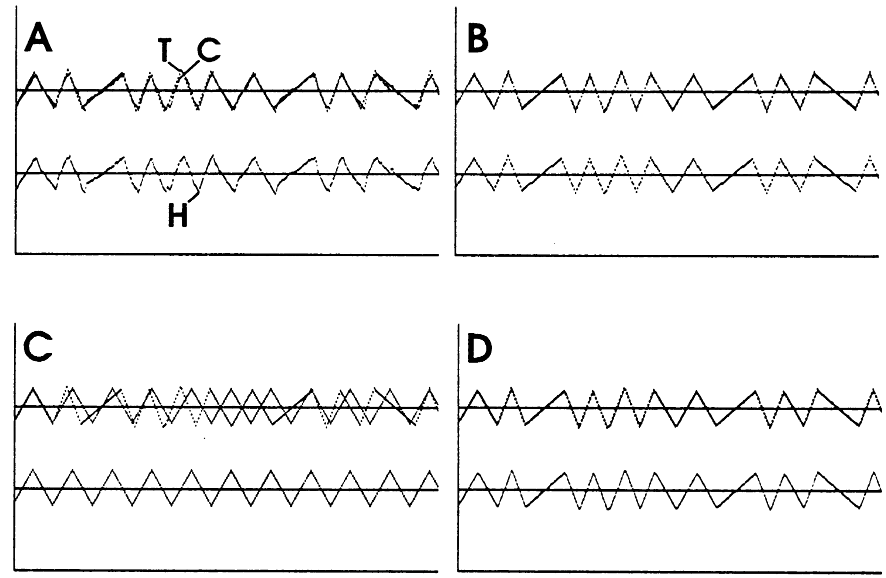

Figure 2. Results of pursuit tracking in Phase 2: data from the person (A); predictions of the data by the stimulus-response model (B); by the cognitive model (C); and by the control-system model (D. In A, H = handle, T = target, and C = cursor. For target and cursor, “up” in the figure is toward the top of the computer monitor; for handle, “up” is away from the person. The duration of each experiment is 60 seconds.

The reason for this failure is obvious. The cognitive model assesses properties of the environment and computes an action that will have a required result. But now the environment, in the form of target movements, is subject to unpredictable variation. The cognitive model gets no information about the next target speed before it is experienced. Thus, the best that a cognitive “central-process” model can do is command its output to match the best estimate of average target speed; in the present case, that average is the speed that occurred throughout Phase 1, when the motor plan was established. The cognitive model continued to produce a triangular wave of handle and cursor movement that conformed to the average waveform of target movement—a form not like the waveform of the target in Phase 2.

One might think of modifying the cognitive model so that the central processor re-assesses the environment's properties on an instant-by-instant basis. That would solve the problem, but only at the expense of converting the cognitive model into a control-system model intent on making its output match its input: the new model would be a control-system model acting like a stimulus-response model. The core concept of a cognitive motor plan would be abandoned.

Prediction of the control-system model. Fig. 2D shows the results for the control-system model. The program steps from Phase 1, using the same values for the parameters k and p*, successfully predict the person's handle and cursor positions. The correlation between modeled and actual handle positions is .981; between modeled and actual cursor positions, also .981.

Summary of Phase 2. The person performed the tracking task with reasonable accuracy, and simulations of the S-R and PCT models produced results like those for the person. However, the cognitive model continued to make its output follow the path 'learned" during Phase 1; consequently, its cursor did not follow the now-erratic waveform of the target. After this round of simulations, only the S-R and PCT models remain reasonable as explanations of the person's performance.

Phase 3

Now the three models predict behavior under a radical change of conditions. The target still moves up and down at randomly varying speeds, as in Phase 2, but for every time-interval, a new value of a random disturbance is added to the position of the cursor. Now, with the handle held still, the cursor wanders randomly up and down. When the handle moves, the net movements of the cursor are determined by the sum of handle movements and disturbance changes.

In both previous phases, the "d" in the cursor equation, c = h + d, was zero. Now it varies unpredictably, although not rapidly (the bandwidth of variations is about 0.2 Hz). This new disturbance enters after the motor outputs of the person and the accompanying handle movements, "downstream" in the causal chain. The cause of the disturbance is hidden; the only evidence the person has about the disturbance is the deviation of cursor position from the momentary equivalent of the handle position. At any moment, there is no practical way for the person to know the degree to which either the position of the handle or the value of the disturbance affects the position of the cursor.

The person's data. As we show in Fig. 3A, the person still made the cursor track the target (mean distance between cursor and target = -1.0 screen units, S.D. = 3.0), despite the unpredictable variations in target speed and the unpredictable interference of a disturbance. Had the person not moved the handle, the correlation between positions of the cursor and momentary values of the disturbance would have been + 1.0;

that between positions of cursor and target, near 0.0. Instead, the correlation between the disturbance and cursor was only .101, while that between cursor and target was .940.

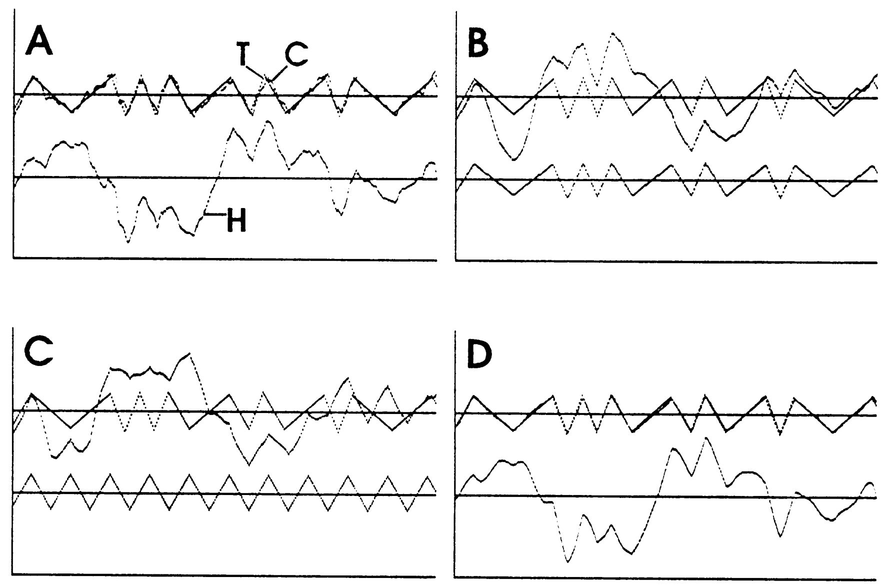

Figure 3. Results of pursuit tracking in Phase 3: data from the person (A); predictions of the data by the stimulus-response model (B); by the cognitive model (C); and by the control-system model (D). In A, H = handle, T = target, and C = cursor. For target and cursor, "up" in the figure is toward the top of the computer monitor; for the handle, "up" is away from the person. The duration of each experiment is 60 seconds.

In Phases 1 and 2, the handle alone determined the position of the cursor: the correlation between handle and cursor was + 1.0. But in Phase 3, the person moved the handle any way necessary to cancel the effects of the random disturbance on the cursor: the correlation between positions of handle and cursor is only .294, that between positions of the handle and the disturbance that moved the cursor away from the target is -.992.

Prediction of the S-R model. As we show in Fig. 3B, the S-R model failed: the correlation between modeled and actual handle positions is .296; between modeled and actual cursor positions, .385.

Successful simulation can no longer be attained by moving the handle in synchrony with target movements. That is why the person moved the handle in a pattern that deviated radically from the pattern of target movements; the deviations were exactly the ones needed to counteract the effects of the new disturbance. But the S-R model responded to the target stimulus just as before, and moved the handle proportionately to any movement of the target. The simulated cursor, now subject to an independent disturbance, did not follow the target.

To salvage the S-R model, one might propose that the cursor, too, be included in the definition of the stimulus. However, the person's data in Fig. 3A show that the cursor moved in nearly the same pattern as the target, but neither pattern resembled what the handle did. To include the cursor in the definition of the stimulus, we might conclude that the difference between the target and cursor positions is the stimulus. On further examination, we would find that this difference does not match the handle movements, either, but its time-integral does: perhaps the time-integral is the stimulus. That change is acceptable, but if we adopt it, we are left with the fact that cursor position depends, simultaneously, on handle position and the independent random disturbance: now there is no true independent variable in the causal chain, and the core premise of any model of stimulus control over behavior is abandoned. Neither the cursor nor any relationship between the cursor and any other variable can be described as a pure independent variable, because it is also, at every moment, a dependent variable.

Prediction of the cognitive model. Fig. 3C shows that the prediction by the cognitive model failed. The model followed its plan learned in Phase 1 and moved the handle to conform to the average behavior of the target. It should have moved the handle in the erratic pattern produced by the person, shown in Fig. 3A. The correlation between predicted and actual handle positions is .119; between predicted and actual cursor positions, .151.

Even if we gave the cognitive model more practice in the new situation (and the ability to learn), it would revert to essentially the same actions. The average deviation of cursor speed from 2.8 seconds per excursion is zero. The average amount of disturbance applied to the cursor closely approximates zero. Neither the next speed of the target nor the next variation in the disturbance is predictable. No matter how smart one wants to make the central processor when it comes to predictions, we can always make the disturbances still more random. Any cognitive model must compute output that is calculated to have a desired effect. It can base its computations only on experience with properties of the external world. When those properties contain significant instant-by-instant irregularities, as they do in our simple experiment, the core concept of the cognitive model cannot work. Unless, of course, it is modified to compare its plan of the world against its momentary perceptions of the world and to act so as to eliminate any discrepancy, but those modifications would make the model a control-system model.

Prediction of the control-system model. As we show in Fig. 3D, the control-system model produced precisely the outputs required to maintain a pre-selected target-cursor separation, despite two kinds of random variation that called for pronounced changes in the output pattern. The PCT model faithfully predicted the person's behavior. The correlation between actual and predicted handle positions is .996; between actual and predicted cursor positions, .969. Correlations as high as those here, between tracking behavior and predictions by PCT, are commonplace, even when the interval between predictions and behavior is as long as one year as is reported by Bourbon, Copeland, Dyer, Harman, and Mosley (1990).

To avoid drawing this paper out any longer, we omit analyses of other variations that the person and the PCT model can handle, with no change in the model's parameters. Both the person and the control-theory model continue to track accurately if we alter the scaling factor that converts handle movement into cursor movement; if we add a third or a fourth or a fifth independent source of disturbance to target speed or cursor position; if we put nonlinearity into the connection between handle and cursor (the person and the model still move the handle in an inverse nonlinear relationship to target and disturbance); or if we make the ratio of handle movement to cursor movement time-dependent (at a reasonable speed). None of these variations can be handled by the core concepts of the S-R or cognitive models. Yet all of these variations, as well as those shown in the three phases of our experiment, are commonplace in the real environments where real behavior must work.

Discussion

We attempted to determine if core assumptions about the immediate causes of behavior in three different models of behavior are consistent with what is known about the world in which behavior occurs. We compared specific predictions made during simulations of the three models with the performance of a person for three phases of a simple task. We concluded that the causal assumptions in a control-theoretic model are consistent with what is known about the world, while those in any pure stimulus-response (stimulus-control) model, or any pure cognitive-control (neurological-control) model, are not. The control theory model assumes that, when organisms act, they produce correspondences between their immediate perceptions of selected variables in the world and internal (to the organisms) reference states (reference signals) for those perceptions.

We did not ask whether reference signals exist in any particular physical form, or, if they do, whether they are "gained" through interaction with the world, whether animate, inanimate, or social, or are inherited as part of a "genetic code." Robinson (1976) wrote of this issue in a discussion of Aristotle's concept of "final cause," which refers in part to a person's goals or intentions: "The issue is not *how* a given goal or intention was established. Rather, the issue or proposition is that outcomes are never completely understood until the final cause is apprehended, no matter what 'caused' the final cause" (p. 91, emphasis in the original). In our simulations, by hypothesizing and estimating the magnitudes of "reference signals," whatever their origins, that function in the manner of "final causes" within a control-system model of a person, we can understand and predict the outcomes when the person controls selected perceptions of parts of the unpredictably variable environment.

Modeling as a proper test of theory. The success or failure of our simulations immediately revealed the robustness, or lack of robustness, of alternative models of behavior. Other behavioral scientists recognize the importance of comparing the simulated behavior of models against the actual behavior of organisms. In a critique of conventional statistical methods in psychology, Meehl (1978) remarked:

> In my modern physics text, I am unable to find a single test of statistical significance. What happens instead is that the physicist has a sufficiently powerful invisible hand theory that enables him to generate an expected curve for his experimental results. He plots the observed points, looks at the agreement, and comments that "the results are in reasonably good agreement with the theory." Moral: *It is always more valuable to show approximate agreement of observations with a theoretically predicted numerical point value, rank order, or function form, than it is to compute a "precise probability" that something merely differs from something else" (p.* 825, emphasis in the original).

Similarly, Dar (1987) wrote:

> In physics... theories are tighter and lead to precise predictions. As a consequence, (a) if the numerical result is as predicted (that is, close enough to the predicted point value or curve), it will be very difficult, in contrast to the situation in psychology, to offer a reasonable alternative theory for that. This is because it is difficult to imagine alternative states of nature that will lead to the exact same curve or numerical result. (b) If the experimental result is not as predicted, some serious revision of the theory would be required. This is because a tight theory simply does not allow for significant (I do not mean "statistically significant") discrepancies from predicted outcome (p.148).

And in his review of a book on cognitive theory, the behaviorist Shimp (1989) declared:

> A theory that behaves, that produces a stream of behavior, would seem in an intriguing way to fit better with Skinner's chief criterion for a good theory than do many more common sorts of behavioral theory. Skinner has argued that a good behavioral theory is a theory on the same level as the behavior itself. What is closer to the level of a behavior stream of an organism than a behavior stream of a theory? (p. 170).

We could not say it better. On any given experimental run, our simulations produced multiple simultaneous streams of behavior, altogether comprising thousands of predicted data points. The levels of agreement between the simulations and the behavior of a person allowed us to immediately assess the adequacy of the three models of behavior and of their implied models of the world.

The worlds implied by the models. For all three models, the results reported here would be general. Within its physical limits, any S-R system could make its movements match any target input, no matter how unpredictable. But, as happened with the cursor in Phase 3, if the consequences of those movements were disturbed, they would always deviate from the target by an amount equal to the variations in the disturbance.

Upon its first encounter with a new pattern of input, no cognitive system could compute commands to immediately make its behavior match the input. After some time, of course, an appropriately endowed cognitive system could search for a new pattern of commands. But if the input followed an unpredictable path or were presented only once or too few times for the system to "compute" an appropriate plan, learning would be impossible. Furthermore, if the consequences of its actions were continuously and randomly disturbed, no command-driven cognitive system could compute behavior to keep the consequences in any consistent relationship with the input. To do that, the behavior must deviate from its original pattern by precisely the amount needed to cancel the effect of the disturbance, but the source of the disturbance cannot be sensed in advance to allow anticipatory compensations in the commands for behavior.

The only ways to salvage the traditional models, short of turning them into control systems, rely on whimsical assumptions about the world. For example, the S-R model might still work if it were only necessary that changes in stimulation result in corresponding changes in behavior, with no regard for the consequences of behavior; and the cognitive model might still work, if it were only necessary that movements repeat, while their consequences were allowed to change at random. But those assumptions contradict any reasonable understanding of behavi r and its role in survival: behavior is functional, and its consequences matter. An alternative defense is to assume that the antecedents of behavior never change, or that they conveniently change across a small enough set of discrete options so that we can always recognize which one is present and perfectly match it with computed outputs-either that, or we must anticipate the changes by "precognition." And nothing must ever disturb the consequences of behavior. The world demanded by those assumptions is not the one we know.

In contrast, within broad limits, any perceptual control system would vary its behavior to keep its perceptions of a controlled variable at the value specified by a reference signal, even if both the target event and the consequences of the system's actions were subject to unpredictable variations.

We live in a changeable world, in which organisms with behavior determined solely by environmental stimuli or solely by internal commands could not survive; but theories of behavior that postulate control by stimuli or by commands have survived for centuries largely because they are not systematically exposed to the test of modeling. To modify cognitive or S-R models so that, like living systems, they might thrive amidst change, we must abandon the core concept that behavior is at the end of a causal chain, wherever the chain allegedly begins. We must give each model an internal standard and a process for comparing present perceptions against that standard. But then the models would all be control systems, each controlling its input.

Conclusions. The sciences of life reflect a three-century commitment to linear models of cause and effect, with behavior as the final step in a causal sequence. If we are to advance our understanding of life, we must question those venerable models, however plausible they seem. We can no longer embrace them, knowing that they presuppose nonexistent worlds. To question our traditional models raises the specter of difficult change; but if we retain them, with their fanciful worlds, we risk the trivializing and decline of our science.

The search for alternative models of behavior can begin with a simple change in the question we ask, from "Why is behavior the way it is?" to "Why is the world the way it is?" The answer to the new question includes a long-elusive answer to the old one: the behavior of organisms controls many variables in the world.

Acknowledgments

We thank Andrew C. Papanicolaou, Philip J. Runkel, Gregory Williams, and William D. Williams for their critical reviews of several early versions of the manuscript, and Glenn Millard for valuable assistance in preparing the figures.

References

Bizzi, E., Mussa-Ivaldi, F. A., & Giszter, S. (1991). Computations underlying the execution of movement: A biological perspective. *Science*, 253, 287-291.

Bourbon, W. T., Copeland, K. C., Dyer, V. R., Harman, W. K., & Mosley, B. L. (1990). On the accuracy and reliability of predictions by control-system theory. *Perceptual and Motor Skills*, 71,1331-1338.

Churchland, P. M. (1986). Some reductive strategies in cognitive neurobiology. *Mind*, 95, 279-309.

Churchland, P. S. (1986). *Neurophilosophy: Toward a unified science of the mind-brain*. Cambridge, MA: MIT Press.

Dar, R. (1987). Another look at Meehl, Lakatos, and the scientific research practices of psychologists. *American Psychologist*, 42,145-151.

Donahoe, J. W., & Palmer, D. C. (1989). The interpretation of complex human behavior: Some reactions to *Parallel distributed processing*, edited by J. L. McClelland, D. E. Rumelhart, and the PDP research group. [Book review]. *Journal of the Experimental Analysis of Behavior*, 51, 399-416.

Hayes, S. C., & Brownstein, A. J. (1986). Mentalism, behavior-behavior relations, and the behavior-analytic view of the purposes of science. *The Behavior Analyst*, 9, 175-190.

Hershberger, W. A. (Ed.). (1989). *Volitional action: Conation and control.* Amsterdam: North-Holland.

Kihlstrom, J. F. (1987). The cognitive unconscious. *Science*, 237, 1445-1452.

Kugler, P. N., & Turvey, M. T. (1987). *Information, natural law, and the self-assembly of rhythmic movement*. Hillsdale, NJ: Erlbaum.

Marken, R. S. (Ed.). (1990). Purposeful behavior: The control theory approach. [Special issue]. *American Behavioral Scientist, 34(1).*

Meehl, P. E. (1978). Theoretical risks and tabular asterisks: Sir Karl, Sir Ronald, and the slow progress of soft psychology. *Journal of Consulting and Clinical Psychology, 46,* 806-834.

Newell, A. (1990). *Unified theories of cognition.* Cambridge, MA: Harvard University Press.

Pavloski, R. P., Barron, G. T., & Hogue, M. A. (1990). Reorganization: Learning and attention in a hierarchy of control systems. *American Behavioral Scientist, 34,* 32-54.

Powers, W. T. (1973). *Behavior: The control of perception.* Chicago: Aldine.

Powers, W. T. (1978). Quantitative analysis of purposive behavior: Some spadework at the foundations of experimental psychology. *Psychological Review, 85,* 417-438.

Powers, W. T. (1989). *Living control systems: Selected papers of William T. Powers.* Gravel Switch, KY: Control Systems Group.

Powers, W. T. (1992). *Living control systems IL Selected papers of William T. Powers.* Gravel Switch, KY: Control Systems Group.

Real, L. A. (1991). Animal choice behavior and the evolution of cognitive architecture. *Science, 253,* 980-986.

Robinson, D. N. (1976). *An intellectual history of psychology.* New York: Macmillan.

Runkel, P. (1990). *Casting nets and testing specimens: Two grand methods of psychology.* New York: Praeger.

Skinner, B. F. (1953). *Science and human behavior.* New York: Free Press.

Skinner, B. F. (1989). The origins of cognitive thought. *American Psychologist,* 44,13-18.

Shimp, C. P. (1989). Contemporary behaviorism versus the old behavioral straw man in Gardner's *The mind's new science: A history of the cognitive revolution.* [Book review] *Journal of the Experimental Analysis of Behavior,* 51,163-171.

The Control Systems Group is a membership organization which supports the understanding of cybernetic control systems in organisms and their environments: *living control systems.* Academicians, clinicians, and other professionals in several disciplines, including biology, psychology, social work, economics, education, engineering, and philosophy, are members of the Group. Annual meetings have been held since 1985. CSG publications include a newsletter and a series of books, as well as this journal. The CSG Business Office is located at 73 Ridge Pl., CR 510, Durango, CO 81301; the phone number is (303)247-7986.

The CSG logo shows the generic structure of cybernetic control systems. A Comparator (C) computes the difference between a reference signal (represented by the arrow coming from above) and the output signal from Sensory (S) computation. The resulting difference signal is the input to the Gain generator (G). Disturbances (represented by the black box) alter the Gain generator output on the way to Sensory computation, where the negative-feedback loop is closed.

Currently available:

LIVING CONTROL SYSTEMS

Selected Papers of William T. Powers

ISBN 0-9624154-0-5, 1989, 300 pages, $16.50

Introduction to Modern Psychology

The Control-Theory View

Edited by Richard J. Robertson and William T. Powers

ISBN 0-9624154-1-3, 1990, 238 pages, $25.00

LIVING CONTROL SYSTEMS II

Selected Papers of William T. Powers

ISBN 0-9624154-2-1, 1992, 277 pages, $22.00

MIND READINGS

Experimental Studies of Purpose

Richard S. Marken

ISBN 0-9624154-3-X, 1992, 212 pages, $18.00

Prices are postpaid. KY residents please add 6% sales tax. Order from: CSG Book Publishing, 460 Black Lick Rd., Gravel Switch, KY 40328.

Closed Loop

Threads from CSGnet

Spring 1993

Volume 3

Number 2

Closed Loop

Threads from CSGNet

Spring 1993 Volume 3 Number 2

Edited by Greg Williams, 460 Black Lick Rd., Gravel Switch, KY 40328

CONTENTS

Members of the Control Systems Group receive Closed Loop quarterly. For more information, contact Ed Ford, 10209 N. 56th St., Scottsdale, AZ 85253; phone (602)991-4860.

CSGnet, the electronic mail network for individuals interested in control theory as applied to living systems, is a lively forum for sharing ideas, asking questions, and learning more about the theory, its implications, and its problems. The "threads" in each Closed Loop, stitched together from some of the net's many conversations, exemplify the rich interchanges among netters. Some issues of Closed Loop also feature research reports by netters, in hopes of initiating new conversations.

There are no sign-up or connect-time charges for participation on CSGnet. The Internet address is "CSG-L@UIUCVMD" while CSG-L@UIUCVMD is the Bitnet address. Messages sent to CSGnet via these addresses are automatically forwarded to over 120 participants on five continents, as well as to hundreds of NetNews (Usenet) sites where CSGnet can be found as the newsgroup bit.listserv.csg-l. CSGnet also can be accessed via CompuServe, AT&T Mail, MCI Mail, or any other computer communication service with a gateway to Internet or Bitnet. For more information about subscribing to CSGnet, contact Gary Cziko, the network manager, at G-CZIKO@UIUC.EDU, phone him at (217)333-8527, or send a FAX to (217)244-7620.

Portable PCT Demonstrations

In "A General Feedback Theory of Human Behavior: Part II," by W.T. Powers, R.K. Clark, and R.L. McFarland (*Perceptual and Motor Skills 11,* 1960, 309-323), the authors describe a series of demonstrations showing the negative-feedback closed-loop operations of a "Portable Demonstrator"—that is, a normally functioning human being. Powers offered additional demonstrations which illustrate the principles of his theory of psychological control systems (now known as perceptual control theory, or PCT) in his 1973 book, *Behavior: The Control of Perception* (Aldine, Chicago; currently available from De Gruyter /Aldine, Hawthorne, New York). In particular, his 'Parable of the Rubber Bands" (pp. 241-244) has provided other perceptual control theorists with a foundation upon which to build both provocative and enlightening "hands-on" illustrations of the control-system basis of human behavior. For example, Philip J. Runkel has used rubber band demonstrations to introduce PCT concepts to behavioral scientists in his *Casting Nets and Testing Specimens: Two Grand Methods of Psychology* (Praeger, New York, 1990), while family counselor Edward E. Ford has written in his *Freedom from Stress* (Brandt Publishing, Scottsdale, Arizona, 1989) of using rubber band demonstrations to teach PCT concepts to his clients.

The following thread from CSGnet explores some possibilities for using the rubber band demonstration and other portable PCT demonstrations—most needing no fancy equipment—as practical means to understanding perceptual control theory and its implications. To introduce the thread, two sets of instructions for rubber band demonstrations are included. Readers who are unfamiliar with "rubber banding" are especially encouraged to read the brief instructions given immediately below (written by Dag Forssell for CSGnet, based on suggestions by Powers and Runkel) and to try the demonstration themselves before proceeding to the thread. Chuck Tucker's instructions for a more elaborate rubber band demonstration originally appeared in *Continuing the Conversation: A Newsletter of Ideas in Cybernetics* (12), Spring 1988, 16-19, and are reprinted here with the author's permission.—Ed.

A Do-It-Yourself Demonstration of the Phenomenon of Control
by Dag Forssell

You can demonstrate control in action to yourself and others, wherever you are, with the simple prop of two rubber bands joined by a knot. Just get a friend to help you play a game. This game will illustrate all of the elements of human control, their interactions and functional relationships.

Get two small rubber bands and join them in a knot. You hook a finger into the end of one rubber band, and your friend hooks a finger into the other. Tell your friend something like: "You are the experimenter. Move your finger as you like. Watch what I do. When you can explain what is *causing* me to do what I do, let me know."

When you sit down with your friend, place yourself so that the knot joining the rubber bands lies above some mark you can see, but which your friend probably will not notice a small mark on a table top or paper, or a piece of lint on your knee, something like that. As your friend's finger moves, move yours so that the knot remains stationary over the mark.

By agreeing to keep the knot over a target, you have adopted a standard for the position of the knot as your want. When something acts to disturb the position of the knot, you will restore the knot to its position over the mark. You will move in any way necessary to do that.

Of course, you can't keep the knot stationary if your friend moves faster than your natural reaction time can handle. Some people playing this game seem to want to move abruptly, too fast. If that happens, ask your friend to slow down. The lessons to be learned will be much more obvious to both of you if you are able to keep the knot continuously over the mark. You might say: "Don't move so fast. I can't keep up with you."

Your friend will soon notice that every motion of his/her finger is reflected exactly by a motion of yours. When your friend pulls back, you pull back. When your friend moves inward, you move inward. When your friend circles to his/her left, you circle to your left. You must do that, of course, to keep the knot stationary. Your action illustrates very plainly the phenomenon of control—that we act in opposition to a disturbance to maintain a perception we want.

Notice that you perform many different acts to maintain your perception of the knot remaining over the mark. You move your finger to the left, to the right, forward, backward, diagonally at varying speeds.

Most people, when they announce that they can explain what is causing you to do what you do, will say that you are simply imitating what they do, or mirroring it, or words to that effect. Some will put it more forcefully: that whatever they do, you are acting in opposition to it. Almost all will say or imply that they are the cause of your behavior.

A few people will notice that the knot remains stationary. That is an excellent observation, but not quite an explanation of cause. Agree, but keep asking: 'What is causing me to do what I do?" Most people will say that your intent is to do something in reaction to them. But then you deny that. They will eventually give up and ask: "All right, what is

causing your behavior?" Then you explain that you have been keeping the knot as close to the mark as possible, and that any deviation caused you to do what you did.

No, you tell your friend, your purpose has not been to oppose his or her intention. Your purpose has not been to frustrate him/her. If, instead of his/her finger, a machine had been hooked to the rubber band, you would have moved as you did. Your purpose was to keep the knot motionless over the mark. That's all.

You moved to oppose any motion of the knot away from the mark, not to oppose him/her. Your motivation had nothing to do with what he/she might have been trying to do; you did not care. You watched only the knot and the mark. Indeed, if you had not been able to see your friend's moves, your actions would have been identical. Watching the knot and the dot, you could not pay any attention to your friend's movements.

Reactions of "experimenters" will vary widely. A few will accuse you of being devious and go away grumbling. Most will be surprised, even dumbfounded, to have missed the obvious. A few will find many of their previous ideas so shaken that they will think about it for days or weeks afterward.

Play the game with your friend. Play it with several friends! Suppose you played this game with 10 of your friends. Let us say that one was in fact able to explain (without coaching) that you were only holding the knot steady over the mark and acted to keep it there. That means that nine out of 10 failed to recognize the phenomenon of control when it was right in front of them. They have never been shown what control is or how to recognize it. Without a paradigm of control, they are quite literally blind to a phenomenon that is fundamental for all living organisms.

Let us play the game again, with more visibility for both you and your friend. This time, you experiment on your friend and play the game on a piece of paper with a clearly marked target. Ask your friend to record his or her movements by holding a pen against the paper as he or she moves in response to your disturbing influence on the knot. Now we can focus on your friend's visible behavior and ask the question: 'What can a reasonable observer conclude about your friend based on what the observer can see of your friend's behavior?" What is your answer? Would you agree that you cannot draw any conclusions about your friend from his/her behavior? Your friend's behavior is clearly a product of what your friend wants, combined with the disturbances acting on what your friend is controlling. His/her behaviors are what they have to be under the circumstances, given all of the other elements and their influences.

This demonstration clearly recognizes wants and perceptions, the difference between them, thoughts that provide instructions for action, the variable we control, the actions themselves, other influences on the variable, and the extraneous, sometimes confusing byproducts of our actions.

Demonstrating Control Theory
by Chuck Tucker

In this paper I present the procedures that I have used in classrooms and conferences for demonstrating Powers' control theory. These procedures are derived from his discussion of experiments in *Behavior: The Control of Perception* (1973, 241-244). I have modified them only to the extent that I have written explicit instructions to be used by the demonstrator and a volunteer. I have found the demonstration to be a powerful tool for explaining the fundamentals of control theory. I will present the demonstration exactly as I have done it and mention some implications and possible modifications at the end of this paper.

Materials

This demonstration requires: (1) six sheets of poster paper; (2) twelve 6″ pieces of masking tape, to attach the poster paper to a smooth wall or chalkboard; (3) two short pencils of different colors (I have used black and red); (4) two large rubber bands tied together with a knot; (5) a marking pen; and (6) 5″ x 8″ index cards with instructions. An easel with a pad could be used instead of the poster paper.

Introduction

I think it is very important to get the members of the audience or class involved in the demonstration, so I begin by reading this statement from an index card: "I will, with the help of another person, perform a series of demonstration exercises to illustrate the basic principles of Powers' control theory. The demonstrations are slight modifications of those found in the book *Behavior: The Control of Perception.* I want all of you to take part in these demonstrations. It will not be useful to you unless you do take part. For each demonstration, I want each of you to watch and listen to the volunteer, and answer the question: What instructions or directions is he/she using to perform the movements in this demonstration? The volunteer will be asked to read and follow some directions, and your job is to figure out what instructions are being followed by him/her. I will give you a sheet of paper to write your answers on after each demonstration."

Then I hand out a single sheet of paper to each person, which states:

CONTROL MODEL DEMONSTRATIONS

INSTRUCTIONS: There will be six different demonstrations of a control model. For each demonstration, answer this question about the volunteer: WHAT INSTRUCTIONS OR DIRECTIONS ARE BEING USED TO PERFORM THE MOVEMENTS IN THE DEMONSTRATION? You must watch each one carefully and answer the question for each demonstration after it is completed and before the next one begins. THANK YOU.

DEMONSTRATION I

The instruction(s) used by the volunteer is (are): ________________
____.

The remainder of the sheet has a separate question for each demonstration.

Beginning the Demonstration

I begin the demonstration by reading this statement from a card: "I want someone to volunteer for some demonstration exercises. It will not be harmful to you, and all that is required is that you can read and follow directions. If you wish to volunteer, please raise your hand." I then motion to one of the persons with a raised hand to come to the front of the room, while I say 'Please come to the front of the room." Then I say "Thank you for volunteering." I introduce myself (if necessary) and have the person introduce himself/herself to me. Then I say "Please take these cards and read the top one and follow its directions." I then hand the volunteer a stack of index cards with printing on them.

The Exercises

The first card in the volunteer's stack states: "DEMONSTRATION EXERCISES—MOVE THIS CARD TO THE BACK OF THE STACK AND READ THE NEXT CARD."

The statement on the next card is: 'There are several cards, each containing a different set of directions. Read each card carefully before doing the exercise. I will ask 'Do you understand?' and you should say 'Yes' or 'No.' If you say 'Yes,' I will ask 'Are you ready?' You say 'Yes' or 'No.' If you say 'Yes,' we will do the exercise. Now move this card to the back of the stack and read the directions on the next card."

The next card states: "MOVE THIS CARD TO THE BACK OF THE STACK AND READ THE NEXT CARD."

The next card states: "EXERCISE I—In this exercise, you will be given a pencil. Take it in your hand and place it on the paper, holding it steady until you are asked if you are ready. When you are ready, you will move the pencil so that you draw the *same* diagram that I am drawing. Do this at *the same time* that I am drawing my diagram. Move this card to the back of the stack when you understand this exercise."

While the volunteer is reading this card, I take a sheet of poster paper, write "I" in its upper right-hand corner with a marking pen, and tape it to the wall or chalkboard with a piece of tape on each corner. When the volunteer moves the card, I say, "Do you understand?" If the volunteer says, "No," then I say, 'Please read the card again." After the volunteer has read the card, I again ask, "Do you understand?" I repeat this until the volunteer answers, "Yes." (I have not had to ask a volunteer more than once to reread a card.) After the "Yes" answer, I give the volunteer a black pencil, and I take a red pencil. The different colors allow the audience to distinguish between my drawing and the volunteer's. I then ask: "Are you ready?" and when the volunteer answers, "Yes," I say, "Let's begin."

Standing in front of the poster paper, I slowly begin to make a drawing from an index card without letting the audience or the volunteer see the card. Although my drawing is complicated, it need not be for the demonstration. I try to have a drawing that has straight, sawtooth, and curved lines. I make the drawing about a foot square.

I begin with a vertical line, then make a 90-degree horizontal line, and then several squares which do not overlap. These are followed by several arcs and a sawtooth line, another horizontal line, another vertical line, concluding with an s-shaped line. The drawing is done at a slow pace, and none of the lines repeat the same path, although they do intersect one another. When I finish my drawing, I remove my pencil from the paper, turn to the audience, and say, 'Please answer the question on your answer sheet for Demonstration I." Then I turn to the volunteer and say, "Please read the next card."

The next card states: "MOVE THIS CARD TO THE BACK OF THE STACK AND READ THE NEXT CARD."

The next card states: "EXERCISE II—In this exercise, you will take a pencil in your hand as you did in Exercise I and hold it steady on the paper until you are asked if you are ready. When you are ready, you will move the pencil in the same *directions as my pencil, always keeping your pencil at a distance of one foot (12") and on the same level or same plane as my pencil.* Keep your pencil on the paper at all times. Move this card to the back of the stack when you understand what you are to do in this exercise."

While the volunteer is reading the card, I remove the poster paper for Exercise I from the wall or chalkboard and put up a new sheet marked

"II." Then I ask the same questions that I did for the first exercise, and I stand in front of the paper when the volunteer is ready.

For Exercise II, I make the same drawing as I did for Exercise I. When I complete the drawing, I ask the audience to answer the question for Demonstration II, and then say to the volunteer: "Please read the next card."

The next card states: "MOVE THIS CARD TO THE BACK OF THE STACK AND READ THE NEXT CARD."

The next card states: "EXERCISE III—In this exercise, you will take the pencil in the same hand as in Exercise II, but you will place that hand through a rubber band. Always hold the pencil on the paper so a mark is made by it. I will place my finger through the other rubber band and then on the paper. *Watch the knot between the rubber bands, and always keep it on the same 'spot' or place on the* paper. The knot will move, but keep it in the same place. Move this card to the back of the stack when you understand."

While the volunteer is reading this card, I remove the paper and replace it with another sheet marked "III." Then I ask the familiar questions of the volunteer about his/her understanding. I show the volunteer how to hold the pencil and the end of the rubber band at the same time, and then I stand in front of the paper with my pencil, and begin my drawing.

My drawing for this exercise is quite different from that in the previous exercises. Again I have it on a card, and I look at it while drawing. I begin with a vertical line, then make a right angle with a line toward the volunteer, then make another right angle with a vertical line, and then a horizontal line. I follow these with several arcs, then a horizontal line toward the volunteer, ending with a vertical line and an s-shaped line. I remove my pencil from the paper when I finish my drawing, and I say to the volunteer: 'Please read the next card." I then ask the audience to answer the question for Demonstration III.

The next card states: "MOVE THIS CARD TO THE BACK OF THE STACK AND READ THE NEXT CARD."

The next card states: "EXERCISE IV—In this exercise, you will hold your pencil on the paper in the rubber band as you did in Exercise III. I will make a 'dot with a circle' on the paper. Your task is to keep the knot of the rubber bands exactly over the 'dot' inside of the circle, even when the knot moves. Always keep the knot over the 'dot.' Move this card to the back of the stack when you understand."

While the volunteer is reading this card, I remove the paper and replace it with another marked "IV." In addition, with the marking pen, I make a dot surrounded by a circle in the middle of the paper. Then I ask the same questions of the volunteer as before, regarding his/her understanding. (By this time, no one has ever had any problems following the instructions.) Then I show the volunteer again how to hold the rubber band and the pencil, and I proceed to make a drawing different from those in the previous exercises.

I begin this drawing with several arcs toward the volunteer, then I draw several arcs moving away from him/her. This set of lines is followed by a horizontal line away from the volunteer, a vertical line at a right angle, a horizontal line toward the volunteer at a right angle, and a short vertical line. When finished, I remove my pencil from the paper and ask the audience to answer the question for Demonstration IV; then I ask the volunteer to 'Please read the next card."

The next card states: "MOVE THIS CARD TO THE BACK OF THE STACK AND READ THE NEXT CARD."

The next card states: "EXERCISE V—In this exercise, you will hold the pencil in the rubber band as you did in Exercise IV. I will make the same 'dot with a circle' diagram as I did in the last exercise. This time, your task is different. Your *task is to keep the knot of the rubber bands exactly over the vertical line ABOVE the 'dot' even when the knot moves.* Always keep the knot over the place where the line and circle intersect *above* the 'dot.' Move this card to the back of the stack when you understand."

While the volunteer is reading this, I remove the paper and replace it with another, marked "V." On this paper, with the marking pen, I make a dot surrounded by a circle, with four small lines on the circle, 90 degrees apart from each other. This configuration looks like a target.

When the volunteer understands and is ready, I make the same drawing as I did for Exercise IV. When finished, I ask the audience to answer the question for Demonstration V, and then I say to the volunteer: 'Please read the next card."

The next card states: "MOVE THIS CARD TO THE BACK OF THE STACK AND READ THE NEXT CARD."

The next card states: "EXERCISE VI—In this exercise, you will hold the pencil in the rubber band as you did in Exercise V. I will then make the same 'dot with a circle' diagram as I did in the last exercise. This time, your task is different. *Your task is to keep the knot of the rubber bands inside of the circle even when the knot moves.* Always keep the knot within the circle. Move this card to the back of the stack when you understand."

While the volunteer is reading this card, I remove the paper and replace it with another, marked "VI," then I draw a "target" on the paper. When the volunteer understands and is ready, I make the same drawing as I did for Exercise V. When finished, I say to the volunteer: "Thank you, we have finished all of the exercises. You did very well." I ask the audience to answer the last question on their sheet for Demonstration VI.

When discussing this demonstration, I put the drawings for each exercise up in full view. I have the volunteer standing next to me at the front of the room. After I put up each drawing, I ask, 'What was the instruction he/she used to make this drawing?" and members of the audience are called upon to read their answers. I initially focus on the answers which are in error, and then I mention those which are correct. After getting a few answers which are in error, I ask the volunteer to read the actual instructions. As I discuss each drawing, I follow the same procedure. When I have finished discussing all of the exercises, I then use each exercise as an illustration of control theory.

The drawings for each exercise are designed to highlight different aspects of control theory. The drawings for Exercises I and II are the same, while the drawing for Exercise III is different, and those for Exercises IV, V, and VI are the same, but different from the others. Exercise I is supposed to demonstrate the "classical" stimulus-response (S-R) model, in that the volunteer imitated my drawing. But it should be pointed out to the audience (the volunteer will usually agree on these points) how slowly the volunteer moved, since he/she had to "see" my drawing before he/she could move. The volunteer usually agrees that this task was very difficult to accomplish. But the most important point to make is that the volunteer could not have done anything without the instruction "draw the *same* diagram that I am drawing." This instruction had to be used by the volunteer as a reference state to control his/her own conduct. This point can also be made for Exercise II, since the drawings are the same.

When comparing the drawings made in Exercises I and II, the audience might judge the reference states to be the same. It should be noted that the volunteer was better able to draw the one in Exercise H in (small) part because of previous experience, but that the instruction for the reference state was very specific. The point to be made is that two instances of similar behaviors can be generated with two different reference states, *but* that the different precisions of the instructions will make a difference in the two actions. Again, although the stimulus-response model seems to be relevant, it can be pointed out that it could not account for a similar behavior resulting from two different instructions; the S-R model would predict similar behavior, due to similar stimuli. These first two exercises, when explained with control theory principles, can counter most arguments for the stimulus-response approach.

The drawing for Exercise III has some lines similar to those in the drawings for Exercises I and II. This was done to illustrate that the volunteer will have a similar drawing even when the reference state, perceptual signals, and sensory signals are quite different. The volunteer could not have "carried over" the entire drawing from the previous exercises. It also can be pointed out that the volunteer's action was quite shaky, due in part to lack of specification of the "dot" and comparative sensory signals. This information can be used when this exercise is compared with the next one.

I made the drawings for Exercises IV, V, and VI the same for several reasons. First, these movements seem to work best for using rubber bands; sawtooth and vertical lines do not produce much movement by the volunteer. Second, I wanted to find out if audiences judge the reference signals for these three exercises to be the same from similar drawings and the target on the paper. Finally, I wanted to see how much "carry over" there might be from practice with different reference states. I use these exercises to show the effect of different instructions and reference states on the perceptions of the volunteer. I have never had a volunteer fail to report the importance of these differences.

The volunteer does *a* much better job with the drawing in Exercise IV, because there is an actual dot on the paper, rather than an "imagined" dot as in Exercise III. Some, but not many, members of an audience are able to distinguish between the instructions for Exercise III and those for Exercise IV. There is very little "carry over" for these drawings, because the volunteer is concentrating on the target instead of my drawing actions. But the instructions for Exercise VI provide a very interesting illustration of control theory.

Exercise VI specifies a reference state with a wide range of movement and very little possibility for error. If the volunteer follows the instructions properly, he/she will not have to move at all. I make my movements in such a way as to keep the knot within the circle at all times. The difference between the drawing for Exercise VI and those for the previous two exercises is usually quite noticeable. Many members of the audience say that the volunteer was confused or made an error. But this exercise is important to illustrate that a reference state (certainly at the higher levels) can be specified as a "range" where a variety of actions can occur before any negative feedback is noticed by the person.

You might think of other ways to treat these diagrams. Remember, even if the volunteer does not use the reference state that is specified in an exercise, he/she will use *some* reference state. In most instances, it is rather easy to determine what the volunteer controlled for in an exercise. I have rarely been wrong when I have guessed the reference state of a volunteer in these exercises.

These exercises, although clearly borrowed from Powers, have some distinct advantages over his for instructional purposes. Among the advantages are: (1) a record (trace) of the movement behavior of both the demonstrator and the volunteer, offering the possibility of precise

comparative measurement; (2) reference state instructions are known only to the demonstrator and the volunteer, not to the audience, which takes away the "obviousness" or "oh sure" audience response; and (3) the use of different exercises allows a comparative approach to control theory.

Possible Modifications

One could use a clear plastic board with clear plastic sheets for drawing, allowing the audience to see both the demonstrator and the volunteer from the front. Or a computer and a large screen could be used with a program which would make the drawings while the volunteer was following the instructions by using a joystick. This procedure would also allow for precise measurement of the volunteer's movements, with a printed record of the drawings. I am sure that other modifications could be made to increase the utility of these exercises.

Rick Marken: What we need to do as control theorists is develop more demonstrations of the *fact* of control.

I think that some of the best evidence of hierarchical organization in behavior comes from experiments showing one (or more) control systems operating within the time-frame of other control systems. This is the beauty of some of Bill Powers' "Portable Demonstrator" experiments. The simplest is when E's hand pushes down on S's hand to signal S to move his/her hand down from a fixed position. S's initial reaction is *always* an upward push before downward acceleration—the position control system reacts to the disturbance to position before the higher-order system can treat the disturbance as a signal to change the reference for the position control system.

I think it is important to get people to understand the *phenomenon* of control before pushing the *theory* that is designed to explain it. Telling psychologists that control theory is beautiful and powerful and revolutionary and humanistic and whatever just won't cut it. Theories are interesting to the extent that they explain what you want explained. And control theory explains control; so it would be most useful to show how control is involved in the behavior that psychologists are typically interested in. If psychologists are interested in cognition, then figure out demos that show how control is involved in cognition (we've done some of this, but not nearly enough). In some areas, like operant conditioning, the existence of control is fairly easy to demonstrate. In other areas (like language production), it might be more difficult to show how control is involved. But this must be the approach to promulgating control theory; because people cannot be expected to get interested in a theory if they have no idea what it's for. Indeed, I have more of a problem dealing with people who love control theory *qua* theory (they like the negative feedback and circular causation and all that) and have no idea what phenomenon the theory is designed to explain. I think there is a name for this latter approach to control theory; it's called "religion."

Gary Cziko: You can play with your tongue to see how speech disturbances are corrected. For example, keep the tip of your tongue against the inside of your bottom teeth and talk. I found this very easy to do with almost no sound distortion. Even sounds which normally require the tip of the tongue to move to the top of the mouth (/t/, /1/, /n/) are no problem—the middle of the tongue just comes up instead. For some reason, "gluing" the tip of the tongue against the bottom teeth is much harder, but still intelligible after a little practice. But watch yourself in the mirror if you want some laughs. The facial compensations that I use make me look like I'm snarling. Vowels are quite easy either way. Mustn't there be real-time perceptual control for this to work? Seems so to me.

Bill Powers: I get a very strong sense of the imagined auditory feedback by just mouthing "hello" without any sound (not breathing in or out). I don't actually hear sounds (no intensity or sensation) but the mouthed "hello" is still very plain to me as an imagined auditory experience. Does this work for anyone else? (Of course any other words will do—that's just the one I tried a moment ago.) It's the same imagined auditory experience I get from *reading* "hello." (Come to think of it, there's also an imagined kinesthetic experience in reading "hello" or "hello?" Even more so with "rouge," in French. Next thing, I'll be moving my lips when I read.)

Gary says: "For example, keep the tip of your tongue against the inside of your bottom teeth and talk. I found this very easy to do with almost no sound distortion." Brilliant! Yes, it's easy! There is some distortion of the final result, but I'll bet that if you used Crazy Glue to keep the tip of your tongue fastened to your bottom teeth for a month, you'd be talking essentially normally at the end. What you would be saying is another matter. Any volunteers?

Gary also says: "For some reason, 'gluing' the tip of the tongue against the bottom teeth is much harder, but still intelligible after a little practice." I presume you meant "upper teeth." *Yes,* it's harder—you have to use the lateral margins of the tongue to make a "t," and the vowels get distorted. But it's still quite intelligible.

I just love this kind of simple portable demonstration. It's a complete refutation of the idea that articulation consists of producing a preset pattern of motor outputs, and anyone can do it in two seconds. Absolutely ingenious, Gary.

In the Coin Game [a portable demonstration of PCT's Test for the Controlled Variable], the Subject lays out four coins on a table so that some pattern the subject has in mind is contained in the layout. The Experimenter disturbs the arrangement. If the Subject can no longer perceive the intended pattern, he or she moves one or more coins so that the pattern is again visible. If the Experimenter's move left the intended pattern still visible, the Subject just says "no error." The game is finished when the Experimenter can create disturbances that predictably call for corrective action, and disturbances that predictably result in "no error." The criteria can be adjusted as suits the players and their degree of skepticism.

An example of such a pattern is "at least one right angle." An experimenter unaccustomed to this test for the controlled variable might take half an hour to discover this pattern. Of course, it is possible to devise patterns that are undiscoverable, if you like wasting time. It is impossible to discover what pattern the Subject is controlling if the Subject keeps changing the reference pattern during the game. In such a case, only the Subject and God know what the controlled variable is at any moment. Even in this kind of case, however, the Experimenter can go up a level and approximate a higher-level controlled variable: "Either you've got an extremely obscure pattern in mind, or you keep changing it every time I get close." Playing this game will teach one a lot about how to find controlled variables, and the pitfalls of assuming that "insights" into what another person is doing have any relevance. The Test is all about eliminating wrong hypotheses. When you get systematic about doing that, you can guess very efficiently. If you get hung up on a clever hypothesis, you might take forever to find the right controlled variable.

It is also instructive to discover that a perfectly good verbal definition of a controlled variable that passes every test is not the one the subject used. You say, "It's a zig-zag!" The subject says, "No, it's an 'N' on its side." Of course, it's *really* a "Z." In *fact,* it's the perception that the Subject intends to reproduce, not its description.

Dag Forssell: I would like to build a simple control system demo to fit in a briefcase and be powered by regular house current. How about a joystick-controlled rheostat providing a reference signal to represent the aiming of a cannon? I would like the reference signal to be visible, perhaps in the form of a voltmeter. The gun barrel would be hooked to a rheostat and the signal made visible in the same way; the error signal thus created also would be made visible. An amplifier (with output visible on a watt or amp meter or perhaps through the glow of a light-bulb or tone (pitch, volume) of a small speaker) would drive a motor left or right with proper drive of the gun barrel.

It occurs to me that if the perceptual, reference, and error signals were shown on strings of light-emitting diodes, then the subtraction could be made graphically quite visible.

One of the things I want to come alive is the (rapid) conformance of the perceptual signal in response to a rapidly changing reference signal.

Surely some of the people on the net have thought about this and can suggest designs, components, and sources of supply. Please give me some ideas.

Later, we could create visible conflict between two units and also build a hierarchy.

Wayne Hershberger: Dag, you might want to consider the control systems used to pilot radio-controlled model aircraft. They are battery powered. You would not need the radio, just a servo and a device to provide the servo's reference input. You can purchase a servo in an R/C hobby shop for about $25. You might need to special-order the servo driver. Ace R/C Inc. (Box 511, 116 W. 19th St., Higginsville, MO 64037; phone 816-584-7121) calls their driver a "Servo Cycle." It costs about $30. Ace R/C sells their products as kits or as assembled units. If you were to get everything assembled from Ace R/C, I estimate that it would cost you about $85. You would need: Ace R/C Bantam servo: $26.45; Servo Cycle (with connectors for Ace servo): $32.95; nicad battery pack (4.8 volts DC, 500 milliamps): approximately $20.00; battery charger for nicads: approximately $8.00; plus postage: $3.00.

Gary Cziko: Wayne, could you tell me what such a system as you outline for Dag would actually do?

Wayne Hershberger: Gary, the radio control systems employ *position* servos. Each servo, encased in a small box displacing less than one cubic inch, controls the position of a small arm extending from the case. Four servos are customarily used in model aircraft, with the output arm of each servo linked to one of the four flight controls: throttle, rudder, elevator, and ailerons.

Disassembling a position servo, one finds three components:

(1) an effector—a geared electric motor that drives the output arm about its axis of rotation,

(2) a position sensor—a potentiometer whose wiper is attached to the output arm's axle of rotation, and

(3) a circuit board comprising electronic components that receive the reference position (signal from a transmitter), compare that reference position with the sensed position, and amplify the error signal to drive the motor.

By wiring electrical meters into the circuit, one can see that the servo draws very little current (about 10 milliamps) when its arm is idling in the reference position, but over 100 milliamps when a load tries to displace the arm.

Gary Cziko: Wayne, are the R/C servos strong enough to interact with a human? That is, could I grab hold of the arm (if only delicately with two fingers) and disturb it and feel it fighting back? For a good demo, it should have enough loop gain and "muscle" so that I can feel it resisting, but not so much so that I can't even budge the arm.

Of course, it would be nice to have it move in more than one dimension to make it seem even more alive.

Bill Powers: Wayne, I've sent for the Ace catalogue. I had always thought that those servos are just up-center-down or on-off. There should be all kinds of neat demonstrations we can come up with using a pre-packaged position servo as the core device. You could use two of them to play the rubber-band game in one dimension. Maybe you could make a balsa-wood jointed arm. More toys!

Dag Forssell: Thanks, Wayne, for the info on servos. I had in fact visited a hobby store and concluded that these servos were stepping motors and not suitable for demo-building. Based on your post, I sent for the Ace catalog. I hope that the ready-made, inexpensive, and certainly compact servos offered by Ace will be suitable.

Bill Powers: Yes, if those little servos can produce enough output force, they will certainly make suitcase demonstrators much easier to achieve. But let's not give up on finding components for designing our own.

Joel Judd: [Excerpts from a manuscript titled "Second Language Acquisition as the Control of Perception":]

I need a volunteer for a harmless demonstration. [What follows is the rubberband demo—a volunteer is asked to hook one finger into a loop, and I hook one of mine into the other. I tell the volunteer to move as he sees fit, avoiding exaggerated or extreme movements, and when he knows what I am doing, to stop.]

What do you observe? (Responses.) You will notice that as long as the rubberband was taut, I moved my finger in a coordinated fashion along with the volunteer's. If you think I was simply mirroring his movements, what would you say if I covered up his hand or used a rubber band with three loops? One could run a correlation on the movements of our two fingers, and find a coefficient somewhere in the neighborhood of .95, or even better. This is an incredible correlation for the behavioral sciences. But what does it tell you about me? It tells you that our finger movements correlated. Somehow that knowledge is not very satisfying. Why was I moving the rubberband as I did? What if I tell you that I was trying to keep the knot over a spot on the board? Now what do you know? Does that change your perspective on this experiment? Now you know *why* I behaved as I did—why, as long as the game was on, I counteracted moves made by the volunteer. Did the volunteer *cause* me to move as I did? Well, yes and no. Yes, in the sense that he caused disturbances to my goal of keeping the knot over the mark. No, in the sense that if "keep the knot over the mark" was not my goal, his movements would not have required any action on my part. I was concerned with maintaining a particular relationship between the knot and the mark on the board. He made it difficult for me to achieve that goal, so I had to *do* something to overcome the disturbances.

The principles I would like to emphasize from this demonstration at this time are three. First, I had a purpose in playing the game. My purpose was to maintain a close match between my goal ("keep knot over mark") and what I perceived with respect to that goal (through vision). My behavior was purposeful. Second, while the behavior had a purpose, it was not controlled, it was only incidental. It was the result of a comparison between my goal and my perceptions. If my perception was "knot over the mark," then little behavior was required; if the perception was "knot far from mark," marked behavior resulted. But what I was controlling is the third principle: perceptual inputs. What I wanted was to perceive the knot over the mark. My observed behavior was only one of several ways that I could have achieved the desired perception (I could move just my finger, my whole body, or even the chalkboard itself). I did not concentrate on the volunteer's finger movements, or on my own. What I did was to check my actual perceptions against what I *wanted* to perceive.

Given such an interpretation of events, what can we predict? We can say that I will do what is necessary and possible for me to do in order to maintain my internal reference or goal. You cannot predict what *exact* physical behavior I am going to exhibit; you can predict (knowing my goal) that I will do *something* to maintain that goal in the face of disturbances.

We saw in the demonstration how the observation of my behavior did not lead you to understand why I was doing what I did (or perhaps you hazarded a correct guess). It was obvious that I was moving

in concert with the volunteer, but even noting this (and correlating an extremely high correlation), you learned nothing about my purposes. There are only two ways in which you could find that out. One is to ask me. Of course, I can lie or mislead you, but it is possible to find out one's goals by asking what they are. How often this possibility is overlooked in the social sciences. The other, "purer" way to determine goals is to hypothesize what they are, then apply systematic disturbances to the organism and see if it tries to overcome them—to obtain the goal even though unpredictable obstacles threaten to prevent its attainment. This description of behavior is known as perceptual control theory.

Bill Powers: Joel, if the goal is "keep the knot over the mark," there are several ways to achieve it:

1. Drive a nail through the knot into the mark (no control needed).
2. Shoot a curare arrow into the other person so he/she will stop disturbing the knot.
3. Employ a visual-motor control system and give it the reference-image or goal that we describe as "knot over mark."

Just quibbling.

Gary Cziko: Stand up. Close one eye. Then push on the side of the open eye with your finger. The perceived motion makes you feel a bit unsteady, doesn't it? Now, walk a straight line with one eye open and then, while walking, disturb the eyeball with your finger (give a nice steady push and *hold* it there). You might find that you can no longer walk a straight line. If you are pushing on the right side of your right eyeball, you will likely veer to the right (because your brain thinks your body is leaning to the left, it will "compensate" and send you off to the right). I actually feel that I will fall down if I don't stop walking and don't stop pushing, although some of my colleagues here can continue quite well, but still report how odd it feels.

Let's get some data from others out there on the latest PCT portable demonstration.

Rick Marken: Gary, I damn near knocked myself over when I was walking. What a splendid demo. Everyone should try it—but be *careful*. I am truly amazed at the power of this little demo. I would have imagined that walking could be carried out quite well even if the visual input were disturbed—but no way!

Gary Cziko: To complete the eye-pushing demo, three more steps should be added. First, reverse eyes to see how you now stagger to the other side, to show that this is a systematic effect. Second, walk a straight line with eyes closed and then push on the eyeball. No problem (except for the normal "drift" of trying to walk straight with no visual feedback). Therefore, it is not just pushing on the eyeball which causes the instability. Finally, walk with the one eye open while making lots of saccades. No problem. Therefore, it is not just moving images on the retina which causes the problem. The problem is caused only by having the retinal image disturbed by some "outside" factor.

For even more excitement, try this while riding a bicycle (now I know why I wear a helmet). Doing it while driving a car is also amusing (make sure first that you are insured against PCT demos). Any airplane or space shuttle pilots out there?

Let's see what my list of portable PCT demos looks like now:

1. Powers' classic rubber band demo of keeping the knot of two knotted rubber bands (elastics) over some fixed point with the subject's (controller's) finger in one loop and the disturber's finger in the other. Demonstrator can be either subject or disturber.

2. The demonstration of the levels of the hierarchy involving hand movements.

3. Speaking while keeping the tongue in some relatively fixed position, e.g., tip touching upper or lower teeth.

4. Eyeball pushing and walking.

5. Finger pointing. Close one eye. Reach out with one arm with index finger extended and put the finger where it "touches" some distant object. Keep it there and watch how your arm seems to take on a life of its own as it actively compensates for disturbances caused by breathing, heartbeats, body movement, muscle fatigue, etc. Then decide to move your finger a certain distance above or below where it was. When I change the target, I get a feeling of "willing" which is very different from just maintaining the finger in a certain spot. This lower-level maintenance of finger position feels like it is running on a sort of automatic pilot. I do not feel as if I am actually in control of the little movements needed to maintain the finger on the target. And, of course, in a sense I am not. All (higher-level) "I" can do is specify the reference signal for the lower systems, and then they do their jobs without any further assistance from the higher levels.

Are there other portable demos that I have missed? If we get enough of these, maybe PCT will start taking over classic psychology in introductory courses just because it will be so much more fun!

Martin Taylor: The eye-push demo is very reminiscent of one of psychologist J.G. Taylor's demos. Gary's mention of bicycling brought it to mind. Taylor claimed (and demonstrated) that distortions of vision are corrected if and only if the distorted components are affected by behavior as part of the feedback loop involving that behavior. One of the

distortions was to wear spectacles with some kind of prism, such as inverting spectacles that interchanged left with right, or up with down, or merely displaced things (say) 20 degrees to the left. He mentions in his book the experience of seeing a narrow strip of floor in front of him (on which he would be walking) seem perfectly normal and flat, while on either side (irrelevant to his walking) the floor seemed to be sloping. At the same time, the surface of a table showed no levelling effect during the 13 days of the experiment. At the table, he indulged in no control behavior (he says) that involved gravity, and thus there was no mechanism for the distortion to correct itself (no feedback to test reality).

But more dramatic is a film of Seymour Papert (of MIT) learning to ride a bicycle while wearing left-right inverting spectacles. At first, as soon as he put on the spectacles, Papert would crash—he applied the wrong corrections. After a while, he could stay on, albeit wobbly. But then he would crash when he took the spectacles off while riding. After more training, he could put the spectacles on and take them off quite freely, while maintaining control. But then comes the kicker—Taylor took the prisms out of the spectacles (the frames were quite heavy), or perhaps he substituted non-inverting prisms. At any rate, Papert's view of the world was normal with or without the spectacles. But on putting the spectacles on, he crashed as in phase one. Taylor took this to mean that the totality of sensation involved in a situation was all part of the control system, and this included having reorganized the system to include what amounted to a switch based on the weight of the spectacles on the nose (Papert "knew" what was in the spectacles, but his fully trained control systems didn't).

Papert, by the way, contributed a mathematical appendix to a chapter of Taylor's book, so he was familiar with the theory, though it is hard to see how this could have contributed to the effects. I certainly would not have allowed myself to keep falling off a bike to support someone else's theory of perception!

Gary Cziko: Here's a demo on "apparent social control."

I ask the audience to close one eye and reach out with one hand with finger extended to "touch" my hand. I then move my hand around and they all follow quite nicely. It almost looks as if my hand is connected to all of their hands (like puppets on strings). I then ask them to raise their hands as high as they can. Then, for contrast, I ask them to shout "Go Illini" as loudly as they can at the count of three (the Illini is the name given to football/basketball teams at my university). I then count to three and... total silence.

So they will do as I ask only if *not* doing so would create an error signal (they want to be cooperative, and audiences are certainly used to moving their hands around). But shouting as loud as you can in a classroom would create more of an error signal (or I suppose I should say an error at a higher level) than the error signal created by not following my shouting request. (I wonder what a roomful of Republicans would do if asked to shout 'long live George Bush!")

Bill Powers: Gary, I love the social-control demo. We need a version of it for the portable demonstration collection. How about giving a pin to someone, and giving him/her instructions to push it all the way into a chair cushion, a rug, and the palm of his or her hand?

Gary Cziko: In my continuing search for portable demonstrations of perceptual control theory in action, I was playing with handwriting this morning. Here are some things to try. I will assume that you all write with your right hand. Lefties need to substitute left for right and vice versa.

1. Write a sentence with your right hand. Notice how easy and quickly you can do this.

2. Write a sentence with your left hand. Notice how much more difficult and slower this is. Nonetheless, if you take your time, you can probably still write quite legibly (I can, anyway).

3. Now write another sentence with your left hand, but write *backwards, that* is, from right to left with the letters laterally transposed (the way DaVinci wrote, I'm told). I found this to be even more difficult than 2. From a motor perspective, we might expect it to be easier, since the actions are the mirror image of what is done with the right hand. But I think it is harder since we simply don't have a good idea of what reversed writing should *look* like. I found that I would often "freeze" at the beginning of a word. I did not "freeze" when I wrote in the normal direction with my left hand.

4. Now try writing some individual words (still lefthanded and backwards and laterally reversed), but now write normally with the right hand the same words simultaneously. You should be able to write quite easily backwards this way, since now you have a "motor" reference level which you can use for the left hand.

5. After a little bit of this, you might find that you can just "imagine" the right hand writing normally and then write backwards with the left hand without much difficulty. But you still probably don't have a good idea of what your backwards writing should look like. So even though you're looking at your left hand write backwards, it seems (to me anyway) that the visual feedback is not used very much. It is sort of like writing normally but with your eyes closed.

6. Now, you can switch back and forth between forwards and backwards writing with your left hand. And it feels quite different. Writing normally (left to right) with my left hand, I am using primarily visual feedback, since I know what it is supposed to look like. It is quite slow, but I can make it look quite good if I take my time. Writing backwards with my left hand (after using the simultaneous right hand trick) is much faster and feels just fine, but looks pretty awful since I am using proprioceptive feedback, not visual.

In addition to its use as a demonstration of some key PCT ideas, handwriting might be a good way to do research on reorganization. Only pen and paper are needed, and the subject leaves a permanent record of his or her behavior, with no need for fancy computers and C compilers. In addition to the above tasks, you can see reorganization in action by holding a mirror at the head of your paper and writing so it looks normal in the mirror. This makes you write upside down. This is maddeningly difficult. You can see your runaway streaks of positive feedback as you try to make *a* line go down and it keeps ascending faster and faster the harder you try to get it to descend. This reminds me of Martin Taylor's account of Seymour Papert learning to ride a bicycle wearing reversing prisms as eyeglasses, but it doesn't hurt nearly so much to make a mistake.

Joel Judd: I have wished many times for a "real" example of linguistic control to give to the audience (or perform *on* the audience). Nothing fancy, just something that most would agree demonstrates control. I have been looking for something along the lines of James' (18%) Romeo and Juliet description for purposeful behavior. Is there something in some netter's experience which might qualify as such a linguistic example? Perhaps some child's conversation, or something literary? Once one believes in control of perception, examples are all around—but what about some attention-grabber that would be difficult to explain away in other than PCT terms? Thanks.

Bill Powers: Joel, I don't have any clever examples of linguistic control, but maybe a halting attempt will suggest something to you or others. We want to demonstrate perception, reference condition, error, and action to correct the error. So how about using some phrase that everyone knows and putting errors into it: "Now is the time for all men good to come to the country of their aid."

This demonstrates a number of principles.

1. Even though this is not a proper sentence, it will be recognized as "nearly" right. This shows that perception even of sentences has an underlying continuum of variation that we express with terms like nearly correct, pretty close, not so close, pretty bad, and awful.

2. The fact that you know the sentence isn't right implies that you're comparing it with some standard. The "right" sentence isn't (for most

people) vividly in awareness at the moment the mistakes are detected, but some criterion for correctness must be there in some active form.

Given a sense of error in the sentence, you can then point to the causes of the error: in this case, the placement or ordering of the words in the sentence. So to correct the sentence as a whole perception, it's necessary to alter some of the elements of which this whole perception is made: that shows hierarchical control even without trying to pin down just what the levels are.

There are many other ways for this sentence to be in error besides just the ordering of words. If it starts with "Now is time for...," a word has been omitted. If it starts with "Now was...," a wrong tense is present. "Now is the tome..." contains the right word misspelled (interesting, because "tome" is a perfectly good word, but not in this sentence, so we recognize it as almost "time," but with a letter error in it). Each kind of error exposes some aspect of the sentence that you are monitoring for correctness. It also shows that you hold in memory various kinds of criteria for correctness and apply them all in parallel. If the sentence starts with "Now is tome the for...," two errors of different kinds are sensed simultaneously.

Examples of control while constructing a sentence can sometimes be seen when the listener indicates an unwanted response before the sentence has been finished. You say, "Excuse me, could you tell me where, oh, sorry..." as you realize that the person is smiling, shaking her head, and tapping an ear to indicate deafness. It's too late to edit out the futile "oh, sorry," but your first impulse is to do so. Maybe you could set up some kind of role-play to show how this editing on the fly occurs.

Another simple-minded example is like the method Dick Robertson and David Goldstein used to demonstrate control of self-concept. They tried some complicated questionnaire methods which gave the usual equivocal results, then finally decided just to wait for the person to utter a self-description and reply by contradicting it: "No, you're not like that," or something of the sort. I believe that 25 of the 26 people tested responded by saying something to oppose that disturbance.

If you ask someone to explain how a fan moves air (or something simple like that), the person hearing the explanation could respond by saying, "Oh, I see, the air going through the fan makes the blades turn," or something like that. PCT would predict that this unwanted understanding of the communication will call for more communication that is aimed at getting a more correct statement from the other person, and that when the reflected understanding is judged correct, efforts to change the other person's understanding will stop.

Bruce Nevin: Control theory demos typically concern behavior that is closely matched to environmental variables (I hope my epistemological looseness here is forgivable). There isn't much room for variation in a forehand smash in tennis: either you strike the ball so as to put it over the net to the chosen area of the court with desired speed and spin, or you don't.

In the range of variation that an outside observer might judge to be tolerable, one person will control for a more restricted range, and another player might control for a different part of the total range. An opponent or spectator might label person A an "aggressive" player and person B more "laid back," and so on. On a different occasion, player A might "let up" a bit so as to accommodate a less skilled player.

Social behavior of the second sort exploits the range of tolerable variability of control for purposes of constructing a persona and presenting it to others. A self-image.

It is this that is one of the principal motivators of language change and variability.

So we control for what we are paying attention to, and ECS daemons with alternative interpretations don't get heard. We only notice ambiguity when we are paying attention to ambiguity itself, or more commonly when an expected agreement is not reached (but even then far from all of the time). Looked at this way, I doubt there is anything that is not ambiguous.

Somewhere in there is the germ of a demo. A problem with any demo is that our control of language is so exuberantly pandaemonic (massively parallel, as they like to say). One has to account for other parallel threads of interpretation without being able, as in traditional laboratory protocols, to eliminate them from the experimental setting. Some parallel threads are redundant (feature, segment, semisyllable, autosegment, stress group, etc. in phonology) or partially so, and some are competitive (ambiguity). Such an accounting does not make for a succinct, crisp, yet convincing demo. Too many audience yeah-buts are possible. Maybe it's in answering the yeah-buts that you show how parallel pandaemonic control frames the demo.

Rick Marken: It is very easy to demonstrate the illusion of control—just ask someone to track your moving finger with theirs. If they are willing to do that, then you can control the position of their finger—you will experience control and your actions will look like control to an observer. Interestingly, the "subject" will not feel controlled until your disturbances (finger movements) require action that produces a perception conflicting with their ability to control other variables (Bill's example of moving the subject's finger close to a hot soldering iron comes to mind). In such a case, the subject will probably notice your disturbance as an effort to control him/her, and you will notice a loss of control—especially if you really want the subject's finger to be close to the soldering iron.

Chuck Tucker: I have thought some about how to devise rubber band demos for more than two persons. Actually, the ideas for these came from a discussion at a CSG meeting which was given by Ed Ford and his friend, Jim Soldani, who used such demos for people in organizations to show how conflict can arise if people don't understand one another's purpose. These demos can be used to illustrate PCT (obviously), as well as what many have called "social constraints" or "social structures" (which are, in my view, mainly arrangements that someone devises to make it extremely difficult for a person to accomplish his or her purpose), in addition to illustrating conflict and conflict resolution. In these demos, you can give the instructions in verbal, written, or graphic form, or all of these forms, you can have the participants talk to each other or not, you can be one of the participants, and you can try to restrict the "sensory input" of the participants by using screens, blindfolds, or heavy gloves on their hands, or by having them hold the rubber band with a hook instead of with their fingers directly. All of these variations I can see as attempts to illustrate different aspects of the PCT model.

I'll illustrate with a three person demo. Three rubber bands are each knotted on a fourth rubber band (opened up to make a circle) at equal distances from each other. Using the picture of a circle with 360 degrees and treating a line intersecting the circle anywhere as 0 degrees, tie the three rubber bands at 0, 120, and 240 degrees, equally dividing the fourth rubber band into thirds. Now have the participants make triangles of various shapes, have one participant refuse to make a triangle with the others, have them make a triangle and tell them, "Now, hold that position for 10 seconds and remember how that felt, because I will ask you to do it again without being able to see what you are doing." Then blindfold each of them, place the rubber bands on their fingers, and ask them, "Make that triangle again." (Take a picture of each performance, so you can compare them.) The parameters of the activity can be changed by tying three rubber bands at 0, 160, and 200 degrees, then at 160, 180, and 200 degrees, and so on. When these "structural conditions" are set, the types of shapes that can be made will be restricted *unless* the participants devise ways (like crossing over each other's rubber band) to make the shapes; if they are all required to stay on the same plane (another structural condition), then the shapes they can make will be limited.

Now expand this to four participants and start with the simple setup of four rubber bands at 0, 90, 180, and 270 degrees around a fourth rubber band to make a square or other square-like shapes; ask the participants to make a triangle; to make a circle; to make a hexagon; and on and on and on. Do five participants, six, seven, and on and on as long as you have rubber bands.

I believe that these demos are not only very useful in illustrating the ideas of PCT, but are enjoyable and memorable for the participants.

Bill Powers: In the Coin Game, The Test for the Controlled Variable is done with the subject there. The actions of the subject can be perceived as affecting the environment in many ways, and objectively have many different effects on objects, relationships, etc. in the environment. The question is which, if any, of these effects of the subject's actions is under control. The experimenter devises a disturbance that will alter one of those effects. If the effect changes—if the subject does not change the action in a way that prevents the change from taking place—then that effect of the action is not under control.

Use four coins (same or different, as you please). Two people play, an Experimenter and a Subject. The Subject places the coins on a table such that they exemplify a pattern or condition that the subject has in mind. The Subject privately writes down this reference pattern on a piece of paper, and hides it. The Experimenter is to discover what the controlled pattern is, by means of disturbing the arrangement of the coins.

The rules are as follows. One round of the game starts with the Experimenter doing something that alters the arrangement of coins on the table. The Subject looks at the new arrangement, and if the target pattern can still be seen, says, "No error." If the pattern now differs from the target pattern, the Subject makes any rearrangement of the coins required so that the perceived pattern once again matches the target pattern. After either a "no error" response or a corrective move, it is the Experimenter's turn again.

The game ends when the experimenter can demonstrate three different moves predicted to produce a "no error" response, and three different moves predicted to produce a correction. Then the subject displays the written description of the reference condition. No verbal communication except the words "no error" takes place during the game.

You might think at first that it will be easy for the Experimenter to discover the pattern, and compensate by choosing (as Subject) a complex reference condition. I advise choosing a simple reference condition if you want the game to finish in under half an hour, or not be abandoned.

This game illustrates all facets of The Test for the Controlled Variable. Clark McPhail has been using it to teach The Test (he sent me copies of experimental reports by students—wonderful reading, especially the comment by one student that he really admired sociologists for being able to use The Test in their work, because it is so complex).

Martin Taylor: I appreciate the Coin Game. My thesis supervisor did quite extensive studies of this kind of game, perhaps not structured exactly the same way, but very like it. He was studying perception, not control, but the issue is also of determining by such trials the nature of prespecified relationships. I think it was a popular experimental paradigm at the time. The effects of interactions on the kinds of relations that were readily detected or were detected only with difficulty was the point at issue. He *came* up with the notion of integral and separable perceptual dimensions, which seem to be quite important. I can't give any specific references, but if you want to search for them, look for W.R. Garner in the late '50s or '60s.

Gary Cziko: While I don't think that I will ever design a control system, I couldn't help noticing a "digital proportional radio control system" for $49 in *a* local hobby shop and so figured that this might be a way for me to at least interact with one (an artificial one, I mean; I already have lots of experience with the living kind).

This is a Hitec "Challenger 260" two-channel system that includes a pistol grip transmitter (reference-level manipulator) and receiver connected to two servomechanisms (it is made for controlling speed and direction of model boats and cars). Pulling the trigger and turning the wheel on the transmitter move wheels on the two servos. I replaced the wheels with two four-legged spiders that came with the kit and attached rubber bands to one arm on each.

With either the transmitter or receiver turned off, one can easily move the spiders for a total range of about 90 degrees (it's a bit stiff and I don't know how good this is for the servos). But with the both transmitter and receiver on, they really fight to respect their position reference levels. You can feel them vibrate and fight back when you try to disturb then. While it *is* possible to overpower them, I am quite impressed by how strong and stubborn the two little servos really are—the more you try to push them around, the more they push right back at you (very much like most people I know!).

The rubber band is a *nice way* to add disturbances. I can ask someone to pull the rubber band hard any which way, and it makes virtually no difference to the position or pattern of movement that I am sending with the transmitter. This is a very nice demonstration of why controlling reference levels is the way to go. I let the servo control system worry about the rubber band disturber, and it makes no difference to me, the upper-level reference signal supplier.

While Powers' and Marken's computer demos are great, there is something to be said for the real physical interaction that these servos provide. Also, they provide an easy way to give my students hands-on artificial control system experience. Highly recommended.

Dag Forssell: Imagine that you are playing the rubber band demonstration with a strong machine, programmed to go through a set pattern of motion. You have no difficulty, since the machine only influences the position of the knot, through the tension in the rubber band.

Now, "tightly couple" the knot to the machine by substituting a stick (or a rope, as I like to do when demonstrating conflict between two "pullers"). If you still are connected to the knot by a rubber band on your side, you will pull in vain. If the stick is extended to your hand, you will be pulled along, powerless to do otherwise. You are being controlled by the machine with overwhelming physical force, the only way Bill says you can be controlled.

I believe it is important to remember one of the hallmarks of control systems: amplification. This term does not communicate well. I am shifting my language to "the direction of resources" or something like that, with emphasis on resources. Your heating system at home opens a valve to release (and ignite) a stream of natural gas. The stream is not finely calibrated, but it has a powerful influence on the air temperature. If you have an air conditioner working at the same time, set at 68 degrees, while the heater is set at 75 degrees, the two will pull (with tight coupling) on the air-temperature knot with as much influence as each is capable of. If the gas line has the capability to release more resources to raise the temperature than the air conditioner has resources to lower it, then the air temperature will stay at 75 degrees.

The rubber band is such a marvelous tool, because it shows influence without tight coupling. Try the demonstration with a rope and two dots; one dot towards the left as a target for the left person, and one a little to the right (one foot apart if you are at a blackboard with a four-foot rope, which works best, one inch apart if you are on a paper with a short string) as a target for the person on the right. See which person is willing to pull hardest. This person will pull the knot to his/her dot and keep it there. This illustrates the heater/air conditioner conflict.

On another subject, suppose you pull on your end of the bands to keep the knot over the dot. *You* control! I disturb the position of the knot by moving my end. I provide a stimulus, and you respond. My disturbance is a property of the environment, from your point of view. (I can represent any kind of machine or natural effect disturbing my end; I am not trying to control.) So is the quality of the rubber band which converts your action (and my disturbance) into an influence on the knot. As long as you do control, your action will be *what it has to be* to keep the knot over the dot.

Your action is 100% determined by the disturbance and the nature of the rubber band, which are *properties of the environment*. The only requirement is that you do control somehow. The rubber band experiment illustrates the fact that you *do* control. It tells you *nothing* about how you are organized inside to accomplish this control. Therefore, what you see (your erratic movements) is due to properties of the environment, not of the organism (you). This is most clear when you do the rubber band exercise slowly, allowing nearly perfect control. The knot stays steady over the dot, and your actions are perfect mirror images of the disturbance.

In Phil Runkel's book is an excellent, detailed description of the rubber band experiment that is *more instructive* than the way I was introduced to it. You invite a friend to experiment *on you.* "You are the experimenter. Move your finger as you like. Watch what I do. When you can explain what is causing me to do what I do, let me know."

Phil spells out the typical suggestions of friends. I have confirmed this. Saturday, I had a group of six, with no notions of PCT. I used an easel with the above instruction printed in the center. My rubber bands had a yellow ping pong ball over the knot, to make it visible at a distance. I kept the ping pong ball over one letter. All I got was that I was mirroring the experimenter. Of course, the experimenter "causes" me to do what I do. I kept telling them that that was not the cause and challenging them to come up with a better explanation. No luck.

It is true that people cannot see control even when it is staring them in the face.

Starting the experiment this way makes the paradigm shift stand out. You can point out that an absence of a point of view makes it impossible to see the phenomenon. Your ignorance makes you blind—literally!

Only *after* this sequence do I experiment on my friend by asking him/her to keep the knot over the dot. Later, one can point out that the better the control, the less exciting the appearance. Good control is invisible because nothing happens.

Gary Cziko: Dag, another variation of the rubber band demo I often use when presenting to a group is to ask for a volunteer. I then whisper to the volunteer, "Keep the knot over the dot (or other landmark)," and then I disturb. The audience has to figure out what the subject is doing and make guesses, but the *subject* responds as to whether the guess is right or wrong. I can *then* even have someone in the audience be the experimenter.

It's amazing how difficult it is for some people to find the controlled variable. It seems the more psychology one knows, the *less* likely one is to find the answer. That's understandable. But why the very sharp control systems engineer I tried it on gave up after a few minutes remains a mystery to me.

Bill Powers: Try this: Knot three rubber bands together at a common point. Do the experiment on a large sheet of paper or against a black-

board. Use three positions of the disturbing end of the rubber band measured relative to the known target position of the knot: large, medium, and small distance from the knot. Make these positions only about an inch different from one to the next. The positions can be premarked on the paper or blackboard. The experimenter pulls back to each position and records where the subject's finger goes, marking the positions on the paper or blackboard.

In Experiment 1, the disturber loops two of the rubber bands around his finger, leaving one for the subject.

In Experiment 2, disturber and subject get one rubber band each, the third one just dangling.

In Experiment 3, the subject loops the finger through two of the rubber bands, leaving one for the disturber.

To distinguish the data for the runs, label the subject's finger position marks as la/lb/1c, 2a/2b/2c, and 3a/3b/3c.

In all three experiments, the size and direction of the disturbance is the same small, medium, or large amount. The subject, however, will respond very differently in the three experiments, as can easily be seen during the experiment and by measurements with a ruler afterward.

I'm not going to tell what happens. You should be able to reason it out from elementary PCT principles, then verify that your prediction is quantitatively correct, using the method outlined above.

If you get the right answer, you will realize that you don't even need a subject for this experiment: you can play both parts. All subjects who keep the knot over the dot will behave in exactly the same ways in each of the three experiments. These measurements are not measuring any properties of the subjects. I leave it to the advanced student to say what they are measuring.

Rick Marken: Dag does an excellent job with his wonderful variations of the rubber band demo, showing that extraordinarily complex "behavior" seems to be going on when people are doing nothing more than trying to perceive a simple relationship between configurations—"knot on dot."

I believe that one of the problems confronted by those of us who are trying to "sell" PCT with models and demos is the same as the problem we confront when trying to point out to psychologists that there is a phenomenon (called "control") that is going on in front of their eyes that they have not taken into consideration in their attempts to understand mind and behavior. The problem is that the disturbances, constraints, and calibration problems that make control necessary and obvious are simply invisible. When you point your finger at a target, the pointing just seems to happen; the fact that you can repeat this pointing with great precision seems completely unimpressive. You just point at the target again and again. Disturbances (such as changes in your orientation with respect to gravity), constraints (such as the fixed length of the segments of the arm), and calibration problems (like the fact that a neural signal never produces exactly the same amount of muscle tension) go completely unnoticed. When disturbances are visible (such as movements of the target), they look like stimuli guiding the response. It is, thus, very easy for those who want to, to ignore control.

Basically, the problem with demonstrating models of control is the same as the problem of seeing control in normally occurring behavior; what is most amazing about control is what you can't see. And you can't see the amazing aspect of control (disturbance resistance, constraint satisfaction, and calibration compensation) because control itself prevents these things from having any noticeable effect. So it is the fault of control itself that the process of control is invisible.

In order to see control, you must be the agent of disturbance; you must be able to do something that you know should have an effect on a variable if it were not under control. If you think a person is controlling the position of a limb, then you can literally "push" on the limb to see if the push has the expected effect (movement of the limb). This "test" must be done carefully—not too much disturbance (control systems have limits to the amount of output they can produce), with an appreciation that control of some variables occurs more slowly than others (so the disturbance might seem to have an effect but will be slowly cancelled if there is control).

I don't know if there is any really dramatic way to show control; we keep trying, but we obviously haven't found a real "grabber" that would get psychologists to throw up their hands en masse and cry, "Oy vay, I've been missing the point for my whole career; people don't respond to stimuli or generate outputs—they control! Now I have to abandon all my work and start studying control. Damn, how did I miss that—I guess that guy Powers wasn't just a stubborn, contrary, radical outsider after all."

Dag Forssell: Rick, we do not need more startling demonstrations. PCT tells us that all action is initiated by error signals. What we need is to address the error signals that lurk out there in people. A synonym for error signal is dissatisfaction. We need to reach people who are dissatisfied with what they can accomplish, people with a yearning for something better. A better way to deal with each other.

A dissatisfied person will be open to suggestions and interested in trying a different solution. Much of the debate on this net addresses people (directly and indirectly) who are perfectly satisfied with what they know, proud of it and ready to defend it. Forget it. Ask people what problem they are anxious to solve. Ask if they are willing to think

for themselves and evaluate an alternative. When people refer to authorities, they are not prepared to think for themselves. PCT does not need anything more than a student who is willing to think for himself/herself and make the effort to understand the evidence.

Our challenge is to tell our story so that people become aware of the error signals they frequently deal with, and understand that we have a permanent solution they might like if they spend a little time looking at it.

Gary Cziko: Rick Marken should already know that my all-time favorite experimental report in PCT is his "The Cause of Control Movements in a Tracking Task" (which is included in his book *Mind Readings).* This is such a neat experiment because it yields results which make absolutely no sense without PCT, since it clearly shows how you can get the same "responses" when the "stimuli" are very different. So I was trying to figure out how this could be done with rubber bands. Here's as far as I got.

Use the classic setup of two rubber bands looped together and thus joined by a knot. The disturber inserts a piece of chalk in his or her end of the rubber band, and the controller does likewise. They start out so that the knot is over the reference spot; the controller is asked to keep the knot there. Then the disturber slowly draws a pattern or letter or writes a short word while the controller compensates (controls the knot). Then they move to a different spot on the board and do it again (same disturbance pattern used).

The purpose of this is to show that while the disturbance and response patterns are (essentially) the same in the two runs, the movement of the knot (the "stimulus") is not the same. This is the magic of control.

While it is easy to get a record of the behavior of the controller and the disturber, I don't see an easy way to get a record of the *knot's* movement. Maybe if I use *two* sets of rubber bands and use long pieces of chalk (attached to the rubber bands at the bottom and top), I can join the rubber bands on another piece of chalk and have it leave a record of the knot's movement. I'll have to try this out. Meanwhile, I would appreciate any other suggestions for this portable demo.

Chuck Tucker: Without modesty, I would refer each of you to my rubber band demo paper in *Continuing the Conversation* (12), Spring 1988. I merely transformed Bill Powers' original demo in *Behavior: The Control of Perception* into several sets of instructions so that they could be used with a group (e.g., a class or conference). Instead of using a chalkboard with the rubber bands (or a single rubber band knotted in the middle), I use 3′ x 5′ drawing paper (which I attach to the board with masking tape) to make a record of the trace of pencils or markers (each person has a different color), then I take down each paper after the demonstration. After all of the demonstrations, I can compare the tracings (you could also reduce the tracings to distribute to the group and also use the tracings in papers you might write about the demonstrations). One way to make a record of the movement of the knot is to videotape the entire demonstration, making certain that the center of the frame is the knot. I would also suggest making the knot a different color from the background and drawing a grid around the target. With slow motion, you could see (and even crudely measure) the movement of the knot. You could also use the video to show to other groups as well as improve upon your procedures.

Bill Powers: Our image of "rubber banding" is unfortunate in one respect, because this demonstration has deliberately been made very simple, to illustrate principles. A more realistic example of rubber banding would give the control system one rubber band attached to the knot, and 20 different people 20 rubber bands attached to the same knot. The control system won't have any difficulty in controlling the knot (unless the combined disturbance results in breaking of the control system's rubber band) because only the vector sum of disturbances matters. Control might actually be easier because independent random disturbances will sum to a net disturbance having much less variability than any one of them has.

But any one person acting as a disturbance, trying to influence the control system's hand position, is going to have great difficulties because of all of the other random disturbances that are present. While control still remains possible, it's no longer possible for the disturber to estimate the best direction to move his/her own hand to achieve a correction of the other's hand position, because there is no longer any best direction. And it becomes difficult for the putative disturber to know what disturbance is actually being applied; perceiving one's own rubber band's tension is no longer indicative of the net disturbance on the other's controlled variable. The only way to make sure of applying a known disturbance is to isolate the control system from all of those other influences.

Chuck Tucker: Instructions for Students for a Rubber Band Experiment:

1. Review the "rubber band experiment" described by Runkel in his book (Chapter 10). You need a rubber band knotted in the middle (called RB below), a target diagram (three examples are given and others suggested), and a table or other flat surface.

2. Select a person, P, with whom to carry out a modification of the RB experiment. Instead of having P guess what is reasonable for your behavior, you will ask P to adopt a particular reference signal and perform accordingly.

3. Place a target diagram (see examples below) on the table and ask P to keep the knot in the middle of the RB "over" the center of the target by saying: "Please put your finger in this loop of the rubber band and keep the knot in the center of the rubber band above the 'X' in the middle of the target. I will put my finger in the other end of the rubber band and move my end, but you should keep the knot over the 'X' until I say *stop."* (Pause until both of you have your fingers in the ends of the RB and have placed the RB on the target diagram with their knots over the "X.") Say *"stop"* when you have accomplished one of the purposes described below.

4. Your assignment is to move your loop of the RB such that you can place P's finger (which is inside the other RB loop) over the letters on Diagram A to spell out the word "CONTROL" by having placed his/her finger on these letters in sequence: C, 0, N, T, R, 0, L. When you have done that, say *"stop."* Then ask P: "Do you recall what word you spelled when your finger touched the letters?" Whatever the answer, tell P: "The word you spelled is 'CONTROL.'" Then say to P: "I want you to keep the knot over the 'X' again; I will tell you to stop for each of the letters." Do as above and have P spell "CONTROL," stopping when each letter is touched by his/her finger. *Hint:* If P maintains his/her reference signal (maintaining the knot over the "X"), you should be able to place P's finger over the letters. P cannot control *both* his/her finger *and* the knot: they are connected by P's maintenance of the RB and P's resistance to your disturbance. If P wants to control the knot over the "X," P must resist your disturbance. With ingenuity, you can therefore get P to place his/her loop finger over each of the letters in turn.

5. You can do this assignment with Diagram B (colors) by having P's finger touch a sequence of colors. You can make diagrams of your own with figures (triangles, circles, squares, rectangles) at both ends or with numerals like 1, 2, 3, 4, 5, 6, 7, 8, 9 at both ends, etc., as long as you can get a sequence of moves, remember them, and have P touch them.

6. Repeat this exercise with a total of three persons.

7. Your report should be typed (no more than three pages, double-spaced) and include the following: (a) a description of your procedures, P's verbal and non-verbal actions, and the outcomes for all three persons; (b) based on the observations reported in (a), write a one-paragraph analysis of the hypothesis that you can influence another person only if it serves some purpose which is important to that person.

8. Your could add an appendix to your report describing what happens when you ask P to see if he/she can get you to spell a word, touch colors, figures, symbols, or numerals.

Diagram A

A	R	C	U	L	O	T
N	L	O	C	B	P	F
S	I	Y	O	N	C	V
Q	W	R	T	O	N	D
A	G	E	I	U	Y	S
H	K	Z	L	B	Z	P

X

A	R	C	U	L	O	T
N	L	O	C	B	P	F
S	I	Y	O	N	C	V
Q	W	R	T	O	N	D
A	G	E	I	U	Y	S
H	K	Z	L	B	Z	P

Diagram B

BLACK	WHITE	RED	BLUE	YELLOW
GREEN	ORANGE	BLACK	WHITE	PURPLE
YELLOW	RED	GREEN	BLACK	WHITE
BLACK	WHITE	RED	BLUE	YELLOW
GREEN	ORANGE	BLACK	WHITE	PURPLE
YELLOW	RED	GREEN	BLACK	WHITE

X

BLACK	WHITE	RED	BLUE	YELLOW
GREEN	ORANGE	BLACK	WHITE	PURPLE
YELLOW	RED	GREEN	BLACK	WHITE
BLACK	WHITE	RED	BLUE	YELLOW
GREEN	ORANGE	BLACK	WHITE	PURPLE
YELLOW	RED	GREEN	BLACK	WHITE

Diagram C

@	%	&	*	$	!	?
?	!	$	#	&	%	@
*	&	@	*	?	$	%
@	%	&	*	$	!	?
?	!	$	#	&	%	@
*	&	@	*	?	$	%

X

@	%	&	*	$	!	?
?	!	$	#	&	%	@
*	&	@	*	?	$	%
@	%	&	*	$	!	?
?	!	$	#	&	%	@
*	&	@	*	?	$	%

Instructions for Students for Observing and Recording the Coin Test:

1. Have a person, P, arrange four "coins" on a table (don't use actual money, instead use either quarter-size circular paper disks or same-color poker chips), making a specified pattern which P writes down in advance but does not show you. You might do this by asking a classmate or friend to help you with a class project which will take about twenty minutes and can be done without going to any other place. Besides the "coins," you need a writing instrument, paper, and a table top, flat board, or desk.

Say something like this: "I will give you four disks and a piece of paper. What I want you to do is to think of a pattern in which you can arrange these (name of objects) on this flat surface. But first, I want you to draw and name that pattern on this paper; *do not show it to me.* Next, arrange the (name of objects) in the pattern that you have on your sheet of paper. Then, without having any more conversation between us (except for the announcement 'No change in pattern' when my moving of the disks does not change the pattern) until I say 'Game is over,' I will change the pattern and you are to put the coins back in the pattern after each time I change it. Do you understand this game?" *If not*, repeat the instructions as written above again and clarify where necessary.

2. Your task is to discover what P has in mind without asking any questions or using any verbal communication at all. Your discover pro-

cedure is to change or disturb the pattern among the objects by moving the objects around in some ways and noting what you have done to change the coins. So before you make a change, draw and name the pattern of the objects on a piece of paper without showing it to P; after you make the change, draw and name the pattern that you made with the changes. If your change alters the pattern P has specified, P must correct the error by re-arranging the objects to re-make the original pattern. If your disturbance does not alter the pattern that P has specified, P must announce "No change in pattern."

3. Repeat this process until you are certain that you can: (a) specify and demonstrate three disturbances that will call for P to re-arrange the objects and correct the pattern; this you should be able to discern from what you have drawn on your pad—you will have to change the objects many times to observe three disturbances and three "No change in pattern" instances; (b) specify and demonstrate three disturbances that resulted in P announcing "No change in pattern"; this you should be able to discern from what you have drawn on your pad.

4. Compare your drawings with P's drawing of the pattern and (a) report the extent of agreement, including whether you identified the pattern but named it something other than what P named the pattern; (b) report at least one example of failing to see the pattern to which P returned the objects following a change that you introduced.

5. Your report should be typed (no more than three pages, double-spaced) and include (a) your report of agreement between you and P; (b) your answer to 4(b); (c) a brief statement on what you learned from this experiment; (d) a copy of P's written specification of the pattern; (e) your drawings of three change and three no-change patterns.

In the "coin" test, I tried to imagine what pattern would be "coins" arranged as follows: "All coins exactly where they are and oriented as they are." The only one I can come up with is one where P is controlling for the relationship of the "coins" to the surface; as long as the "coin" is on the surface, then there is "no change." I have tried to eliminate that by having P arrange the "coins" in a pattern and name it, which makes it extremely difficult for P to control for such a reference signal as "relationship to the surface." I have changed the game so that coins are not used, to eliminate "heads and tails" situations, different types of coins, different colors of coins, and other such variables. I would also suggest that one demonstrate this game with those whom you are asking to do it before they do it. In a class, I use either a magnetic board with letters "o" or a felt board with disks. There are ways, in other words, to reduce the variability of patterns for this game. Of course, if you want to introduce a variety of variables that are possible for controlling, then you can make the game more complex.

I am also posting three revised diagrams for the rubber band experi-

ment. These diagrams have the symbols/colors in a mirrored reverse arrangement. This arrangement makes it possible for the E to place his/her finger on a spot on his/her side of the diagram, and the P will have his/her finger on that same spot in mirror image, be it symbol, letter, color, or whatever. The E does not have to watch the finger of the P, but simply his/her own finger. This is a way to show that E can control for a particular symbol, and, if P is controlling for the knot over the "X," then P's finger will be on the same symbol as E's.

Diagram A

A	R	C	U	L	O	T
N	L	O	C	B	P	F
S	I	Y	O	N	C	V
Q	W	R	T	O	N	D
A	G	E	I	U	Y	S
H	K	Z	L	B	Z	P

X

P	Z	S	L	Z	K	H
S	Y	U	I	E	G	A
D	N	O	T	R	W	Q
V	C	N	O	Y	I	S
F	P	S	C	O	L	N
T	O	L	U	C	R	A

Diagram B

BLACK	WHITE	RED	BLUE	YELLOW
GREEN	ORANGE	BLACK	WHITE	PURPLE
YELLOW	RED	GREEN	BLACK	WHITE
BLACK	WHITE	RED	BLUE	YELLOW
GREEN	ORANGE	BLACK	WHITE	PURPLE
YELLOW	RED	GREEN	BLACK	WHITE

X

WHITE	BLACK	GREEN	RED	YELLOW
PURPLE	WHITE	BLACK	ORANGE	GREEN
YELLOW	BLUE	RED	WHITE	BLACK
WHITE	BLACK	GREEN	RED	YELLOW
PURPLE	WHITE	BLACK	ORANGE	GREEN
YELLOW	BLUE	RED	WHITE	BLACK

Diagram C

@	%	&	*	$	!	?
?	!	$	#	&	%	@
*	&	@	*	?	$	%
@	%	&	*	$	!	?
?	!	$	#	&	%	@
*	&	@	*	?	$	%

X

%	$	?	*	@	&	*
@	%	&	#	$	!	?
?	!	$	*	&	%	@
%	$	?	*	@	&	*
@	%	&	#	$	!	?
?	!	$	*	&	%	@

Bill Powers: Here's an example of the bandwidth of a control system. Hold up your forefinger about 18 inches in front of your nose and move it slowly from side to side over a total distance of three or four inches, like a slow metronome. Now, keeping the average position and the amplitude of movement the same, gradually speed up the movement, like a metronome going faster and faster. Keep going faster until you absolutely can't do it any faster. At that point you will be using your whole arm, and you will feel quite large muscular efforts, even though the movement from side to side is still only three or four inches (try to keep it that way).

The fastest movement you can produce is at a frequency essentially equal to the bandwidth of your finger position control system. Obviously, you can perform this back-and-forth pattern at any slower speed (lower frequency) with no great difficulty, right down to zero frequency (stationary finger). But when you try to produce an oscillating movement at a frequency higher than the bandwidth, your control system simply won't obey. You can *imagine* a faster movement, but you can't *produce* a faster movement.

Why is there a bandwidth? One explanation might be that your muscles simply can't reverse the motion of your arm any faster, because they reach the limits of force that they can produce. If that were the only limit, you ought to be able to move your finger faster if you move it over a span of only a quarter of an inch instead of three to four inches. The maximum force needed to maintain an oscillation goes as the square of the frequency, so when you move your finger one-tenth as much, you should be able to oscillate your finger about three times as fast.

In fact, you can move perhaps a *little* faster, but certainly not three times as fast. You can oscillate your finger with an amplitude of, say, four inches or less at about four to five cycles per second, but not significantly faster, *even* for the smallest movements (I assume you're not a concert pianist, and anyway, concert pianists don't have much occasion to practice sideways trills).

If you increase the amplitude to a foot or eighteen inches, you will indeed find a decreasing speed limit set by muscle strength; the force required increases linearly with amplitude in a linear system (which your arm is not). At large amplitudes of movement, you slow down because your muscles won't produce enough force to maintain the same frequency of oscillation you can maintain with a small amplitude. But below a certain amplitude, the speed limit is no longer set by muscle force. Something else is limiting the speed.

When you slowly speed up a small movement, keeping its amplitude the same, you'll notice another phenomenon. At low frequencies, you see a finger waving slowly back and forth. But at the highest frequency you can produce, you can see the finger only at the end of each movement, where it reverses. Between those positions it's just a blur; you can see right through it. Obviously, you couldn't track anything with your finger at that speed, because you couldn't see its movements, much less track irregular movements of something else. What you're seeing is the bandwidth of your visual perceptions of position. The frequency at which your finger just ceases to be a blur and becomes a finger again is the bandwidth of retinal position detection (actually, you have to suppress eye movement by fixating on the background to find the true bandwidth, Which is quite low, only two to three Hz).

It's interesting that the bandwidth or maximum frequency for small movements is higher than the bandwidth for retinal position detection. Something is limiting kinesthetic control at a frequency higher than that at which position control takes place, but at a lower frequency than is set by muscle strength. This probably involves a perceptual limit, too, in that kinesthetic position sensors do have speed limits, but more likely it is caused by temporal filtering that is required in order to make the kinesthetic control systems (that position your finger in the dark) *stable.*

The kinesthetic position control systems contain time delays of something like 50 milliseconds of neural transit time and synaptic delay around the loop. The muscles themselves have viscous damping. The noisy nature of neural signals, trains of impulses, requires that some smoothing take place in order to turn barrages of neural impulses into smooth changes in neurochemical concentration levels. All of these factors mean that there is an unavoidable lag in these systems of about 100 milliseconds, part of it a transit-time delay, and part of it an integrative or smoothing lag. That would imply that to switch as fast as possible from one position to another under kinesthetic control should take a little longer than 100 milliseconds, and to switch back another 100 milliseconds, for a total of 200 milliseconds for one cycle of a repetitive movement. That would give a frequency for continuous switching of four to five Hz, which is pretty close to what you see when you do it. Not bad for a ball-park estimate.

You can easily see the relationship between speed of movement and bandwidth. Try the experiment again, with small movements, only this time switch as fast as possible from one position to another four inches away, pause, then switch back as fast as possible, and pause. You're trying to generate a square wave. At low frequencies, each switch is discrete. You finger blurs over to the other position and is stationary for a while, *then* blurs back again. But as you increase the frequency of the square wave, still making each movement as fast as you can, the movements begin to blend into a continuous movement, so that when you reach the maximum frequency you're back to a continuous

sine-wave movement. In fact, even at the low frequencies, each switch has been like half a cosine wave—a high-frequency cosine wave at just about the bandwidth frequency. So the slow square wave you started with was rounded off a little, and that rounding-off means that the movements actually never exceeded the maximum bandwidth for continuous oscillations.

It is possible for you to generate oscillations at higher frequencies. The only way to do it, however, is to destabilize your spinal control systems, the lowest level of control. If you press your hands together very hard and maintain the push until the muscles begin to fatigue, you might see "clonus" oscillations, at a frequency of about eight to 10 Hz. This results from changing the force-tension curve in the muscles enough to make the control systems unstable. They break into spontaneous oscillation. But you can't produce this kind of frequency voluntarily. (You might see lower-frequency oscillations—the next level might get unstable first. Shivering is probably a clonus oscillation of this kind, produced by destabilizing the control systems in some other way. So climb naked into the refrigerator if you want to see 10-Hz oscillations).

For visual tracking using control of finger position to follow a target, you obviously have to be able to see a finger while it's moving. This means that the bandwidth for following a randomly moving target is about two to three Hz, the frequency at which the finger just stops being a blur. This bandwidth is set by perception and output functions, not muscles. The kinesthetic systems clearly have a wider bandwidth; they can execute faster movements than you can control visually. And the lowest level of kinesthetic control, the spinal reflexes, have the widest bandwidth of all.

What's most interesting to me is that these nested bandwidths are just about what is necessary to maintain stable control at each level. There would be no point in being able to see movements beyond a bandwidth of two to three Hz because the kinesthetic control systems used by a visual-motor control system have a bandwidth only slightly higher—four to five Hz. Therefore, we *don't* see faster movements! In fact, if we could see faster movements, the bandwidth of the visual control systems would be so high that the lags of the lower control systems would be too long for stable control at the higher level. In technical terms, at a frequency where the phase shift of a sine-wave disturbance passing around the loop is 180 degrees, the gain would still be above one. Negative feedback would turn into positive feedback at that frequency, and the whole system would oscillate. Oscillation is not good for control.

Rick Marken has explored several of the higher levels of perception, showing that as the (hypothetical) level increases, the bandwidth of perception continues to decrease. This is only logical, once you do some experiments yourself. For example, while moving your finger back and forth as fast as you can, *vary the amplitude* between, say, a four-inch amplitude and a two-inch amplitude. Obviously, you can't even *see* "amplitude" in a time smaller than the fastest oscillation. And to vary amplitude, you have to have a couple of oscillations of each size. In principle, you could do one large oscillation and one small one, and so forth. In practice, you can't perceive changes in amplitude that fast. So you can't control amplitude as fast as you can control position. Rick's demonstrations are simple and elegant, as usual, showing the effect clearly. So naturally he can't get them published.

The relationships between bandwidths at different levels are, once you understand why they exist, perfectly simple and logical. It seems that bandwidth follows from physical principles and obvious relationships among physical phenomena, such as between frequency and amplitude. It's obvious that you can't change amplitude in less than one complete cycle, because amplitude doesn't even exist until at least one cycle is completed. Ho hum.

But remember that this is a constructed reality we're talking about. This relationship holds because of the way we perceive amplitude as a function of movements. Having constructed a perception of amplitude, we then discover that it has properties, and that it is related to lower levels of perception such as movement and position. The ho-hum self-evident relationship suddenly becomes evidence about how perception is constructed—much more so than evidence about the natural universe. The bandwidth relationships also tell us that higher perceptions must be functions of lower ones, and that higher control systems use lower ones to accomplish their control. The evidence just continues to pile up that we are looking at—and *with—a* hierarchy of perceptual control systems.

When is the world going to wake up to what is going on here?

Gary Cziko: I did it! I just designed the most awesome manual PCT demonstration of all time. And you just need three pencils, four rubber bands, and two pieces of paper taped together (on the other side) to make a long sheet about 22 inches high and 8-1 /2 inches wide (or just use two attached fan-fold computer sheets).

Take the three pencils and attach them to each other like rungs on a ladder using the rubber bands. Now get your long piece of paper and draw a line horizontally across the middle (just above or below the seam of the two sheets). This is the target line. Place the paper on a table and tape down the corners so that it won't slide about.

Take one end-pencil and have your subject take the other end-pencil. Put your pencil point above the target line at the extreme left side

of the paper and have your subject put his or her pencil point below the target line so that (a) all pencils are perpendicular to the paper, (b) the middle pencil point is on the target line, and (c) the rubber band connecting the subject's pencil to the middle ("cursor") pencil is perpendicular (at a 90-degree angle) to the target line. Tell your subject to maintain these two perceptual variables (cursor pencil point on target line and rubber band at 90-degree angle) as you *slowly* trace out an approximation to a sine curve above the cursor from one side of the paper to the other. Make sure that all three pencil points are making contact with the paper and leaving a trace (felt-tip pens leave nice traces with little pressure).

After you've done this once, *do it again,* this time making sure that you, as experimenter, follow the same line as you did the first time.

You will now have before your very eyes a very remarkable piece of paper. Above the target line, you will see an approximate sine wave drawn twice (they will look more like one line if you're a really good disturber). These are records of the two *disturbances.* Below the target line, you will see two mirror images of the approximate sine curve drawn twice. These are records of what the subject *did.* They will probably be more irregular than the disturbances, but there should be an obvious similarity between the two response curves. In the middle, you will have two "cursor" lines, which are records of what the subject saw during the two trials. These two lines should not have any discernible pattern to them. In addition, they will *not be similar to each other* (if they are, this is an indication that you disturbed too fast and the subject lost good control).

This is very strange indeed, since the subject's responses are *similar* on the two trials and yet what he/she saw (the cursor pencil point) during the two trials was very *different.* How can the subject respond similarly on two trials when what was seen (the "stimulus") was so different? If anyone can come up with an explanation of this which does not look like a closed-loop negative-feedback model, please let us here on the net know about it.

Your subject might find it difficult at first to control both the position of the cursor pencil on the target line and the angle of the rubber band, so you might want to let him or her practice first using the eraser ends of the pencils. Alternatively, you can practice yourself and let your participant be the experimenter.

This is a manual (i.e., non-computer) approximation of the task and analysis used by Rick Marken in "The Cause of Control Movements in a Tracking Task," available in his book, *Mind Readings.* Rick showed in a similar task using a computer and game paddle that the correlations between cursor variations (here, middle pencil variations) were usually less than .20, while correlations between response variations were always greater than .99.

Now comes the fun part. If you are a psychology student, show this demonstration to your local non-PCT psychology professor (if these are hard to find, please let us know where you are located) and ask him or her to explain the findings. He or she will most certainly have to say that the two sets of cursor variations are similar, even though they are not. If he or she doesn't believe they are not similar, show him or her Marken's paper with the fancy computer and game paddles and correlations sometimes to *four* decimal points; he reports one correlation between cursor variations of .0032 with a corresponding correlation between response variations of .997. If your non-PCT psychology professor is really sharp, he or she will quickly point out that a correlation of .0032 can be statistically significant with a large enough sample). If he or she is not that sharp (or much sharper), he or she should be quite shaken up.

Greg Williams: Gary says: "This is very strange indeed, since the subject's responses are *similar* on the two trials and yet what he/she saw (the cursor pencil point) during the two trials was very *different.* How can the subject respond similarly on two trials when what was seen (the 'stimulus') was so different?" The two "responses" to the two "very *different" "stimuli" are* not the *same,* only *similar.* Skinnerians would have no problem with the *fact* that (even slightly) *different* "responses" resulted from different "stimuli." And if you show them results where successive "responses" are *identical,* yet the "stimuli" in each case are different, they will talk about "stimulus generalization" or say that the organism can "lump" different-appearing stimuli (to the experimenter) into *one* kind of "discriminative stimulus." But it gets even worse. If successive "responses" are judged as different by the experimenter, they might say that they really are all in the same "operant" set of responses.

Gary also says: "If anyone can come up with an explanation of this which does not look like a closed-loop negative-feedback model, please let us here on the net know about it." At the level of the observed phenomena, it is obvious that the "stimulus" in your experiment is affected by the "response." Skinnerians have no problems with such situations, which they call instances of "self-stimulation." But at the generative-model level (should any of them dare to speak thereof, lest they lose their "Skinnerian" labels!), some of them might argue that the "discriminative stimulus" is "middle pencil point moving (either direction) away from the middle line" and that the (ongoing) "operant" (set of "responses") consists of "actions to move the middle pencil back toward the middle line." Such a generative model makes no explicit reference (no pun intended) to postulated internal (to the organism)

states. (PCTers, of course, will immediately note the *implicit* reference level. Skinner used to argue that bringing in such hypotheticals would add nothing to the "analysis of behavior," and, worse, would tend to distract one from the data. I still think he had a point, *to* a *degree.* Yet, by hewing to that line so cautiously, he was unable to explain the existence of particular "wants"—to which I quickly add that he did not *want* to explain that). Also, the typical Skinnerian would want to claim that the person would "respond" to the "discriminative stimuli" because of *previous* reinforcements having to do with his/her relationship to the experimenter. Regardless, Skinner's notion of "operant" *sets* of *outputs,* each of which result in the same *outcome,* was a significant step toward replacing specification of outputs with control of perceptions.

Bill Powers: Gary, a very nice implementation of Rick's experiment. I'm sure that demonstrating it is much simpler than explaining it in words.

Thinking about your description of what happens, and about Rick's experiment, I first thought that the failure of the "cursor" movement to predict the behavior must be due to the fact that control is good enough to bring the error down to the noise level of the system. This is the general explanation I've been entertaining, for some time, for the phenomenon of low correlation between behavior and the variable that it controls.

There is, however, another possibility: chaos. Most of the control systems we've investigated are modeled best by a system that integrates the error signal to produce output. Integrals are known for hypersensitivity to initial conditions, one brand of chaos. When error is near zero, any slight perturbation will lead to an output that drifts away from the optimal setting one way or the other, which way depending on the sense of the perturbation. The result is that the perturbations due to noise are greatly amplified; the system disturbs itself and these disturbances result in a continuous wandering of the controlled variable in the vicinity of its reference level. So the wanderings are actually much larger than one would predict from the basic signal-to-noise level of the neural signals. They are large enough, in fact, to be comparable to the amount of error required to produce the output that opposes disturbances. I suspect that this is a better explanation of the low correlation between action and the controlled variable.

Rick Marken: Gary says: "I did it! I just designed the most awesome manual PCT demonstration of all time." Nice work! Looks like the end of psychology as we know it, right? Wrong! Greg points out: "Skinnerians would have no problem with the fact that (even slightly) different 'responses' resulted from different 'stimuli.' And if you show them results where successive 'responses' are *identical*, yet the 'stimuli' in each case are different, they will talk about 'stimulus generalization' or say that the organism can 'lump' different-appearing stimuli (to the experimenter) into *one* kind of 'discriminative stimulus.' But it gets even worse. If successive 'responses' are judged as different by the experimenter, they might say that they really are all in the same 'operant' set of responses."

Greg makes an excellent point: psychologists in general (and Skinnerians in particular) are not going to be persuaded by these demonstrations of principles, because it is very easy to say it's just "stimulus generalization" or "discriminative stimuli" or "operants" or whatever. You can't "persuade" people with these demos unless they are (1) willing to be persuaded, and (2) willing to deal with the problem *quantitatively*.

Gary, your demo shows that sensory input is not the cause of behavioral outputs, no matter how ridiculously counterintuitive this seems. But will this demo convince a psychologist who is busily doing research based on the assumption that o = f(i)? *No way.* S/he can always describe the results *verbally,* invoking the shibboleths of scientific psychology: "stimulus generalization," "response generalization," etc., and get back to work.

As Greg said, demos like this are no problem for the scientific psychology establishment. I confidently predict that if you (Gary, or anyone else) try this demo with a standard psychologist, they won't even break stride; they'll have a quick explanation, see no problem, and go off, comfortable in the knowledge that there is no problem at all. I don't think any demo—no matter how clever—will ever wake the psychological establishment from its dogmatic slumbers. Only those who are willing to learn—and are willing to think quantitatively—have any hope.

I say this, Gary, so that you won't be too frustrated when you find that your brilliant demo produces virtually *no* revelations among your colleagues.

Very nice work, though.

If people don't want to believe in control, then they don't have to. I think people throughout the behavioral sciences have a serious commitment to the input-output view of behavior (and it is a "view"—just as much as feedback control is a view). No amount of evidence can "demolish" someone's world view; people who wanted to view the earth as a fixed sphere at the center of the universe had no trouble dealing with evidence suggesting that it was not; moreover, some of the most compelling evidence was in favor of the stationary earth view—just as some of the most compelling evidence is in favor of the input-output view.

Ed Ford: Gary, Chuck, and all of you other PCT demonstrators, I have a different way of demonstrating PCT with rubber bands. Take two rubber bands (I prefer big ones) and knot them together. Ask the participant to hold the ends of the rubber bands, one in each hand, facing you. With his/her hands outstretched, the knot will be directly in front of both of you. Then point your finger at the knot and ask him/her to keep the knot directly in front of the tip of your finger. Begin moving your finger and he/she will automatically look at the relationship of the knot to the tip of your finger. Next, ask him/her to look at his/her left hand and watch its actions while trying to achieve the same goal of keeping the knot at the tip of your finger. Obviously, he/she can't. In fact, there is a strong internal urge to take a look at the knot-fingertip relationship. Thus, he/she will perceive the need for feedback and the inability to achieve goals by watching behaviors. I've found this to be the best way to lay to rest the idea that we control our actions.

An alternative is to get two people to participate. Begin with the two rubber bands knotted together. She holds the end of one rubber band, and he holds the end of the other rubber band, with the knot between them. Again, you point your finger, with the tip being right at the knot. Move your finger around; they have to keep the knot right at the tip of your finger. Now, ask them to achieve the same goal by watching their actions and to concentrate on how they move their hands as they attempt to reach their goal. Or tell them to watch each other's actions. Or ask one of them to close their eyes and the task for the other becomes more difficult. Move your finger about. Again, the internal desires on their parts to look at the knot show the need for feedback to achieve internal goals and the inability to control by concentrating on the actions. I find that when you ask people to switch from watching the knot and its relationship to the tip of your finger to watching their hands move, it becomes so obvious how we control for input, not output. The nice thing about these demonstrations is that you don't need a chalkboard, you can demonstrate most anywhere (with the exception of a phone booth), and you still maintain control over the disturbance.

Bill Powers: Ed, those new rubber band demos will take their places in the fundamental set. I think that Occam's Law, saying that we should choose the simplest and most parsimonious explanation covering the facts, ought to be supplemented by Occam's Economic Law, saying that the way of communicating the explanation should cost as little as possible. You have taken a demonstration that could be done on a $2000 computer and have managed to show every major point using equipment costing about 10 cents. Very, very nice.

Ed Ford: I'd like to add an additional feature to my version of the rubber band demo. You can show others what it is like to try to control another's actions. I had my wife, Hester, hold the two knotted rubber bands, and my son, Joseph, point his finger while Hester tried to follow the finger (as he moved it about) with the knot. I then gently said, "I'll help," and I took her arm to help her achieve her goal. Anyone want to guess the response? "Hey, I can do this myself" is what she said. Her attention immediately went from her chosen goal of trying to follow the pointed finger of Joseph with the knot to getting me off her back. I suggest when you try this that you use someone with whom you have already established a close and warm relationship. I don't want to be sued. Bill, you can't get that out of a $2000 computer.

Bill Powers: Ed, what's nice about the rubber band experiment is that you can demonstrate just about all of the features of PCT with it. Your latest is a fine example.

Gary Cziko: Bill Powers has made some illuminating arguments [not quoted in this thread] about how a feedback system can be expected to be faster than an open-loop, ballistic system. I've found what I think is a simple way to demonstrate this, to which I would like to get some reactions.

You need a large, smooth table, two coins, and a stopwatch. An assistant to work the stopwatch might also be useful. Stand or sit beside the table. Next, take two coins (B for ballistic and F for feedback) and put them on one side of the table, with one coin directly above the other. Now slide one coin (B) from where it is to the other side of the table ballistically; this means giving it a quick push with your finger so that it slides shuffleboard-style across the table and comes to rest where the laws of physics say it must. The time taken from the initial push to the final stop should be timed. Now move the remaining coin (F) from where it is (and from where B started) to the new location of coin B and time how long it takes. Do this as quickly as you can while keeping your finger on the top of the coin as you slide it across the table. Compare the two times.

When I do this, I get times for coin F which are *always* faster than coin B, feedback taking only about 60% of the time of ballistics (for example, 0.86 seconds for coin B, compared to 0.52 seconds for coin F). It seems that the two coins take off at about the same speed, but that coin B starts to slow down immediately after the push, while coin F continues to accelerate until getting very close to the target, where a very sophisticated braking system takes over to decelerate the coin sharply right before the target.

Somehow I feel that the physics of coins sliding on tables doesn't make this a watertight demo showing that closed loop can be faster

than open loop (limbs don't have much friction, do they?), but I'd like people to try the demo and get some reactions from the more physically enlightened people on the net.

Rick Marken: Gary, the coin-sliding demo is absolutely your best yet!! I love it. I hope it holds up to the scrutiny of the physically inclined types. But I think there are all kinds of possible variations on it that can satisfy even the most hard-nosed members of the net. It's really ingenious, Gary; nice going.

Robert Clark The various rubber band demos are fabulous! I particularly like Ed's recent one, taking one band in each hand. Especially the reaction to "help." I am reminded of the leader-follower" demo with the Portable Demonstrator. Do you remember this, Bill? It requires two subjects who first work with ordinary finger-tracking separately with the experimenter. Then one subject, Joe, is asked to lead the tracking, and Pete is asked to follow. After they settle down, the experimenter calls, "Pete," who is now to be the leader, with Joe the follower. Again, after settling down, the experimenter calls, "Joe," and Joe and Pete change roles. Make the changes rather slowly at first, so that it becomes easy—then gradually increase the pace.

It isn't easy to adjust a remembered skill to an unfamiliar situation. I am reminded of the time it occurred to me to carefully apply my ball-throwing sequence, right-handed, to my left arm. I thought I might be able to do this, since I am generally rather ambidextrous. It took careful and detailed reworking of the remembered movements to apply them to the other side. It worked much better than I expected! But I found that I had to have *no* witnesses (distractions? violations of my self-image?), and it felt so strange that I never tried it again! But there are many such reworkings possible. In school, we played with "talking backward." I have found since that this is not uncommon, but we did not reverse the spelling, we reversed the sequence of phonemes! "Nack ooyah cawt zdrawacab?" Not easy! With experimenting to acquire a stock of remembered performances, "Ti zih tahn draha!"

Bill Powers: Bob Clark says: "I am reminded of the 'leader-follower' demo with the Portable Demonstrator. Do you remember this, Bill?" Yes, indeed, and thanks for bringing it up. It's been a long time since I mentioned it, however, and it really does belong in the portable demo collection. Just to expand a little on your brief description: The object of the demonstration is to see how long it takes people to switch roles, namely, from leader to follower. B moves a finger arbitrarily in space while A tries to keep a forefinger aligned with B's finger. This results in B tracing out some pattern in space, while A's finger lags behind it a little, always trying to catch up. Then, on a signal from a third party, the two participants swap roles. Now A is moving a finger in arbitrary patterns, while B tries to track it with a finger as closely as possible. Clearly, it is now A who creates an arbitrary pattern in space, while B's finger lags a little behind it, always trying to catch up.

The third party keeps giving the signal at variable intervals, and the participants keep swapping roles, until they are executing the swap as fast as possible. The claim that Bob and I would make is that the minimum possible time required for this swap is longer than the time taken to change any lower-level control process. The time should be longer, for example, than the time required to correct the error when tracking a regular pattern over and over, with the disturbance being a sudden stop in the target pattern. And that time is longer than it takes to track a target that moves in random jumps to fixed positions, which is longer than the time it takes to respond to a downward push by swinging the arm rapidly downward, which is longer than it takes for a directly disturbed arm to begin to move back toward the undisturbed position. So we would seem to have five nested and demonstrable levels of control with progressively longer reaction times, the fifth being the role-swapping, and the lowest being the position reflex. A proviso is that all of these tasks should be well-learned, so we aren't looking at reorganization along with the control actions.

I just checked this out with my wife, Mary, and it still works. While checking it out, it occurred to us to wonder what would happen if one of the people simply changed roles without warning the other and without any external signals. With different pairs of people, the results might be different, but in our case the results were hilarious. I won't spoil it for you by describing it.

Gary, your coin-sliding demo inspired me to think up another demo that shows a little more of the effect you want. We happen to keep around the house various toy trains, for purposes of grandparenthood. I picked a wooden train car about six inches long, weighing about a pound, with wooden wheels and a convenient hook at each end. To each hook I fastened a string of three rubber bands, fairly weak. I set the car on a table, on its wheels. Then, using both hands, I stretched out the rubber bands so the car came to a balance point between my hands.

There are now two ways to move the car: (1) Move both ends of the rubber bands by a fixed distance to one side and let the car end up where it will; (2) Watch the car and move your hands (keeping tension between them) so as to bring the car to a fixed position.

If you try to move the car as fast as possible by the first method, you can make two marks on the table and move your hands to the marks as rapidly as possible. The car will be accelerated in the direction your hands move, reaching maximum velocity just as the tensions in the

two rubber bands are equal. It will then proceed past the midpoint until its velocity is reduced to zero by the growing tension in the trailing rubber band. It will then accelerate back the other way, and so on in diminishing oscillations to an end-point.

Using the second method, you mark the final position of the car resulting from the first method, then reset the car to its original position. Now you watch the distance between the car and the mark, and move your hands in parallel, maintaining tension between them, to bring the car to the mark. It will move to the mark and stop there with no oscillations. With practice, you can make it do this far more rapidly than you can get the car to the mark the other way.

This would be even more dramatic if the rubber bands were very weak and the wheel bearings good. You would have time to accelerate the car toward the final position by moving your hands far to one side, to get a strong acceleration, and then far in the other direction to slow the car to a stop, your hands returning to the correct final position automatically.

The only way to make the first method work almost as well as the second method would be to generate an arm-movement waveform just right to produce a high initial acceleration, and then at just the right time, a high final deceleration. In other words, provide a central pattern generator of high precision that produces the same arm positions as the control system generates without pre-programming.

This is why people have been driven to proposing motor programs instead of systems that just issue a "position" command. The motor programs are supposed to compensate for the dynamics of the controlled variable, as well as the kinematics of the jointed limbs. Once you start down this trail, still thinking of commanding output, you are driven step by logical step until you fall into the hole. Your basic premise leads you to propose a pattern generator of incredible precision, and a program of equal precision that bases its command outputs on unobtainable data of just as great precision—and it requires you to ignore all long-term disturbances.

The second method is not only simpler and faster, but it can work indefinitely (no cumulative computation errors) and it can achieve good final precision using low-precision output effectors, even in the presence of environmental nonlinearities and disturbances.

Robert Clark: Bill, your more complete description of the leader/follower demo is helpful. I would join in your claim that "the minimum possible time required for this swap is longer than the time taken to change any lower-level control process." The "lower-level control process" consists of tracking the leader's finger. This requires control of muscle variables, position variables, and time variables.

The follower has formed recordings from observations and, perhaps, tracking experiences. He can select one that might produce acceptable results. He uses this to provide reference levels to produce his tracking movements. As the pattern changes, different recordings are needed. It takes more time for these changes than it takes for lower-level (muscles, positions) changes. To the follower, this is still the tracking demo.

The leader also has a supply of recordings available from his experiences, etc. His assignment as Leader calls for him to select one to be tracked by the follower.

Warning! If your subjects are unfamiliar with participating in such demonstrations, there can be some unexpected side effects. For example, some people avoid the role of leader. Being a follower might be acceptable, but being the leader introduces some intrinsically different perceptions. The particular behavior depends, of course, on the specific individual. As switching becomes faster, the participants might become confused and conflicts (internal) might develop.

I suggest that you examine your own—remembered—internal experiences when you have been a participant in this demo.

Ed Ford: More on the rubber band demo: I spent the last week in Michigan, training 32 teachers, counselors, and administrators in control theory (among other things). I showed the rubber band demo, where a teacher held two knotted rubber bands stretched out, with the knot directly in front of her chest. She had to keep the knot right at the tip of my moving finger. When I asked the participant to watch the action of her right hand instead of the knot, I began watching her eyes. I could see her eyes occasionally sneaking a look at the knot. Thus, she was able, by sneaking an occasional look at the knot in relation to the dot, to achieve her goal, but with less efficiency. I wanted to force her to just watch her actions, so I got someone else to take my place by moving his pointed finger in front of the knot. Then I took a piece of cardboard (about 12 x 8 inches) and placed it between her eyes and in front of the knot. Now she couldn't see the knot and the tip of my finger, and her ability to sneak glances was eliminated. Her inability to keep the knot over the dot became far more pronounced; in fact, she couldn't do it at all. That demonstrated clearly that we need feedback to achieve a goal, and that watching our behavior has nothing to do with controlling a variable.

Rick Marken: Here is a portable demo which could easily be turned into a computer demo. Just have the subject track your finger with her finger (I just did this with my daughter) as it makes a regular pattern (an approximately 8-inch-diameter circle seems to work nicely). Move your finger at the rate of about one cycle per second—slow enough for

good control, but fast enough so that knowing the circular movement pattern really helps. Then stop your finger at an unpredictable time. Your subject's finger not only takes a while to stop (about 1 /2 second), but while it is moving, it is tracing out an obvious *curve,* even though there is no target present to track. So the movement after the signal to stop is still controlled relative to a reference circular movement. There is "anticipation" that the target finger will not only continue to move, but that it will continue to move in a circle. (I put "anticipation" in quotes, because this could be modeled without any explicit computation of predicted target position at all—the model just controlling a higher order variable that could be called "relative circular motion.")

Now do the same thing, but use irregular movements of your "target" finger. Try to move you finger at about the same rate at which you were moving it to make the circle. I did it by writing out some words in the air. Now, when you stop the finger, you will find that the subject moves very little *after* the stop. This is because (in theory) the tracking is now being done at a lower level; if target movements are sufficiently unpredictable, there is nothing the subject can control except the distance between target and finger (a configuration). So there is no change in the variable to be controlled when the target finger stops; the distance between target and finger is all that must still be controlled. But when the target was a circle, the stopped target changes the variable controlled from "circular pattern" (probably an event-level perception) to no pattern.

Anyway, it's a nice way to spend a few minutes with your kids. My daughter got a kick out of seeing her finger keep moving in a curve after mine stopped; even though she was trying very hard *not* to let that happen. I didn't mind humiliating her in this way, because she keeps beating me at every computer game I've got.

Bill Powers: I sent the critique below to Dag Forssell after seeing the video tape of his "Purposeful Leadership" presentation to a group of Edward Deming aficionados. [Excerpts from the critique are reprinted here.—*Ed.]*

I would start the actual introduction to control theory by laying out your strategy to the audience and getting their agreement with it. What I would say would go something like this:

"In the rest of this presentation, we're going to go through two stages of development. In the first part, I'm going to teach you the basic principles of perceptual control theory. To do this, it's best to focus on a simple example and make sure you understand every aspect of it, so this phenomenon becomes familiar to you, and so you begin to know what to expect. Please don't worry about what this has to do with Deming's Total Quality Management (or whatever). I promise that we'll get to that. What I say about the Deming approach will make a lot more sense to you later if you just focus for now on grasping certain relationships that are basic to perceptual control theory. I hope you will interrupt, ask questions, ask me to repeat anything, no matter how simple, that you're not perfectly clear about. The better you understand what you see in this segment, the easier it will be for you to see the parallels when we start talking about real life. So for about the next half hour, let's all concentrate on a single goal, together, which is to master some basic principles and make sure that everyone understands them. The payoff will come in the half hour that follows. I want to hear that you're willing to do this: to forget Deming for half an hour, and work only on understanding the basics of PCT. How about it?"

Then I would go directly to the rubber band experiment. When you do it, you should have clearly in mind a sequence of basic principles that you want to demonstrate and explain. Don't worry about what you're going to say (output). Just be very clear at every stage exactly what you want the audience to understand. You are very good at this; you don't need to worry about your words.

The first part can go pretty much as you did in the video. Set up the task with a volunteer, and spend 20 seconds moving the rubber band around while the subject keeps the ping-pong ball on the target. Then ask the audience to explain it, as you did in the video. Be sure to emphasize that the question is, "What was the relationship between me, as the experimenter, and the volunteer, as the subject? What you would say is causing the subject to behave that way?"

With that finished, after no more than a minute, start explaining, and do so in much more detail than on the video. Say: "Watch what my hand does, what his/her hand does, and what the ball in the middle does. Notice that as I pull gradually back on my end, the subject pulls gradually back on the other end. Notice that the ball stays pretty much in one place. Now, as I raise (lower) my hand, notice that the subject lowers (raises) his/her hand. And notice that the result is *always* that the ball remains in the same place."

Then explain, right there, that you asked the subject to keep the ball in a specific place. You did *not* ask the subject to move the hand in any particular way. You told the subject what to *perceive,* not how to *act.* And right at that point, *prove* that the subject is not reacting to your hand movements. Bring along a sheet of cardboard with a notch in one side. Hold the cardboard at right angles against the paper with the notch directly over the ball so the subject can see the ball, but not your end of the rubber band. And demonstrate that the subject can still keep the ball in one spot without being able to see what your hand is doing. When everyone in the audience agrees that the subject doesn't need to see the disturbance, hold the cardboard so the subject can see your

hand but not the ball, and *prove* that control gets much worse when the subject can't see the ball but can see your hand. Check with the subject: "Can you see my hand? Can you see the ball?" Then check with the audience to be sure they get it: that the subject really has to see the ball in order to control it well. You're trying to establish some clear basic facts. This should have taken no more than 10 minutes.

Now you can start drawing your diagram. The ball is there in the environment, so draw a ball. The subject has to perceive the ball, so draw a perceptual function and explain that it creates the perception inside the person of the ball outside the person. Now ask, "Where is the target?" The audience will, of course, point to the target circle on the paper. But you say, "Wait a minute before you decide," and you whisper to the subject to keep the ball six inches to one side of the target, then spend 10 seconds showing the result. Then you ask the audience, "What do you think I told the volunteer?" Most of them will guess right; if they don't, tell them what you said.

Now ask again: "What does the subject want and where is that want?" What you want them to say, somehow, is that it's inside the subject. The subject is perceiving the distance between the ball and the target and obviously wants to see a particular distance, not necessarily zero, as you have just shown. When you look at the *piece* of paper, you see the actual distance, but you don't see the wanted distance. So where is it?

Now you go back to the diagram, and you show arrows entering the perceptual function from both the ball and the target. You label the *perceptual* signal "perceived distance." And now you *can* add the reference signal, labeling it "wanted distance." Emphasize that this wanted distance is now inside the person's head. Then the question is, "So what?" You have here the perceived distance as it is at any moment. You also have here a specification for the wanted distance. Something has to happen right here if there's to be any basis for action. What operation has to be performed? The answer you want to extort is "comparison." Somehow the person has to bring these two things together, the want and the perception, and judge how they are different. If the actual distance is greater than the wanted distance, the action has to make the distance smaller; if less, the action has to make it greater. So the action has to be based on the difference between the want and the perception. It doesn't depend just on the perception; it doesn't depend just on the want. It depends on the difference between them. Draw the comparator box, label the output "difference," and draw the arrow from the comparator to the output function.

Now ask what the rule has to be for converting the difference into an action, just in one dimension. This is not hard to figure out; if the perceived distance is less than the wanted distance, move your hand one way; if it's greater, move your hand the other way. Make sure everyone understands. Everyone should be nodding. If they aren't, ask what the problem is and fix it.

The last step is to close the loop. Notice that when the subject's hand moves, it moves in the right direction, and that the result is to return the ball to the target. You can illustrate this with the stimulus-response demonstration, suddenly pulling back on your end, suddenly relaxing again.

Now apply a very slow change in the disturbance and show that the ball remains near the target. "Notice that the perception and the action happen at the same time. You can't separate out the disturbance, the change in perception, the comparison, and the change in action. They're all happening at once. You understand how each part of this control system works; now when you see them all operating at the same time, you can see that the result is continuous control, in either one (pull back) or two (move up or down) dimensions. Or more."

Now show the relationship between the disturbance and the action. Explain why it now makes sense that keeping the ball over the target requires the subject's end of the rubber bands to move oppositely to yours. The subject is just correcting movements of the ball. This is the illusion of cause and effect. It seems that your hand movements are causing the subject's hand movements. But if you realize that the subject wants the ball to be in a specific place, in between the cause and the effect, of course you understand why the movements are as they are. When you understand what the subject perceives and wants, and what the subject has to do to make the perception match the want, you understand all of the relationships between apparent causes and apparent effects.

At this point, you can test their understanding. Either really, or as a thought experiment, ask them what will happen if your wife, Christine, knots another rubber band near the ball and pulls upward on it. How will the subject's hand behave? And why? If they have any problem, ask what will happen to the ball when Christine pulls, but the subject doesn't respond. Then ask what the subject has to do to get the ball over the target again (pull downward). It would be very nice to get the audience to make the prediction, and then actually do it and show that they are right. "Were you just guessing what would happen?" No. They *knew* what would happen. How did they know? Because they knew where the volunteer wanted—intended—the ball to be.

Now you can say: 'When we started this demonstration, I asked you to explain what the subject did. A lot of suggestions were heard. Now, if I ask you again to explain what the subject did, what will you say? What caused the subject to behave that way?" And you should get nothing but right answers.

At this point, you might take five seconds and ask: "Have you ever heard of a theory of behavior that lets you explain what anyone is doing, however simple, and *know* that you have the right explanation?"

The final point can now be made. Give the subject the pen, as in your video, and get a trace of the subject's actions as you create some random disturbances. Now you have to make this point very clearly, hammer it home, be sure that every person gets it. You say: "If you had just walked into this room, and were told that this is an accurate trace of the subject's exact actions, what could you say about what the subject was doing?"

This is the conclusion of the demonstration. Make sure that all of the people understand why observing actions doesn't tell you either what the person wants or what the person is perceiving. It doesn't tell you what the person is doing—what those actions are accomplishing that the person wants to perceive as being accomplished.

Close by telling them what comes next. "You now understand the basic concept called perceptual control theory. There is a lot more to learn, but what you know now will always remain true. People act in order to make their perceptions of the world match what they want those perceptions to be. You can't understand their actions unless you know what they are controlling, and the specific target. When you do manage to figure out what they're controlling, you can explain an enormous number of cause-effect phenomena—you can see what the cause is disturbing, and how the apparent effect, the actions of the person, are counteracting that disturbance.

"After the break, we're going to start applying what you know to the Deming Philosophy. We will look for parallels between what you saw in this very simple demonstration and what you see people doing in business management situations. We'll do a little role-play first, then apply perceptual control theory to the situation, and then do another role-play later to show how a person who knows PCT will act differently. We'll talk about Profound Knowledge and what all of this has to do with Deming's insights. We can't possibly cover all applications of PCT in the time left, but perhaps you are beginning to suspect even now that these applications will penetrate into every corner of life, in business and outside it. I have been developing my understanding of PCT for several years and still have much to learn. I envy you, because the initial experience of seeing a real theory of behavior unfolding for the first time is one that can't be repeated."

Then the break; give them time to talk about it with each other and let it soak in for a short time.

Robert Clark: Dag, you have changed my view of the rubber band demos. I had realized, of course, that they are useful and supplement the Portable Demonstrator that we used some 30 years ago. But when you introduced the "double ball" version, it added a new dimension. With the single ball, the lower orders of control can be demonstrated by suitable adjustment of the timing and the patterns used by the demonstrator. However, without losing any of them, the second ball makes the subject a full participant. He or she is asked to "select which ball to control"! This can be carried further by suggesting that he or she change, from time to time, which ball he or she is controlling. Another step: switch who is the demonstrator and who is the subject, done on the command of a third party.

Good show Dag, I appreciate the opportunity to know about your activities.

Gary Cziko: Dag and Bill, I was interested in hearing of the new uses of the trusty old rubber band demo. Here's another twist (or rather, slip). In this demo, the subject is told beforehand to keep the knot over some inconspicuous (to the audience) target, and the audience is trying to guess what the subject is "doing" as the demonstrator disturbs by pulling on his or her end of the rubber band.

Instead of using two rubber bands tied together, use three tied together end-to-end. Have the subject hold the end loop of one as usual, but you (the experimenter/demonstrator) hold on to the second knot, *not* the end of the third rubber band. (Got that?) Don't loop any fingers through, but hold on to the second knot between your thumb and index finger (as you would normally hold on to a string), with the third rubber band concealed within your hand.

Now do the demo as usual. When someone from the audience invariably says that the subject is simply mirroring your movements, stop at a position where there is good tension on the bands and then gradually let the third rubber band slip through your fingers and then hold again as the end arrives between your index finger and thumb. While the rubber band is slipping through your fingers, the audience will see the subject move his or her hand toward yours *while your hand remains still.*

So much for the "experimenter's *hand* as stimulus" explanation of the subject's behavior.

Rick Marken: Using the rubber bands, you can show that the position of the knot (p(t)) is always a result of what the subject is doing to his or her end of the rubber (o(t)) and what you are doing to the "disturbing" end of the rubber band (d(t)). You might be able to show (when you move the disturbance slowly) that the position of the knot does not change as you might expect it to if just the disturbance were acting. For example, when you pull gently to the left, the knot might be expected to move correspondingly to the left. But the subject might be able to notice that sometimes the knot is actually moving to the right (due to the added effects of their own actions) while you are pulling to the left. This means that the position of the knot is not a stimulus that "tells" the subject how to pull on their rubber band to correct the disturbance. So the stimulus-response view of control cannot be preserved even when the actual variable (the knot) that the subject is controlling is discovered.

But the deeper point is that perception is just there—it is neither right nor wrong, good nor bad, in error nor not in error—it is not *informative;* it just *is.* The position of the knot is just the position of a knot—but once you have a reference regarding where it should be, then it seems as if some knot positions are definitely "wrong," and one particular one is "right." This is a tough point to demonstrate, because people don't care much about the position of knots, and when you tell them this it seems pretty trivial. But try to explain that this applies *to all* perceptions that are controlled, and you will get some strong reactions. People who are controlling for the neatness of their house have a difficult time believing that the neatness of the house is just a perception—when the house looks "messy," that perception seems just plain wrong. It is difficult to demonstrate that perceptions are just perceptions, and that they only become "good" or "bad" or "right" or "wrong" (i.e., they only become informative) with respect to one's own references for them. That fact is easy to demonstrate with "knot" positions, but a hell of a lot more difficult to demonstrate with political, religious and economic "positions."

Bill Powers: Dag, I will work with you to make the demo section of your 'Purposeful Leadership" presentation an effective teaching tool. How effective it is will depend on how well the audience understands it, and on how well they can relate the principles embodied in the demo to other situations.

The point of the demos is twofold. First, you're just demonstrating a phenomenon of control, which is interesting in its own right, as you have found. Second, you're establishing a way of talking about the elements of and relationships in a control process, so you can use this way of talking later and remind people of what they learned through reminding them of their experiences with the rubber bands. The more clearly you establish what you're talking about in the beginning, the more easily the audience will understand what comes next. I'm going to lay out a strategy for doing this in a period of about an hour. The following segment might seem long and detailed to you. You might worry that the audience will wonder what this is all about, but don't worry. They will be interested because they are learning something.

First, you must carefully show the audience the physical elements of the rubber bands, so they will know what is important to notice:

a. The experimenter's end of the rubber bands.
b. The participant's end of the rubber bands.
c. The ball in the center of the rubber bands.
d. The effect of the experimenter's action on the ball.
e. The effect of the participant's action on the ball.
f. The combined effect of both actions on the ball.

You can do this part alone. You can stand facing the audience with one end of the rubber bands in each hand. Hold one end of the rubber bands still and move the other end, being sure you point out that you can both stretch it and move it up and down. Show that when you move your hand by a certain amount, the ball moves in exactly the same way, but by almost exactly half the amount. Show that this is true when you move either end, holding the other end still. Then return to holding one end still while you move the other end.

Now talk briefly about variables. You can say that you're affecting the ball with your actions. But what is it about the ball that you're affecting? It always has the same color; it's always round; its price is still whatever it is. What you're altering about the ball is its position, either up and down or side to side (illustrating as you speak with the appropriate move). It's only the position that is varying. The position can vary in two ways: up-down or side-to-side.

Here's an example of how the spiel might go: "So we can say that there are two independent *variables* involved. They are independent because you can change the up-down position without affecting the side-to-side position, and vice versa [illustrating as you speak]. So we are really talking about two variables here. If we know both variables, the up-down position and the side-to-side position, we know where the ball is in each of the two ways it can move, and that's all we care about right now.

"Now look at the hand holding the movable end. We say that this hand is carrying out an action. In this experiment, however, we're only interested in the action as it can affect the two variables that define the ball. We're interested in the *position* of the hand. This is a variable, too, and, in *fact, it's* two variables. The hand can move up-down or side-to-side [illustrating as you speak]. So we speak of the action that affects the ball in the same way we speak about the ball: in terms of variables. At any moment the hand variables are set in a certain way. As a result, when the ball variables are in a certain condition, the ball is in a certain position.

"All of this elaborate analysis is meant to let us see something that's not usually understood very clearly: the difference between an influence and an influence. When you understand what that means, you'll

already understand something important about human relationships.

"Look at the moving hand. Obviously, when the hand moves, the ball moves. So would you say that the hand position is an influence on the ball's position? Isn't this like saying that the driver's steering efforts are an influence on the way the car moves, or the teacher's personality is an influence on the students? This is one of the ways we use the word 'influence.' We point at the cause of something else and say that the cause is an influence on the something else. The moving hand is an influence on the position of the ball.

"But now look at the ball. When the ball moves, you would say that it's being influenced by something. You can focus on the effect of moving the hand and call that effect the influence of the hand. What do we now mean by the influence? We mean the behavior of the ball that is caused by the hand. What is the influence of the hand on the ball? Just look at the ball and you can see it: the ball moves. There is the influence of the hand.

"So now we have an influence in two different places: in the thing that's causing the ball to move, and in the movements of the ball. We can say that the teacher's strong personality is an influence, but we can look at how the student's behavior changes, and say, 'That change in behavior is the influence that the teacher had.'

"How do you influence people? Well, in the first place, you don't influence people, you influence variables—you influence something *about* the person that is variable, like the person's behavior or attitude toward you. You can't influence the person's height or age very much.

"Assuming we realize that we're always talking about variables, we influence people by acting in a certain way on them. But does this influence necessarily have any influence? When you apply an action that is supposed to be an influence, is the other person's behavior always influenced? Not by a long shot, and here's the reason. (Now you move both ends of the rubber band around so the ball remains stationary.]

"Look, I'm applying an influence to the ball with my right hand, but its position isn't being influenced any more. The position of my right hand changes, but the position of the ball doesn't. Suddenly my influence on the ball has lost its influence. This is very mysterious. What has happened?"

The audience, of course, can see you moving your other hand. Ask them to explain why your right hand has lost its influence on the ball. Tell them to go ahead and say why, even if it's perfectly obvious. Say it out loud, put it into words. But pin them down to an exact statement. Sure, it's because your other hand is moving the other way. But show them that if your right hand moves to the right, the left hand moves to the left; if the right hand moves up, the left hand moves down. Show them again what would happen if the left hand didn't move (the ball moves to the right), and then what happens when the left hand moves (the ball moves back to the left).

Then explain that each hand has a variable position, and each hand affects the variable position of the ball in each of the two possible ways. The only way for the ball not to move is for the variations in left-hand position to be exactly *equal and opposite* to the variations in the right-hand position. Only that will leave the ball in the same position, if the two rubber bands are identical. The *influences* of the two hands on the ball are equal and opposite, with the result that there is no influence on the ball.

"So the next time you try to get a vendor or an employee or a customer to behave in a certain way, you will think of this, won't you? What you say or do might be an influence on the behavior of the other person, but it might not have any influence. Why not? Because there might be another equal and opposite influence coming from somewhere.

"Now we're going to find out where the most important equal and opposite influence comes from. May I have a volunteer from the audience, please?"

Now you turn to the easel with the paper on it, draw a target circle, and take the volunteer aside and whisper the simple instruction. You can explain out loud that you and the volunteer are going to keep your hands lightly touching the paper. Assume the position.

When you apply disturbances, apply them very slowly and smoothly. Adjust your speed so the volunteer can keep the ball over the circle very accurately. Don't let transients occur; they're confusing at first.

'Now watch. I pull back, using the influence I have on the ball to make the ball move. I move my hand up, influencing the ball to move up. I move down, around in a circle, all different ways. And you can see the influence on the ball that my hand is having, right? [Turn to the audience and raise your eyebrows and ask, inviting an answer, "Right?" Get the audience to point out that you're not having much influence.] Wrong. So even though I'm varying my hand position up and down and side to side, the ball isn't varying that way. Why isn't my influence having any influence? [Audience, even if you have to drag it out of them: "Because Jim is moving his hand the other way."]

"Yes. I'm applying an influence to the ball, but the ball isn't moving because Jim is applying an equal and opposite influence to the same ball. It's just the same as when I had hold of both ends of the rubber bands, but now Jim is playing the part of my other hand.

"Why do you think Jim is doing that? [Audience: "Because you told him to."]

"Yes, but what exactly do you think I told him to do? What would you guess the exact instructions were?"

Now there is a period of discussion while people volunteer guesses. Some will guess right, some will guess wrong. Just let the guesses accumulate for a minute or two, without commenting.

"Ok, you've told me your guesses, and you've heard other people guessing. Is there anyone who wants to change the guess now? OK. Jim, what did I ask you to do? ['Please keep the ball as exactly over the circle as you can."] Thank you. Some of the people out there think you're a liar, but I know you're not.

"I didn't tell Jim how to move his hand. I asked him to produce a certain effect on the ball, and he evidently agreed to try. He evidently succeeded very well. But *how* did he succeed? What was he doing, inside, that created the result you saw? Now we're looking for something besides just a description of what we all could see happening. We're asking how Jim could be organized so he was able to do what you saw him doing. We're looking for an explanation of what we saw.

"You've all heard explanations of human behavior, according to one theory or another. You've probably found some explanations more convincing than others. I'd like to find out now what sort of explanation you think would apply to this little experiment. How do you think Jim works, which would explain what he was doing? For example, how many of you think that Jim could keep the ball over the circle with his eyes closed? [Get a show of hands.] Nobody thinks you could do it, Jim. Let's get into position, and you close your eyes and carefully follow this instruction; listen carefully: keep the ball exactly over the circle. [Jim closes his eyes, and you start moving your end of the rubber band around. This will provoke a bit of laughter.]

"Well, it's pretty obvious that Jim can't follow the instructions with his *eyes* closed. We have made a great discovery: when Jim closes his eyes, he becomes deaf.

"All right, if that's not it, what do we know now? Why did Jim have to see what was going on? [More comments from audience.]

"Let's try to get very specific. What exactly did Jim have to see in order to do what he did? [Get more guesses—your hand, the ball, the rubber bands, whatever.]

"Well, let's test a couple of these ideas. If Jim had to see my hand, then it wouldn't make any difference if *he* couldn't see the ball, right? So we can just dispense with the rubber bands and the ball, and Jim can move his hand the way he thinks he needs to move it when I move my hand. When I say 'freeze,' Jim, just rest your hand on the paper and hold it there, and I'll do the same. Here we go. [Perhaps it would be good for you both to have dry markers, to mark the position.]

"Freeze. Now, with my other hand, I connect the rubber bands the way they were, and let's see where the ball is. [This is *one* reason for making sure that Jim can control very easily and accurately.] Well, not too bad. Are you satisfied with that, Jim? If not, go ahead and put your hand where you think it should be. (Jim corrects remaining error.]

"Now, some other people said that Jim was looking at the ball. Suppose that's true: he can see the ball, but not my hand. I'll hold up this piece of cardboard with a notch in it, and Jim, you position yourself so you can see the ball but not my hand or arm. Ready? Here we go. [Experiment proceeds: use slow, large disturbances. The piece of cardboard should be large enough to conceal entirely your half of the playing field.]

"All right, we have the evidence now. What's your conclusion? [Get some conclusions.] Of course, we can use the last resort: Ask. Jim, while you're keeping the ball over the circle, are you looking at my hand or at the ball? ['The ball."]

"Jim has served us well, but it's time to see if he's the only person in the world who can do this task. Let's thank Jim, and ask for another volunteer. [New volunteer.]

"OK, just a, quick check: keep the ball exactly over the circle, Jane, while I hold up the cardboard—be sure you can't see my hand or arm. [A few seconds of demo.] Good, you work the same way Jim does. Would you like to try it with your eyes closed? No, I didn't think so.

"Can we agree now that watching the ball is sufficient? In other words, Jane doesn't *have* to see my hand, and it probably wouldn't make much difference if she could, because she could hardly keep the ball centered any better. Jane, why don't you sit down here for a little while, because I want to draw a diagram before we go on. [Draw the rubber bands and ball with the target circle a little off from the ball.]

"We've established that Jim and Jane look at the ball during this task. So they were looking at something in this region. [Draw a circle around target circle and ball.] Jane, did you also need to see where the target is? ["Yes."] Jim, you too? ["Yes."]

'Now, what does 'seeing' mean? We see with our eyes, of course, but what gets into our eyes has to get into the brain, too, before any perception happens. So let's draw a box up here, with an arrow representing light rays coming into the box, and an arrow coming out that represents what the brain knows by way of these light-waves. Right at the end of the arrow coming out of the box, I'll draw what the brain would be seeing right now, based on how the diagram looks. Here's the ball, and here, away from it a bit, is the circle.

"Jane or Jim, or both: if this is what you saw, what would you be trying to do? [Reply: Get the ball over the circle.] How would I draw a picture of that? [Reply: Draw the ball inside the circle.] Like this? [Above and to the right of the picture of the perception, draw two concentric circles.] So here we have a picture of how the ball and the circle actually look right now [indicate perception], and here we have

a picture of—what? Jane or Jim, or anyone? [Wait for: How they are supposed to look, etc.]

"Would it be accurate to say that this [reference picture] is how you wanted them to look? ["Yes."] Is this how they always looked? ["No."] Well, then, how did you know how they were supposed to look? Before you answer, Jane, will you come up here again and do a short run with me? [This time, move your end just rapidly enough so that the ball wobbles all around the circle]. Now, how did you want the ball and circle to look? [Jane tells you or points to picture.] Most of the time, how *did* it look? [Indicates perception somehow. If she doesn't point to the pictures, you do it.]

"OK, you knew it should look like this? [Point to reference picture.] And most of the time it actually looked more like this? [Point to perceptual picture.] Good. Well, if most of the time it looked like this [perception], how did you know about this? [Point to reference picture.]

"Let's switch to another example for a moment. Most of you drive cars. When you are going along a straight road, you steer the car to keep in its lane. What are you seeing out the windshield in front of you? [Get descriptions.] Now, consider: How do you know where the car is in its lane? [More.] And finally, how do you know where it *should be* in its lane? [Etc.]

"All of this is building up to a point that a lot of you might have seen by now. The remaining question is: *Where* is this knowledge of the way the car and road, or the ball and circle, should look? ["In your head."] In your head. Can all of you imagine a ball centered in the circle, right now? Can all of you imagine the way the car and road look when you're in the right position on the road? And where is that imaginary picture, right now? In your head—or at least, not anywhere in the room outside you. Even if you don't actually see an imaginary image, there's knowledge, somehow, of how the *scene* should look when it's right. Right?

"You're now ready to understand the theory of human behavior that's behind this presentation. Just a few more steps.

"First, let's start using some consistent terminology. This arrow in the brain, up here, that shows how the ball and circle actually look right now, we'll call the *perception*. Notice that we don't call the actual ball and circle down here, the real ones, the perception. The perception is what the brain, up here, knows about the world, down here.

"If the picture of the actual situation is the perception, then what can we call this other [reference] picture? It's not a perception of the actual ball and circle. It's an imagined perception. We judge the perception of the actual ball and *circle* with reference to this other picture, which just sits there unchanging, telling us how the actual perception should look, not how it does look. So let's call this other picture the 'reference perception.' Or we could say 'the reference condition of the perception,' or just 'the reference condition.' The key word is *reference,* because it's with reference to this [reference picture] that we judge this [perception).

"Now, I ask you: Is this [perception] the same as this [reference]? How do you know that? What would you call the process you carry out in order to decide that they're not the same? [Hope to get "comparison."]

"We call it comparison, and when we draw models, we draw a box right here, which receives information from the perception and from the reference, and compares them. We call it a *comparator.* And what comes out of the comparator? [Draw arrow.] Information about the difference between the perception and the reference. If there's no difference, no information comes out. If the perception is different, this arrow carries just the information about the difference. We can call this arrow a difference signal—in control theory it's called an *error signal,* and you can use that term, too, as long as you understand exactly what it means. It doesn't mean mistake or blunder, it just means that there's a difference. If there's any amount of error signal up here, we know that the real ball, down here, is not in the same position as the circle—or at least it isn't perceived that way.

"While we're at it, let's identify this other box down here. It's called an input function or a *perceptual function.* It receives light-rays or other physical information about the world and converts it into some sort of representation in the brain. It creates a perception, or as we sometimes say, a *perceptual signal,* that continuously indicates the state of the outside world. Right now, your brains contain some perceptual signals that indicate how my words are sounding and how I look as I stand up here. Obviously, everything in this region of the diagram is the brain [draw a big circle], and the rest is outside the brain.

"So, way down here, we have the actual circle and ball. Information conies from them into this perceptual function, creating this perceptual signal that always indicates the relationship of the circle and ball. Up here we have another signal, the reference perception or condition that's showing how the perception *should* be. And here is the comparator receiving both of those signals, comparing them, and spitting out a signal that represents how much difference there is—how far from the reference condition the perception is, and in what direction. These so-called signals are simply currents flowing through nerve fibers in the brain. But we don't have to worry about neurology here; this is about organization.

"Now, if the perception looks like this, and the reference looks like this, what should Jim or Jane do? Obviously, move the arm so that the ball goes this way, toward the target. It would work equally well if the arm could make the target move the other way, toward the ball. And

where is the information that tells which way to start moving? Right here, in the error signal coming out of the comparator.

"All we have to do is hook up this difference or error signal to Jane's arm muscles in the right way, and the arm will automatically move the ball, and keep moving it until there's no more difference signal to tell the arm to move some more. Let's watch it happen.

"Jane, if you'll assume the position....

"Center the ball. Thank you. Now we'll do this a little differently, in stop-motion. First, close your eyes. [Move your end of the rubber bands to move the ball.] Now open your eyes and make what you see look right. Now close your eyes again. [Move in a different direction.] Open again. Close again. Open again. [Etc.] Thank you.

"By stopping the motion, we can see what's going on. Each time Jane opens her eyes, she sees a different picture of the ball and circle. The reference condition is the same, so the comparator puts out a different error signal each time. This results in a different motion of the hand each time, and it's always in the right direction to make the perception move toward the reference condition. [Point to the right places on the diagram as you talk.]

"When we stop the motion like that, we see what looks like a series of stimuli followed by responses. But when we do it in the natural way (Jane, one more time, please, with eyes open), you can see that there are no stimuli and responses. The difference or error is never allowed to get very big, unless I start moving this end of the rubber bands too fast. In fact, Jane is acting continuously to keep that error or difference signal from ever getting very large.

"By doing that, she is keeping the perception of the ball and circle very close to this reference picture. It takes only a very tiny error to make Jane's arm start moving to correct it; the effect of the movement is always just right to keep the error small.

"This is how you drive a car, isn't it? You don't wait for the wind or a tilt in the road to put you in the wrong lane, and then steer back. As soon as you can detect any difference between where you perceive the car to be and where you want it to be, you alter your steering efforts just enough to prevent that change from getting any bigger. So your car hardly wanders at all. At least that's how I hope you drive. These little corrections are quite automatic. You don't have to know about these signals in the brain or how they're hooked up. All you have to do is pick a reference condition. This little circuit here will then make sure that what you perceive matches what you intend to perceive. This little circuit is called a negative-feedback control system. This reference signal is where you put your intention in.

"One last look. Jane, I'd like you to go into slow motion. Do everything just the same, but slow down your actions as if you have to push your arm through heavy syrup. Let's try it. I pull back on my end, and you slowly bring the ball back to the circle. You don't have to wait for my motion to finish; you can start acting right away, but make your action very slow. [If this doesn't work, you can change roles.]

"Now, you can see how disturbing the ball creates a little error, which starts the arm moving the right way. After a while the error is gone again. While my arm is moving, there's a continuous error, which is keeping her arm moving the other way; when my arm stops, she catches up and the error disappears. Thanks, Jane, it's been great.

"That was like seeing a slow-motion film of a control system in action. There's always a little error, a little lag, but not very much. The action is always pretty much equal and opposite to the disturbance, and the error is always pretty close to zero.

'Think back now to where we started, almost an hour ago. Jim got up here and moved his end of the rubber bands around, and you saw what he was doing, but did you understand what you were looking at? Now we have a model to explain what's happening. You can see why Jane or Jim's arm seemed to be mirroring my motions, as if imitating them. You can see why Jim acted to prevent me from having any influence on the position of the ball. You can see that what mattered was not how my arm moved, but how the ball moved. The actions of Jim and Jane were controlling the ball, not just reacting to my arm movements. They didn't even need to see my arm or what it was doing to the rubber bands. All they needed was to see where the ball was, and know where they wanted it to be. That explains everything you saw.

'When engineers work with systems organized like the one in the diagram, they bring all sorts of complications into it. Things like differential equations, Laplace transforms and z transforms, Bode plots, sampling theory, and even information theory. But they're talking about the same system you see here, behaving just as you saw it behave, organized exactly as you see it organized in this diagram. A closed circle of cause and effect. Perception, comparison, and error driving an output—although, of course, they wouldn't talk about perceptions. You now understand the essence of this sort of system in just the way an engineer might understand it, and if you've followed the presentation, your understanding, you can be sure, is correct.

'The last thing we have to do is bring in a few more terms, and then we will be armed and ready to tackle the application of this concept to human behavior in the areas that interest you.

"At my end of the rubber bands, we have something we will refer to as the *disturbance.* We call the position of my end of the rubber bands the disturbance because it disturbs the ball, or *would* if there were no other influences acting on the ball.

"At the other end, we have the person's action. The term 'action' means just what the person's muscles are directly causing to happen, positioning the hand. We can talk about the action without talking about any other effects it might have. The action is also an influence on the ball, but as you have seen, the behavior of the ball isn't the same as the action itself.

"And in the middle, we have the *controlled variable.* In this case, the controlled variable is the position of the ball relative to the circle. We call it a variable because it is capable of varying. We call it controlled because the actions of the person control it. The actions bring the controlled variable to a specific condition, and they vary in whatever way is needed to keep that variable in the same condition. That's what we mean by control.

"So, in the environment of the person, we can see a disturbance, a controlled variable, and an action that is producing the control. In our model of what goes on inside the brain, we can see a perception that represents the controlled variable, a reference condition or signal that represents the intended state of this perception, and an error or difference signal that drives the action. Put all of these elements together, and they add up to an explanation of the behavior you have been seeing. Put them all together, and you have a revolution in the behavioral sciences, which we're soon going to begin applying. Any comments or questions? We can take 10 or 15 minutes for them if you wish. I could go on with this presentation for about three days, so don't worry that we'll fail to *meet a* schedule. We'll just get as far as we can. The most important thing is that you understand, not that we finish an agenda."

After the talk and milling around is done: "Now, let's talk about what behavior is. I need another volunteer just for a couple of minutes. You? Good, come on up. You will see that perceptual control theory, which is what we're talking about, gives a person a lot of confidence. It works with any randomly selected person.

"Here's a dry marker. Hold it against the paper while you move your end of the rubber band, so it leaves a trace. Keep the ball exactly over the *circle,* right. Now, just keep it there for a while. [Put in a slow but broad pattern of disturbances.]

"Thank you—that's all. Now, suppose that someone had just come into this room and heard me say, This trace was created by Pete's hand in the experiment you just saw.' What might that person conclude about Pete's behavior?

"You can't say that Pete's behavior didn't produce this wavering and wandering trace. It did. But is that what Pete was doing? Was he really just making this squiggle on the paper? There's a saying among the adherents of PCT (which is what we call perceptual control theory) that goes: 'You can't tell what a person is doing just by looking at what the person is doing.' Here's a beautiful example of that. What Pete did was to move the dry marker around and leave this trace. But what he was *really* doing was keeping the ball over the circle. You, who know about the controlled variable that Pete was concerned with, understand that. But the person who came in late didn't see the controlled variable. The only evidence left is the record of Pete's actions, which tells us exactly nothing about what Pete was controlling by means of those actions.

"So, you can't tell what a person is controlling just by looking at that person's actions. This is a profoundly revolutionary idea. In most ordinary aspects of life, we look at the people around us and we think we can see what they are doing. We look at their 'behavior,' in quotes. But what are we really seeing? We are seeing their actions. We are not seeing the variables that are being perceived by those people, and being controlled so that the perception is kept near some reference condition. Only the person we're looking at knows what perceptions exist, and what state of those perceptions that person would prefer to experience. Only that person can see the relevance of the action to maintaining control over a particular perception. We, observing from the outside, can't see the purpose of the actions.

"Imagine that we went through another session with this demonstration, but held a big piece of cardboard up so the audience couldn't see the ball and circle. You could see my hand on one side, and Pete's hand on the other side, and you could see them moving, but that's all. Wouldn't it seem that Pete's hand movements were being caused by mine? It would look as though Pete were watching my hand movements and responding with symmetrical hand movements of his own. If you had to draw a diagram of what was going on, you'd draw it like this: [Draw the rubber bands and ball. Draw a line from the disturbing end to a box and from the other side of the box to the action end.] The box is Pete. The movement of my end of the rubber band is sensed by Pete, and this stimulates him to move his end of the rubber bands. We have a nice simple cause-effect diagram, and Pete is just a link between the cause and the effect. If you grind that concept into your mind and really come to believe in it, what will happen when we take the piece of cardboard away? You'll see that the stimulus not only makes Pete's hand move, but tends to make the ball move because of the connecting rubber band. You'll see that Pete's hand movement also tends to make the ball move, but the other way. What an odd coincidence! The ball doesn't move at all, or hardly at all.

"Now, if keeping the ball directly over the circle were vital to Pete's health and safety, you might begin to wonder how the stimulus knows that it should cause Pete to move his hand in just the way that's in his own best interests. You'd try to find an explanation that seemed less outlandish, one that didn't make it seem that Nature was being altruistic. So you might propose that keeping the ball over the mark was

reinforcing to Pete. Whenever Pete didn't move the right way, the reinforcement wouldn't happen, so that wrong behavior would die out. Only the response to the stimulus that happened to keep the ball over the circle would be reinforcing, so that response would eventually be the only one left.

"You can see how it goes. Once you get a model firmly in mind and decide to believe it, all of your explanations from then on have to fit that model, even though they leave you with other mysteries. Just why should a ball being over a circle be reinforcing to Pete? You can't answer that question. All you know is that this explanation seems to work.

"We now have here a roomful of people who understand the control-theory explanation of what *we've* seen. You can compare the PCT explanation with the one we've just been through. While you're doing the comparison, consider this.

'The reinforcement explanation and the cause-effect model are the ones in which nearly every scientific psychologist has believed for most of this century. It's the one you learned in school. It's woven into our language and beliefs in ways that are so taken for granted that they're almost unconscious. Have you ever thought that by applying incentives to someone, you can get that person to behave differently? Have you ever explained your own behavior by pointing to something in your environment, and saying, 'That's why I did it'?

"Long ago, before anyone in this mom was born, the great minds of psychology and biology held up a big piece of cardboard. They said, 'ever mind what's behind this piece of cardboard. Just look at this end of the rubber bands and that end of the rubber bands. Isn't it obvious that movements over here are causing Pete to move his hand over there? You don't need to talk about purposes and intentions and desires and wants and wishes. All you need to do is observe what causes what. Then you will be able to predict and control human behavior.'

"Everyone in this room who has studied Total Quality Management knows what is wrong with that. People are not simply boxes with inputs and outputs, devices that can be made to act in certain ways by applying the appropriate stimuli. People have goals and desires and wishes and purposes and hopes and intentions. You ignore them only at great risk. The principles that Dr. Deming has given us are based on a deep awareness that people are not the kinds of devices that conventional science has told us they are.

'People are control systems. Deming realized this without having any formal understanding of why he knew they are as they are. He knew psychology was an important leg on which his approach stands—but he also knew that the psychology he needed was not the one that existed.

"PCT explains human behavior in a way that completely contradicts all conventional concepts, but which completely agrees with Deming's intuitive assessments. Perhaps even knowing only what you have learned in our simple little demonstrations, you can begin to get a feeling for how PCT is going to alter the psychological approach to management, and, for that matter, to getting along with people in general.

"Let's take a stretch and have some coffee for a while. When we come back, I'll give you just a brief look at some of the ways PCT could be applied, and is being applied. Don't expect to become experts in the final half hour. All I hope for is to stimulate your imaginations, so you will begin to see what lies ahead. You can probably guess that learning how to turn this new understanding into practical action takes more than an introductory session. But I'm sure that by the time we finish, you'll be able to go home and work out a lot of the implications for yourself, and start putting PCT to work."

Dag, that was more or less a role-play—what I'd say if I were doing the demo part of the presentation. Of course, I'd speak differently from the way I write. The things to pay attention to are the pace and the plan. One thing at a time, always aimed at the next thing, and all working toward the final conclusions. A lot of patience and details, with demonstrations of everything. A lot of interaction with the audience. Always demonstrating exactly what you mean, never just generalizing. What you want is for the audience, at the end, to understand what they have seen in every detail, and to make the connections between the specific things they've seen and the parts of one elementary diagram. You want certain terms to be familiar—it doesn't matter if the terms are technical, there's no need to search for the magic word that will make it easy for them. You show them what each word means, and they'll understand.

I advise you to study this presentation, so you see how points to be made later are prepared early on, and how one idea leads to the next logical idea. Notice carefully that the only generalizations are at the very end, after all of the specific hard-core ideas are laid in. And they are very sparing.

You're welcome to use any aspect of this material in any way that will help you. I hope you'll try it out, and try to develop the sense of single-minded development toward one rather simple and specific goal: getting the audience to understand the organization of one simple control behavior. Once they understand that, they will grasp everything else you have to say very easily.

The Control Systems Group is a membership organization which supports the understanding of cybernetic control systems in organisms and their environments: *living control systems.* Academicians, clinicians, and other professionals in several disciplines, including biology, psychology, social work, economics, education, engineering, and philosophy, are members of the Group. Annual meetings have been held since 1985. CSG publications include a newsletter and a series of books, as well as this journal. The CSG Business Office is located at 73 Ridge Pl., CR 510, Durango, CO 81301; the phone number is (303)247-7986.

The CSG logo shows the generic structure of cybernetic control systems. A Comparator (C) computes the difference between a reference signal (represented by the arrow coming from above) and the output signal from Sensory (S) computation. The resulting difference signal is the input to the Gain generator (G). Disturbances (represented by the black box) alter the Gain generator output on the way to Sensory computation, where the negative-feedback loop is closed.

Currently available:

LIVING CONTROL SYSTEMS

Selected Papers of William T. Powers

ISBN 0-9624154-0-5, 1989, 300 pages, $16.50

Introduction to Modern Psychology

The Control-Theory View

Edited by Richard J. Robertson and William T. Powers

ISBN 0-9624154-1-3, 1990, 238 pages, $25.00

LIVING CONTROL SYSTEMS II

Selected Papers of William T. Powers

ISBN 0-9624154-2-1, 1992, 277 pages, $22.00

MIND READINGS

Experimental Studies of Purpose

Richard S. Marken

ISBN 0-9624154-3-X, 1992, 212 pages, $18.00

Prices are postpaid. KY residents please add 6% sales tax. Order from: CSG Book Publishing, 460 Black Lick Rd., Gravel Switch, KY 40328.

Back cover

Closed Loop

Threads from CSGnet

This reproduction of *Closed Loop* was created by Dag Forssell in 2009. Addresses and phone numbers have not been updated. Most are obsolete.

Posted at www.pctresources.com

Spell checked, but not proofread. Please report any errors to webmaster. Email address at website.

Summer 1993

Volume 3

Number 3

Front cover

Closed Loop

Threads from CSGNet

Summer 1993 Volume 3 Number 3

Edited by Greg Williams, 460 Black Lick Rd., Gravel Switch, KY 40328

CONTENTS

Members of the Control Systems Group receive Closed Loop quarterly. For more information, contact Ed Ford, 10209 N. 56th St., Scottsdale, AZ 85253; phone (602)991-4860.

CSGnet, the electronic mail network for individuals interested in control theory as applied to living systems, is a lively forum for sharing ideas, asking questions, and learning more about the theory, its implications, and its problems. The "threads" in each Closed Loop, stitched together from some of the net's many conversations, exemplify the rich interchanges among netters. Some issues of Closed Loop also feature research reports by netters, in hopes of initiating new conversations.

There are no sign-up or connect-time charges for participation on CSGnet. The Internet address is "CSG-L@UIUCVMD" while CSG-L@UIUCVMD is the Bitnet address. Messages sent to CSGnet via these addresses are automatically forwarded to over 120 participants on five continents, as well as to hundreds of NetNews (Usenet) sites where CSGnet can be found as the newsgroup bit.listserv.csg-l. CSGnet also can be accessed via CompuServe, AT&T Mail, MCI Mail, or any other computer communication service with a gateway to Internet or Bitnet. For more information about subscribing to CSGnet, contact Gary Cziko, the network manager, at G-CZIKO@UIUC.EDU, phone him at (217)333-8527, or send a FAX to (217)244-7620.

Visions and Revisions

Bob Clark: Receiving *Closed Loop* led me to realize that I really should get back in touch with the current state of the work that Bill Powers and I developed together some 30 years ago.

I have been using, applying, and developing these early concepts on my own during this period. In beginning to post to CSGnet, I offer some comments on issues (3) and (4) of *Closed Loop.*

My main comment on (3) is about the lack of discussion of the hierarchical levels that could /would be involved! Instead, much of the discussion recapped the standard arguments for/against governmental control—especially of the economy. It seems to me that a more precise and accurate definition of the original hierarchical Orders would be very helpful.

Closed Loop (4) seems to me to go around and around because of insufficient recognition of the role of the engineer in designing and operating his/her system (whether it be "open loop" or "closed loop"). Indeed, many discussions omit the critical role of the observer, the experimenter, engineer, etc. I once co-authored a paper on a related subject: "A Systems View of Psychophysiological Experimentation," presented to the New York Academy of Sciences in 1964. Not a very good paper, but it points out some of the levels of interaction normally omitted from discussion.

The key in (4) seems to me to lie in the phrase—mentioned several times in the discussion—"point of view." The professional engineer takes his/her own role for granted: it is not part of the system he/she is designing. In his/her design work, he/she has, as noted in the discussion, full access to all aspects of his/her work and hence can use the terminology as he/she sees fit. But the moment the engineer is included, most of his/her hierarchical structure is in action!

Greg Williams: I, for one, am happy that Bob Clark has joined the net. His experience of several years of ruminations regarding living control systems—in parallel and independent of Bill Powers' ruminations—should enrich the dialog (two eyes are better than one, and all that). Go to it, Bob!

Gary Cziko: I have had several interesting phone conversations with Bob Clark over the past few months. I am very pleased that two of the three original developers of Perceptual Control Theory (PCT) are now on CSGnet (the third, Robert L. McFarland, has passed away).

Bill Powers: Welcome, Bob. You're right that the discussion of control-system engineering would have benefited from speaking about the levels, but the main problem as I saw it was trying to get a control-system engineer to see that his diagram and the PCT-type diagram are really the same. Obviously, my attempt didn't work. Our friend eventually signed off the net, wishing us luck but saying that as an engineer he simply *had* to think of controlled variables as outputs. Sound familiar to you?

Bob and I developed the basic control-system model in the years from 1953 to 1960. I have never properly acknowledged his part in this development, which was major.

I spent a couple of years trying to find a hierarchy of perceptions written as words, and layers of words, and dependency and inclusiveness relationships among words. Bob knows all about this: he went through it with me, putting as much energy and ingenuity into it as I did (Two Years Before the Blackboard—remember, Bob?). It was Bob who finally characterized what we were doing as "castle building"—building hypothetical dream-castles in the air, out of words. The insight that put an end to this futile project came out of the air between Bob and me: it was not the words we should be looking at, but the perceptions to which they point, which are not words. Control systems control perceptions, not the names of perceptions (unless one is specifically controlling for the construction of sentences and so on). And even a *word* is a nonverbal perception, in the final analysis. It is just a signal, distinguishable from other signals but having no inherent meaningfulness or special properties that other perceptions don't have.

This is what finally put us on the track of the hierarchy. The relationships and typings we were looking for were not to be found in the words we used, but only through looking at words as pointers and trying very, very hard to become conscious of the experiences, the nonverbal experiences, to which the words referred. Bob and I were both highly verbal people, used to communicating in words and pretty good at it, so this was a very difficult and drawn-out exercise. It required wrenching apart the words from their meanings, so that the meanings could be apprehended alone, without the words. I don't think either of us really understood at first what it was that we were trying to do. But when we came up with a "level," it was a nonverbal level for certain. Since then, I have seen that what we then took to be unitary levels merged types of perceptions we had not yet distinguished from each other. I think Bob considers my 11 levels too many; I suspect they are too few. But that's another subject. The central point is that separating words from meanings is no easy task and carries no assurance of being right in any final sense, but it is the essential *kind* of thing that must be done even to know what is meant, in PCT, by the term "perception."

Bob and I came up with general principles, but they were principles of control and hierarchical relationship, not generalizations about any particular perceptions (or words) at any particular level. The specific structure of the hierarchy that was proposed did not come from generalization, but from detailed examination of real examples as we discovered and explored them. There is no generating principle that will tell you, given knowledge of the nature of level n, what level n + 1 ought to be. The content of each level, insofar as we were able to characterize any level, was found empirically.

Bob Clark: My orientation is one of trying to clarify and simplify the basic structure of the hierarchical array of perceptual control systems. In the process, I noticed the discussion of thermostatic systems in *Closed Loop* (4), which seemed to end rather inconclusively. Further thought about thermostats led me to notice that they not only provide a convenient illustration of a control system (I've used them repeatedly for this purpose), but also can provide a tie-in illustration of a hierarchical structure of systems, couched in terms familiar to most people these days. As I developed these relationships, I found that I was using some familiar words that are usually taken for granted but could well have their basic definitions clarified—especially from the hierarchical-structure standpoint.

I selected thermostatic systems because they are familiar to many people, and because they include the basic elements of negative-feedback control systems. That is, they include: 1. a means for detecting a variable; 2. a means for affecting that same variable; 3. a means for subtracting the magnitude of that variable from some "preset" value, with a resulting positive difference acting to produce a positive output from the second "means." This is the usual combination of components composing a negative-feedback control system. Such a system need not have a continuous output to achieve its result. It is interesting to observe that "continuity" is, in part, a matter of "viewpoint." Thus, if a thermostatic system is observed over a period of several hours, its control approximates continuity.

The thermostatic system is a "one-way" system, as usually presented with a furnace, etc. That is only one of several limitations it suffers. Another "one-way" system is a living muscle fiber! It can only pull, not push.

The thermostatic system can also be used to illustrate other aspects of control systems—and other forms of control system. Before thermostatic systems were developed, people kept warm in the winter. My father had a coal-fired furnace that had a damper that adjusted its operation. Too cold, open the damper; too warm, close the damper. And it was a fairly continuous operation.

Where was the control system? Clearly, the situation was livable, although not as convenient as one might like. Obviously, there was a control system in operation, where the temperature (where?) was the controlled variable, even though the control was accomplished by adjusting the rate of heating.

But where was the control system? Without a person, the temperature was not controlled—but also, without a damper and a fire box, the temperature was not controlled. Some person *decided* whether action was needed, and in which direction. He/she then used his/her (lower-order) muscle systems to affect his/her *environment* according to his/her *understanding* of that environment.

What about "his/her environment"? This usually refers, perhaps vaguely, to the physical surroundings outside his/her skin. But someone else might have been available and have been asked to "open the damper" or "turn up the furnace."

Then where was the control system? A person could "do it himself or herself" or "ask someone else." Having made his/her *decision,* he/she used his/her lower-order systems to get his/her desired result. Were there two (or more?) levels of control involved, the "other person" and the "furnace"? Thermostatic systems now take the place of the "other person," and they are much more efficient and convenient.

For those who are familiar with thermostats, most of this is unnecessary. But what about those whose *environment* does not include *understanding* of control systems? You must have seen people turn the setting up higher and higher when the room doesn't warm up fast enough. The furnace was already at full speed, so raising the setting had no immediate effect. Later, however, the room was too warm, and the setting was reduced. This is "over-control"—the system is oscillating!

Notice the importance of the *time* scale of the person versus the "response time" of the system.

This illustrates the difference between regarding the assembly of parts as a control system rather than simply as a group of connected parts. Such a difference in *viewpoint* can result in a difference in behavior. Often these differences have little effect, but sometimes they are very important!

Both viewpoints are valid and result in the same mathematical representation. However, one is more useful for using the system, the other for modifying the system.

Several words have been italicized above: *decides, understanding, environment, choice, decision, viewpoint, time.* These words and their associated concepts are used routinely and seem to be readily accepted. However, each of them is very important and merits closer examination.

And how does each relate to a hierarchy of control systems?

Let us turn to the simplest description of a control system: the "logical box." This is a "black box" without the paint. The box can take many forms, but careful examination reveals that it has four unique connections between the world outside the box and whatever is within the box:

Number 1 is normally taken for granted. It is the cord that plugs into the wall outlet. When unplugged, the box is inert; it responds in no way to any kind of treatment—or mistreatment. The box requires some source of energy. It can be outside the box or inside the box. It is essential.

Number 2: This connection, when "disturbed," does nothing. However, it is observed that Number 4 changes its condition in a manner related to the disturbance of Number 2. The nature of the relevant disturbance at Number 2 (temperature, pressure, voltage, radiation, etc.) might be hard to determine. Likewise, the nature of the related changes in number 4 might be hard to identify.

Number 3: This connection similarly affects Number 4 and can respond to a disturbance different from the one affecting Number 2.

Number 4: As long as the box retains its energy source and *both* Number 2 and Number 3 remain constant, this connection (Number 4) is very difficult to disturb—perhaps impossible to disturb, short of destruction of the box.

As described, this box is not a control system. It does nothing until some specific disturbance is applied to either Number 2, or Number 3, or both. However, if some connection is provided in the *environment* of the box, such that the operation of Number 4 serves to reduce the disturbance of Number 2, then we have a control system acting to control (or at least tend to control) the disturbance of Number 2. In this case, we find that the same disturbance applied to Number 3 results in a corresponding change in Number 4, such that the disturbance of Number 2 is (very nearly) the same as that applied to Number 3. Of course, by suitable modification of the *environment,* the operations of Number 2 and Number 3 could be interchanged or otherwise modified.

This stylized and abstract description is consistent with the usual descriptions of negative-feedback control systems. It is presented here to emphasize that such a *system is defined exclusively* in terms of its *input* and *output* characteristics as it interacts with its *environment.* In addition, to act as a control system, it must respond to another *input,* serving as a *reference signal,* originating outside the system. Note that this says nothing about the nature of any of these three connections or the environment, nor about how they might be interconnected inside the box.

This box acts ("behaves") by controlling its perception, which consists of the disturbance to which it responds and controls.

The Test for the Controlled Variable applies directly to the box.

As presented here, some important items are italicized. These will be clarified later by discussion of their relations and implications in terms of the concept of a hierarchical structure of negative-feedback control systems.

Bill Powers: Bob, in your generalized four-port black box, rather than defining input 2 as affected only by a disturbance, I recommend saying that it is affected by the state of some external variable, so the output of port 4 depends on the state of the variable at port 2. Now we can distinguish the case in which there is an external link from port 4 to the variable at port 2 from the case in which there is a link from an independent variable to the variable at port 2. This also lets you define a "disturbance that doesn't disturb"—the independent variable might or might not affect the state of the variable at port 2, because the output from port 4 could change enough to cancel the effect the independent variable would have when acting alone. So you can distinguish the existence of a physical causal link from the kind of effect obtained by varying the cause.

Bob Clark: Bill, regarding the four-port black box, of course one can describe the situation in terms of its being "affected by the state of some external variable." However, if one is curious about some particular box, it must be intentionally "disturbed" if one is to discover the nature of one (or more) external variables it might be controlling. While discussion in terms of "states of variables" is certainly possible, I find a more experimental approach more useful. I also think some people are unfamiliar with the word "state" in this sense.

You suggest a difference between the case where there is an external link between terminals 4 and 2 and the case where an independent variable is applied to terminal 2. The experimenter, being outside the box, can observe any external connections. And any disturbance from the experimenter is (really, by definition) an "independent variable."

I am intrigued by the question of the boundary between the box and its environment. To the experimenter, it is quite clear. The environment includes everything except what is inside the box. But what does the box perceive? Its behavior is the control of its perception. But what does it perceive? Only the disturbances of those connections that cross the walls of the box. And its output actions reflect the differences between disturbances affecting terminals 2 and 3. In addition, some kind of connection must exist outside the box, not perceptible by the box, but observable by the experimenter. If these relations are not found, the box does not act as a negative-feedback control system.

These considerations seem to me to be applicable to many nonliving systems, and the use of nonliving systems as extensions of human systems is impressive.

Since I have been out of touch for over 25 years, much has been done that I've not seen. I have Bill's important book *Behavior: The Control of Perception, Introduction to Modern Psychology: The Control-Theory View,* and the collection of Bill's papers. Reading between the lines of some of the posts on the net, I get the impression that certain key concepts need examination.

I have been feeling the need for revisions and additions to the original hierarchical structure for some time. As I remember it, we expected the early publications (by Powers, Clark, McFarland, and others) to lead to further developments and modifications. Bill further elaborated and presented those concepts. It is my impression that Bill did not, and does not now, consider *Behavior: The Control of Perception* to be complete and final. However, there seems to have been little discussion of possible changes in the original hierarchy. Rather, there seems to have been discussion of various interesting and important applications and related ideas.

Here, I am summarizing some of the ideas that have interested me over the years. Many of these points can and should be developed and, perhaps, modified. Certainly, I consider all of them open to review and discussion. I recognize that they are based on—and limited by—my own experiences and conclusions. To me, much of this material is self-evident. My purpose here is to organize these observations and relate them to the basic concepts of hierarchical control systems, making them available to others for further development. I hope that these ideas and approaches will be found intriguing, leading, in turn, to further modifications, elaborations—and alternatives!

I am minimizing explanations so as to emphasize the overall structure, viewpoints, and modifications of the hierarchy suggested here.

I begin with a *Decision Making Entity* (DME). This concept seems to have been overlooked, except for an implied inclusion in the "reorganizing function." That function, so far as I know, has never been analyzed in terms of its structure, capabilities, limitations, and relation to the hierarchy. Rather, "under what conditions" and "how to improve results" have been studied. These are important, of course, but where and how these activities occur and how they relate to a hierarchical control system are the subjects of this presentation.

Making *decisions is* an everyday occurrence for most of us. Most are routine ("Do you want cheese on your hamburger?" "What's the best way to Chicago?"), but some involve complex analysis ("How do I get funds for this project?" "Who will be willing to act as editor?") and can reveal unexpected conflicts. "Who," or "what," makes decisions and "where" they are made have received little or no attention.

It is tempting to identify this Entity as the "self" or the "ego," or by a similar label. However, such terms tend to include additional concepts such as "personality," "character," and other aspects of the individual. They might include guidelines commonly used by the individual in making his decisions. These anthropomorphic concepts tend to be derived from previous decisions leading to conclusions and assumptions used as the basis of decisions and actions. These are important, of course, but are excluded from the concept of the DME.

The Decision Making Entity, as here understood, can act without being bound by past decisions. It frequently uses them because they are readily available and alternatives may be overlooked. It has the ability to be arbitrary. It can change past decisions if they are accessible to the DME. Access can be limited by a combination of previous decisions. Consistency among decisions is not intrinsic. It is capable of contradictory actions!

Definition of the DME: The Decision Making Entity is defined in terms of its connections (inputs and outputs) and its capabilities.

Connections available: Input A, information about the current condition of physiological systems, including (a) information about the operating condition of the organism, and (b) information about conditions outside the organism. *Input B,* information about past events—memories, recordings, whatever. This includes both verbal and nonverbal events. The distinguishing feature of this connection is that the information is all from past time, although "past" can be very close to "present." These events can range from the remembered mosquito bite to the remembered discussion of "Real Reality," and so forth. *Reference Levels C,* information specifying the acceptable operating condition of the organism, A(a). These are the "intrinsic levels" of other net discussions. *Outputs D,* information acting throughout the hierarchy. Usually, but not necessarily, outputs act by selecting inputs to the higher-order levels of the hierarchy, leaving the details to the remaining lower-order structures. These outputs can be considered from two different viewpoints: as outputs from the DME, and as inputs to the many parts of the hierarchy. Thus, they act as reference levels throughout the hierarchy.

Capabilities: Directing *attention,* including (a) selecting the information, Input A, to be controlled; (b) selecting information from past events, Input B, for comparison with the current situation; (c) comparison of the current and projected ("anticipated") situations with acceptable magnitudes of the variables selected for control—especially intrinsic variables, C, when they are relevant. Everyday situations usually do not directly involve intrinsic variables. Decision making, including (a) selecting outputs (D) to be used by the DME as reference levels for the hierarchy; (b) activating the selected outputs for controlling the selected systems.

Conditions required: In order to direct its attention and make its decisions, the organism must be conscious. Unconscious means that the DME is unable to receive information from its inputs. However, the remainder of the systems can be functional, operating on the basis of the most recent settings of their reference levels. There are several interesting situations that can occur: sleep, coma, paralysis, trauma, etc. These, and others, are worth separate discussion.

These connections define a negative-feedback control system. The stated capabilities are unique to the DME and critical to its operation.

In this view, the feedback signals include two categories of information: A(a) about the current operating condition of the physiological systems, and A(b), about the surrounding environment. These signals are compared to levels selected from memory, B, and applied as inputs to the hierarchy. In addition, the first group, A(a), is compared with the intrinsic reference levels, C, for possible action. The output function consists of the entire hierarchy, D.

The DME is able to direct attention to any group, subgroup, or combination of available memories and compare the projected results with any other combination of available memories, as well as with any related intrinsic levels. It is able to combine selected memories for application as reference levels throughout the hierarchy, as required for the selected action. The entire set of feedback signals available to the DME is the set of *perceptions, controlled by behavior,* as discussed by Bill in *Behavior: The Control of Perception.*

Each of the following contrasting views is useful—"correct," if you please depending on what your purposes might be.

First, *the DME's view of the world*—a stimulus-response (S-R) view. The DME looks "down" its hierarchy for ways to maintain and improve its well-being. The DME acts, like any control system, when it detects a difference between current perceptions and reference perceptions. It examines available alternatives, based on current data combined with projected results of alternative actions. It selects and then applies its selections as reference levels where needed throughout the hierarchy. The DME has no need to "know" anything about the details of the control systems it is using. It merely applies its output signal(s) where needed, and the systems respond. This applies not only within the organism, but equally to using other individuals or groups of individuals as means to accomplish the selected results. This can be as simple as making requests or giving orders—if the others have already internally decided to accept and act on such requests/orders. The DME acts like an individual trying to maintain and improve his/her circumstances.

Second, *the hierarchical control system's view of the world*. Viewing the world in terms of hierarchies of control systems covers an amazing range of observations and leads to additional study and analysis. The concept of higher-order control of lower-order systems through setting their reference levels is particularly simple and useful. In this view, "behavior" consists of counteracting, or opposing, any disturbance of its controlled variables. That is, it consists of detecting a difference between current perceptions and reference perceptions. A disturbance, uncontrolled at one level, tends to result in disturbance at a higher level. However, this structure has no way to change its reference level—nor does it have a way to change its organization. Its memories are retained in the form of established structures and fixed reference levels. It cannot examine memories with a view to selecting alternative ways to achieve its control. Its high-order reference levels are based on "remembered" events, but there is no way to "project" or "anticipate" alternatives for possible application to a given situation.

This *problem* arose in our early discussions as we sought to define higher levels. How can the changing behavior of an individual be described when he or she is blocked? Analysis working upward through the lower orders assumed (implicitly) a set of fixed reference levels, especially at the higher levels.

How could these be changed? How could the system be "reorganized"? An ad hoc "reorganizing system" was proposed. Without actually being stated, its definition amounted to "whatever is needed in order obtain these results." Bill has discussed this concept in several places, but it seems to me to be incomplete. I am familiar with this view, since I was deeply involved in the early developments leading ultimately to *Behavior: The Control of Perception,* as well as to papers and discussions among others. In my own life, I have found and continue to find this viewpoint very useful in many ways. But it is not the only view that I find useful.

In fact, what seems to be needed is the DME as suggested here. I find that combining the DME approach with the hierarchical view provides some additional answers and leads to some revisions of the hierarchy.

The overall objective is control of perception, as genetically required. A hierarchy of control systems is the means to that end. In the following suggestions, the guiding concept of the hierarchy, "higher-order goals are accomplished through setting reference levels for lower orders," is retained.

I. *Zero Order—Intrinsic Systems.* These are the physiological systems underlying the Decision Making Entity. They include all systems providing neural inputs directly to the DME. They report the operating condition of the organism for comparison with intrinsic levels for control action through the hierarchy. Some of these systems might, themselves, be feedback control systems (I have in mind some of the hormonal systems), but they are controlled only indirectly through the

hierarchy. Zero Order systems also include direct neural signals representing the conditions of the external environment—typically, the usual five senses.

II. *First Order.* Control of individual muscles (or muscle fibers, if you prefer). This remains essentially the same as the original First Order. The signals, of course, are neural intensities serving as feedback signals derived from the tensions of the individual muscles. Bill *(Behavior: The Control of Perception,* pp. 82 ff.) discusses this in depth from several standpoints. From the DME's view, this Order controls individual muscles. It is a "follower" system—it simply reproduces (within its capabilities) the reference signal(s) provided.

III. *Second Order.* Control of "configurations." At this level, they are considered "static," that is, temporal variables are unspecified. Combinations of muscle systems are typical examples, but this Order need not be limited to muscle systems. Our original concept, elaborated by Bill *(Behavior: The Control of Perception,* pp. 99 ff.) "Second Order, Sensations" appears to combine parts of "Zero Order" and "First Order" as presented here. In the present treatment, Second Order is pretty much the same as Bill's "Third Order, Configuration Control" *(Behavior: The Control of Perception,* pp. 115 ff.). Muscle systems are convenient and typical examples. Bill includes the perception of "objects" within this category *(Behavior: The Control of Perception,* pp. 125 ff.). He notes the "invariance" of combinations of sensations that can be perceived as "objects." "Invariance" implies, at the minimum, some degree of short-term memory to provide continuity—invariance. From the DME's view, this Order does not include control of objects. In general, each configuration would be multidimensional, expressible in vector or matrix terms if desired. As output systems controlled by the DME, these are "follower" systems like First Order systems.

IV. *Third Order.* Control of sequence. Bill *(Behavior: The Control of Perception,* Chapter 11) assigns this to Fifth Order, much as we did originally. Bill places a Fourth Order, "transitions," ahead of his Fifth Order. In the present treatment, Third Order controls the sequence of Second Order, static configurations, much like the frames of a movie. The frames could be rearranged by the DME by means of Third Order systems. Of course, there are intrinsic limitations, but the concept remains. These are also "follower" systems.

V. *Fourth Order.* Control of temporal variables. As I recall our discussions, these variables were never made explicit. They seem to have been included within sequences, although no direct statement was made to that effect. Bill (Behavior: The Control of Perception, Chapters 10 and 11) seems to include these variables implicitly without recognizing them. His concept of "transitions" also seems to include temporal Variables implicitly. Here, "temporal variables" refers to such items as "fast," "slow," "tempo," "frequency" (of oscillations), and the like. For example, the DME can apply the same tempo to a variety of situations—it appears to be an independent parameter of systems in action. The importance of these variables seems to be generally taken for granted—but otherwise ignored. I have found it very useful to pay attention to, and control, this Order of variables.

VI. *Fifth Order.* Control of—and selection of—skills. Typically, muscle skills. Skills consist of temporal sequences of configuration produced by combinations of muscle tensions. Control of variables of speed, tempo, and other temporal variables is important. This Order concerns individual motor skills. This is also where the perception of "objects" belongs. The world and its contents are treated as a multiplicity of inanimate objects. For example, bowling requires a ball and an alley, while using a sequence of positions performed with selected timing. Simpler skills include walking, running, etc. These are motor skills where nothing is needed but the physical equipment and the DME's decision. This is a general characteristic of Fifth Order systems.

VII. *Sixth Order.* Control of interpersonal relationships. This Order recognizes the differences between inanimate objects and independent active entities. This includes animals and, especially, people. The DME seeks its objectives through controlling these independent entities. Often, it acts as though they were stimulus-response systems. This frequently works, since many of these otherwise independent entities have decided to accept suggestions and requests as commands. When it doesn't work, the DME seeks alternative methods to reach its goals.

Control of communication could be considered for designating Sixth Order. However, this would focus on the skills used, Fifth Order, rather than the goals of Sixth Order. Communication is essential to the control of interpersonal relationships. Without some form of communication, other individuals are treated as inanimate objects.

Study of the content of communications, whether nonverbal or verbal, can help clarify the Orders of the hierarchy as well as perceptions of the DME. Topics discussed by people in everyday conversation and items published in the media are useful for this purpose.

Modes of Sixth Order. The topics communicated can be grouped according to the levels of the hierarchy:

Zero Mode of Sixth Order. Illness and similar topics are very common. People often have little knowledge of their own anatomy and physical structure. But they talk about it a great deal.

First Mode of Sixth Order. Aside from reports of sore muscles, there seems to be little direct discussion of muscle systems.

Second Mode of Sixth Order. Configurations appear as comments on "posture," positions needed for various skills, and the like.

Third Mode of Sixth Order. Here is discussion of the sequences of

configuration needed to obtain desired results. This includes sequences of positions forming movements required for a skill.

Fourth Mode of Sixth Order. Variables of tempo, rhythm, etc. Aside from discussions of sports events, musical concerts, and the like, these variables seem to receive little explicit attention.

Fifth Mode of Sixth Order. Much attention is directed to all sorts of muscle skills. Generally, several lower-order considerations are discussed, although not always explicitly. "How-to" books are very popular, usually involving most of First through Fifth Modes of Sixth Order. This Mode includes all concepts of the nature of the physical world. Math and theoretical analysis are also here. Everyday discussions commonly show very little understanding of present day physical science, math, and experimental methods. Well-known errors in logic are commonly accepted as valid.

Sixth Mode of Sixth Order. This Mode concerns methods and topics of communication among individuals and groups. Discussion of these topics, of course, uses the lower Modes as needed. Illustration from personal experiences are frequent (Third, Fourth, and Fifth Modes). Reports of activities of public and private individuals and organizations are common. Examples, analogies and similes are used very frequently. Public speaking and teaching skills and methods lie in this Mode. Rules and regulations used to establish acceptable performance appear here. Games, organizations, social customs ("social controls"), laws, police, etc. are within Sixth Mode of Sixth Order. As Fifth Mode does for the Physical World, this Mode includes all theories of behavior—whether magical, mystical, intuitive, or scientific, whatever that means! Everyday routine communications reflect the concepts of popular behavioral theories. Perceptual Control Theory also is in this Mode.

Seventh Mode of Sixth Order. This Mode concerns control of one's own behavior to accomplish higher-order objectives. It uses the concepts and methods of Sixth and Lower Orders for these purposes. Although there is relatively little discussion of these subjects, it does occur. This Mode has a corresponding Seventh Order of control.

VIII. *Seventh Order.* Self image and DME. Self image includes all aspects of the individual's capabilities and organization. To examine one's self image requires review of one's remembered actions and interactions as they relate to one's view of individual behavior. This review would tend to include, but perhaps not "require," extrapolation to possible future situations and events. Such imagined results can be compared to objectives at all levels, with underlying emphasis on intrinsic levels.

This discussion leaves a great many questions unanswered, and equally many fascinating subjects for investigation. I hope that this condensed outline and analysis will be found useful.

Greg Williams: Bob says: “The DME is able to direct attention to any group, subgroup, or combination of available memories and compare the projected results with any other combination of available memories, as well as with any related intrinsic levels.” Why might the DME direct attention to certain memories, rather than others, at some particular time? Bob, do you have a theory of attention “selection” other than the broad viewpoint that the DME tries to “improve its well-being”? Is there some calculus for tradeoffs among various possible ways to “improve” (more or less)?

It seems to me that your proposal would require the DME to run “imagination connection” trials on the alternative actions at a particular time to “examine” and “select” some of them for actual performance. Maybe the DME wouldn’t need to “know” details about the parts of the hierarchy which would then actually be used, but it appears that it would have to be able to “see” the results of such use “in imagination,” prior to actual performance, in order to have a basis for decision making. Or do you have different notions about how the selection process occurs? I’m trying to understand the basics of your model at this point; perhaps I’m headed in the wrong direction. Please clarify.

It appears to me that the DME is basically directed (not completely random) reorganization. Is that a fair characterization? In the past, I’ve been attracted to the idea that there are *both* random and directed types of reorganization possible in humans—the former can get you to a solution (eventually, usually) when the latter has no clue on how to direct, but when it works, the latter is usually quicker. The problem has been in figuring out a working model for directing—hence, my questions above. Fleshing out the mechanism(s?) of your DME’s decision making would be very helpful.

Rick Marken: Bob proposes an addition to the Hierarchical Perceptual Control Theory (HPCT) model: a Decision Making Entity (DME). I don’t know what data motivate the addition of a DME, but perhaps it has to do with Bob’s claim that “Making *decisions* is an everyday occurrence for most of us.... ‘Who,’ or ‘what,’ makes decisions and ‘where’ they are made have received little or no attention.”

I prefer to look at decisions as the conscious results of conflict. So the cause of decision making is already a part of the model. So is the means of dealing with conflict: reorganization. We flip a coin and do one thing (produce one perception), and we tolerate the error resulting from not doing the other (producing the other perception). A better way to solve such conflicts is to “go up a level,” one of the great therapeutic experiences (speaking subjectively) one can have, and a sure cure for the everyday conflicts (decisions) that are the natural result of never achieving a perfectly organized control hierarchy.

The study of decision making has been popular in conventional psychology because it is an inherently statistical phenomenon. If you offer people choices between almost equally attractive perceptions, coin-flipping (statistics) is the only approach (if you don’t go up a level and see the choice itself as arbitrary; but if you did that, you would be kicked out of the experiment). There might be something interesting to be learned about hierarchical control and reorganization through the study of decision making (conflict). But I think we must have very good models of the “elements” of decision making—conflict, in particular—before we can make a coherent stab at decision making (which, as I said, is probably reorganization—of the conscious variety—to settle, not necessarily resolve, a conflict).

I think that I could get a better grasp of Bob’s DME proposal if he (or anyone) could propose some experimental tests for evaluating this addition to HPCT.

Bill Powers: Bob, you say: “It is my impression that Bill did not, and does not now, consider *Behavior: The Control of Perception* to be complete and final.” Right you are. You go on to say that “there seems to have been little discussion of possible changes in the original hierarchy. Rather, there seems to have been discussion of various interesting and important applications and related ideas.” I’m glad to see you opening up the discussion. ‘The” hierarchy is a figment of my imagination, building on *our* imagination. For most of the levels I’ve proposed, the only backing for the specific definitions is anecdotal and subjective. As far as I’m concerned, these or any other levels won’t be “facts” until we have put them to experimental test.

I’ve always felt that defining the levels scientifically is a large project which should begin by experimentally verifying that people can control variables of many different kinds—anything anyone can think of, without regard to levels. Even the most obvious variables should be put formally to The Test, just so we can write down the parameters of control and say that we have in fact observed such-and-such a variable to be under control by a human being. This would be a beautiful thesis, or series of them. On the other hand, maybe it should be the kind of project to which all control theorists contribute, the way astronomers put in some duty-time measuring double star angles and separations (the three well-spaced observations required to determine the orbital elements, needing, in many cases, 1000 years to complete).

Once we have a base of hundreds of certified controlled variables, we can begin to try to put them into order. If there really is a hierarchy, the variables will fall into classes, and the classes will be related in a hierarchical way. That is, *in* order to control a variable of one level, it will be necessary to vary a controlled variable of a lower level. And of course the only way to vary a controlled variable arbitrarily is to alter the reference signal for the system that's controlling it. Showing that this is the case leads to a new series of experiments.

When this project is done, will become a science. We will have advanced from Galileo to Newton.

In the meantime, of course, we can argue. But without experimentation, arguments are just a pastime.

I agree with Rick Marken about decisions: they represent conflicts. Unless there were at least two competing goals to satisfy, there would be no need to make a decision. You would just do whatever is required to achieve the single goal. More commonly, there are multiple goals involved in behavior, but we have learned to organize our actions (as a result, largely, of resolving conflicts in the past) so that all of the goals can be satisfied at once. When that is the case, again no decisions are needed.

At the level I call "programs," symbol-handling processes occur which I characterize as a network of choice-points. There are tests for conditions, with the choice of a branch being determined by a rule applied to the results of the tests. The term "choice" seems to imply a decision, but in fact there are no decisions at this level either. The conditions encountered at each choice-point, plus the rules, completely determine the path to be followed next. Only when there is ambiguity or when the rule is self-contradictory (calling for more than one mutually exclusive path to be followed) is anything like a decision required. If you have an algorithm for making decisions, you don't have to decide anything!

Note that operations occurring between choice-points are sequences, lists of reference levels to be brought about in order. Sequences are the next level below programs. Programs are concerned *only* with applying rules to select branches, as I use the term here. They involve "flow control," as they say in programming manuals. The parts of computer programs that consist only of one instruction following another belong at the sequence level here, not the program level.

If we eliminate programs—the execution of algorithms for choosing paths—from decision-making, what is left? As far as I can see, only the cases in which for some reason we wish to do two contradictory things at once. At that point, we must reorganize or simply suffer the paralysis of conflict.

Bob, I think that you and I agree on this, at least to an important extent. You say: "The Decision Making Entity, as here understood, can act without being bound by past decisions. It frequently uses them because they are readily available and alternatives may be overlooked. It has the ability to be arbitrary." This arbitrariness has the flavor of reorganization. But so far, at least, I have not considered systematic re-

organization. Anything that could be called systematic, it seems to me, belongs in the already organized hierarchy. At the level of logic, systematic consideration of previous choices and possible alternatives is an algorithm. As such, it can be reduced to rules governing selection of paths connecting sequences or lists of behaviors, where by behaviors, I mean controlled perceptions of the consequences of acting. When we remove all algorithms by putting them into the program level of the hierarchy, all that is left of decision making is the arbitrary part: making a change for no reason.

You basically say that the DME receives perceptions either from lower systems that are in the normal mode, receiving information that comes ultimately from interoceptive or exteroceptive sensors, or from lower systems that are in the imagination mode, deriving their perceptual signals from memory. Then you speak of *"Reference Levels C,* information specifying the acceptable operating condition of the organism, A(a). These are the "Intrinsic Levels" of other net discussions." This makes the DME look even more like the reorganizing system, with reference signals specified genetically. You also speak of *"Outputs D,* information acting throughout the hierarchy." This is typical of the reorganizing system as I perceive it. However, I think that your DME includes both learned hierarchical systems and the unlearned system that I call the reorganizing system. When a "decision" is reached, it must entail some sort of action, and to create any systematic action, a higher-level system must adjust reference signals for lower-level systems. Furthermore, since nature never trusts an organism's output to do what it is supposed to do, the consequences of the action must be perceived by the level issuing the reference signals, so that the reference signals can be varied until the perceived result is the intended one. If this control process takes place in an organized way, it must be due to a learned system.

In my concept of the reorganizing system, I have extracted the arbitrary non-systematic kind of action from the hierarchy as a whole and given it a separate existence of its own as a built-in aspect of the organism that functions from the beginning of life. We used to call this the Negentropy System. I gave up the word because it implies things I don't believe. I now just call it the reorganizing system.

Regarding capabilities, I will accept as part of the reorganizing system the direction of attention. The rest I have incorporated into the hierarchy itself. I am not sure what attention is for. We need to do some experiments to find out. The comparison of current and projected magnitudes of variables with acceptable magnitudes is simply the operation of any control system at any level ("projected" magnitudes require the imagination connection). That kind of operation is adequately handled by the "canonical" control-system diagram and, when intrinsic variables are not involved, is simply the operation of the learned hierarchy of control systems. I allocate intrinsic variables and intrinsic reference signals strictly to the reorganizing system, whose actions are arbitrary and random and serve to alter connections and weights in the learned hierarchy. That is how the learned hierarchy becomes organized.

With regard to your "Decision making, including (a) selecting outputs (D) to be used by the DME as reference levels for the hierarchy; (b) activating the selected outputs for controlling the selected systems," I handle all of this in the higher levels of the hierarchy, but I leave decision making (as an arbitrary process) out of it.

You say: *"Conditions required:* In order to direct its attention and make its decisions, the organism must be conscious. Unconscious means that the DME is unable to receive information from its inputs. However, the remainder of the systems can be functional, operating on the basis of the most recent settings of their reference levels. There are several interesting situations that can occur: sleep, coma, paralysis, trauma, etc. These, and others, are worth separate discussion." I have formed a similar idea of consciousness (beginning with our discussions of 35 years ago). However, I begin with awareness (which I think you include). Awareness is the capacity of the reorganizing system to receive information, regardless of its kind. I have proposed that when awareness is receiving information selectively from a portion of the hierarchy, the result is what we call consciousness. This allows us to distinguish between one phenomenon that remains the same no matter where it is applied—awareness—and another that changes its form depending on the source of perceptual signals received in awareness—consciousness. Consciousness always takes on the character of the control systems to which awareness is connected.

Thus an apparent rule that seems to fit experience: you cannot be conscious of systems that are in the conscious mode. Instead, you are conscious of the lower-order world of perceptions received by those systems, and you experience those perceptions with the conscious interpretation typical of the level (or levels) at which awareness is connected. This interpretation appears to be an objective property of the world.

Any system in the hierarchy can operate in the conscious or unconscious mode. In the conscious mode only, it is subject to reorganization.

Implied by this model is the possibility that awareness can be selectively connected to particular levels in the hierarchy. When that is the case, you experience the world consciously as that level perceives it, but you are unaware of applying any interpretation to the perceptions. Instead, you see those perceptions simply as part of the world. If you are operating in the logic or program level, you see the world as full of choice-points and alternatives, with natural rules that define a path through the choice-points. On the other hand, if you are operating in the relationship level, you see a world in which everything is related in some way; you see the constraints that make independent objects and events maintain a certain constancy of interaction.

And while you are attending from the viewpoint of relationships, you are not aware of any higher levels of perception and control. They are still operating, and if you ask yourself why you are paying attention to relationships, you will come up with higher-level reference signals—what you hope to accomplish by attending to relationships. That is, you can often "go up a level" and realize that higher-level control processes were active all the time, even when not in consciousness. But as soon as you do that, you are no longer seeing a world of relationships. The nature of the conscious world changes as you move awareness from level to level.

I think that this proposal is related to your concept of "modes." However, I do not see these modes of consciousness as being modes of just one level, your DME. I see them as resulting from awareness moving from one place in the hierarchy to another. When one is attending to a lower level of perception, higher processes are still operating but they are not operating consciously. By your postulate, all modes would entail consciousness of the highest-level processes. Maybe you're right. But I think experience argues against this view. At any rate, I think your picture is worth trying on for a fit.

You say: "This *problem* arose in our early discussions as we sought to define higher levels. How can the changing behavior of an individual be described when he or she is blocked? Analysis working upward through the lower orders assumed (implicitly) a set of fixed reference levels, especially at the higher levels." This is no longer a problem in the hierarchy as I currently conceive it (since 1973). Higher reference levels are no longer fixed, except at the highest level. At intermediate levels, lower-order reference signals are varied as needed to provide a higher-level system with the perceptions it needs to match its own reference signals—which in turn are being varied as required by still-higher systems. You might ask Rick Marken for his spreadsheet demonstration of this arrangement; it will run on Lotus 1-2-3 or Excel. It shows how a three level hierarchy with six systems at each level can simultaneously control three levels of perceptual variables, despite random disturbances from the environment, and despite considerable interaction among the controlled variables at each level.

You ask: "How could these be changed? How could the system be 'reorganized'?" They (reference levels) no longer require reorganization to be changed. Reorganization is now needed only when the learned systems are not capable of maintaining intrinsic variables at their reference levels (as a byproduct of their actions). Since the model now includes many "intellectual" functions such as classifying, ordering, reasoning, application of principles, and control of system concepts, all of which are learned, the reorganizing system does not have to carry out any rational processes.

You say: "In fact, what seems to be needed is the DME as suggested here. I find that combining the DME approach with the hierarchical view provides some additional answers and leads to some revisions of the hierarchy." I think you will find that the levels I have added (categories, sequences, programs, principles, and system concepts) contain much of what you want to put into the DME. I agree that such functions are required. I have simply broken them out into specific levels of functions, while reserving the arbitrary reorganizing part to a separate non-hierarchical system. I don't say that's right. It's just what I have done.

I have several arguments with your proposed levels, but will pass them up for just one clarification concerning your Fourth Order, temporal variables. For quite a long time after we parted, I considered just the sequence level in the position where you put it, above configurations. Then I realized that there are really two kinds of sequence variables, one exemplified by the second-hand of a dock, and the other by the notes of a melody.

The second-hand of a clock gives rise to a perception of continuing angular motion, d/dt(angle). With angle as a configuration perception, the new perception is simply its time derivative. As such, it has a value at all times, in present time, which can change in magnitude as the angular (or other) velocity increases or decreases.

This is quite different from the temporal progression in the successive notes of, say, "Taps," which can be played slowly or more quickly. In the case of the melody, there is no simple motion signal, but the sense of a specifically ordered progression of different sensations, one following the other. What matters is not so much the speed, but the ordering in time—which note follows which.

On realizing this difference, I introduced the "transition" level, which is basically derivatives (and perhaps derivatives of one variable relative to another). This level went just above configurations, and is where stroboscopic as well as continuous motion or change is perceived. That left the sequence level to cover just the temporal ordering of lower-level variables, including transitions. I called this the "event" level, where an event was supposed to be a short familiar temporal pattern of perceptions of transitions, configurations, sensations, and intensities (you omit intensities).

Only a couple of years ago, Gary Cziko brought up some more examples of temporal variables in which only the ordering is important—in language, for example, the ordering of words. Here the temporal pattern is not evident, for an ordering is quite independent of how long it takes elements to occur and of the spacing between elements. This struck me as different from an event, in which there is a stereotyped unitary pattern that forms a single package in perception. So the sequence level ended up being split once again, the event level now meaning only brief "packaged" temporal patterns recognized as single things like the bounce of a ball, with pure sequential ordering—lists—being moved to a higher level.

We can discuss the rest of your proposals for levels later. I expect that others will have questions and comments, too. I am glad to see the subject opened up again, because I don't like the sensation of having my hypotheses converted into Gospel. I think that by trying to boil down all propositions to the basic underlying operations and connecting them with experience, we can arrive at an agreeable set of levels for experimental test. Maybe the reason that there has been so little questioning of my definitions is that nobody saw any conflicting alternatives and thus didn't feel compelled to make a decision!

Bob Clark: Greg, as usual, it is much easier to raise questions than to answer them. This is to be expected, since the questioner regards the world from his or her own viewpoint combined with his or her available store of ideas.

Also, Rick prefers "to look at decisions as the conscious result of conflict.... the means of dealing with conflict: reorganization." He notes: "A better way to solve such conflicts is to 'go up a level."' Rick further seems to accept the view that decision making "is an inherently statistical phenomenon."

In response to your questions, comments, and remarks, let me point out that I am concerned with the process of decision making. Statistically, the only common element among people is that they all make decisions. Methods, reasoning, procedures, etc. can differ drastically from one person to another.

To me, decision making is a peculiarly individual matter. For this phrase to have meaning, there must be at least two alternatives available. This implies at least a minimal conflict in that they cannot both be selected. The alternatives need not be particularly important (although they could be). There must be some way in which they can be examined. There must be some basis for selection. And there must be some entity capable of putting all of this together.

We already have all of these elements except the Decision Making Entity—which is implied in Chapter 15 of *Behavior: The Control of Perception,* since something somewhere must operate the diagrammed switches. In fact, my view of the relations among the DME, current perceptions, and memories might be regarded as an extension of the concept Bill illustrated with two single-pole double-throw switches. Using this diagram, these switches are controlled by the DME—one of its major functions.

In addition, the Recording Function (would this term be better than "memory"?), can be considered a multi-dimensional recorder, including not only all perceptual signals, but also all consciously imagined combinations, projected conclusions, and decisions. When examined (imagination), the memories are much like multi-dimensional video tapes. These memories are not necessarily logically related, nor otherwise coherent. They might arise simply through some accidental event that provides some connection (a "reminder") to the specific remembered event. It could be an odor, a face, a sound, an idea, a word, etc. Or it could be a problem ("conflict"?) with recognizable aspects bringing related memories to mind.

The Recording Function is mostly undirected, but the DME can make it more easily available by consciously assigning some kinds of labels to suitable recordings. How do you learn the name of someone you have just met? However, many "labels" are acquired more or less accidentally. Thus, the word "chocolate" easily brings an image (images?) to mind. But there are many forms of labels: the appearance of a house, a date, a period of time, and so on. There are many ways to locate specific memories.

Most memories are inactive most of the time. (What a confusion, otherwise!) They appear to be "forgotten" until some form of "reminder" occurs. (I have been intrigued by the questions that "pop" into mind in response to answers on the television show "Jeopardy.")

A great many of the decisions needed are simple, requiring very little attention or analysis. A very rapid (perhaps on the order of a few milliseconds) switching might occur between alternative memories when little analysis is needed and anticipated conclusions are quickly formed. These alternatives are selected for their relevance, often simply by being "reminded" of similar situations. But the DME might find more thorough investigation necessary in seeking a satisfactory selection. This tends to be related to the level in the hierarchy involved—as implied by Rick Marken's remark about "going up a level." Thus, in ordinary situations, the DME can make its selection quickly.

So the Recording Function records conclusions and decisions. These are more likely to pertain to high-level situations, where a previously formed conclusion/decision can provide a quicker response. "I'll push the button the instant I perceive a flash of light," rather than, "There's a flash of light, what do I do now?" Or, "The moment the light turns green, I'll hit the throttle," which results in a fast response.

The DME's attention is directed by the need to select among alternatives. The characteristics of the situation serve to remind the memory

which recording to present. It is interesting to observe that the DME cannot direct its attention to its own acts as it is performing them. Its only information about its own activity (self-knowledge?) is through examining the relevant memories. These are not necessarily readily available.

By and large, these are mostly ordinary and familiar observations, but they seem to have been left out of much behavioral discussion.

Greg, you note: "It appears to me that the DME is basically directed (not completely random) reorganization." In its origins, the Reorganizing Function was proposed to explain the observation that individuals change their behavior when faced with a conflict—especially if it is hard to resolve. However, that is an "outsider's" viewpoint. To the DME involved, it is not intrinsically different from any other decision-making situation. Available alternatives (including, perhaps, violent movements or whatever) are reviewed, and the behavior offering the most promising anticipated results is put into operation. To the outsider who might not even suspect the alternatives available, this will tend to appear "random," that is, "unpredictable."

Bill, you start by saying: "The' hierarchy is a figment of my imagination, building on our imagination." And you require an "experimental test.... verifying that people can control variables of many different kinds." Here we have two concepts (at least): "variables" and "control." In *Behavior: The Control of Perception,* "variable" is defined as "an entity identified by characteristics of the location at which it is measured, and having a number or a continuum of detectably different states. A meter-reading associated with a physical phenomenon." My *Webster's Unabridged Dictionary* includes (after eight definitions as an adjective) "noun. 9. something that may or does vary; a variable feature, factor, or the like, 10. Math. a. a quantity or function that may assume any given value or set of values. b. a symbol that represents this. 11. logic. (in the functional calculus) a symbol for an unspecified member of a class of things or statements." Also: "vary" specifically includes the concept of "change" or "alter."

Here we have two distinctly different definitions: one, a "feature," "characteristic," or the like that can have a changing "value" or "magnitude"; and two, a "member of a class." The first implies continuing identity of the variable, the second is not concerned with possible changes in the individual "member," as long as it qualifies for its "membership."

All of these verbal structures (theories) are critically dependent on mutually accepted definitions. I think we both, with physical science and math backgrounds, tend to use the first definition. However, those with backgrounds in psychology, sociology, etc. tend toward the second.

Here I think we are considering a "variable" to be (from the Glossary in *Behavior: The Control of Perception)* a 'Perceptual Signal: The signal emitted by the input functions of a system; an internal analogue of some aspect of the environment." I think you would include "combinations of perceptual signals represented by neural signals" as additional variables. I would also include signals arising from memory. All of these are generally available to the DME.

Can people "control" such signals? In some ways, this is the very heart of the subject of behavior—we see people doing things to themselves, to their surroundings, and to other people, followed by assorted consequences. Note the temporal implication of "doing things." Is this "control"? Perhaps, but it seems to me that "control" implies "intent" to produce some change. The act might be inappropriate, or ineffective, but if the result is somewhere near the intention, some degree of "control" appears to have been achieved. This seems to be circular, and "begging the question." That is, if the result has not been achieved, there has been no control. On the other hand, if the result occurs, it could have been purely coincidental.

How do we measure "intent" so that it can be related to "result"?

Correlation alone is insufficient evidence. There are many examples of strong correlation between variables, without need to assume "intent." Consider the phototropism shown by many plants. The plant bends toward the light and its well-being improves. Fairly well-identified laws of biochemistry and physics are sufficient explanation. "Control" here? I don't think so. It appears to me that "control" is a peculiarly human concept that becomes pertinent only after certain stages of development. Interestingly, it seems to me that a great deal of history, philosophy, ethics, and psychology (in the generic sense) relates very directly to matters of "control." Thus we have agriculture for food, clothing for insulation, weapons for game (and warfare!), politics, science, and theories. Always someone is seeking to act in such a way that the things he/she cares about (his/her "perceptions," of course) become, and remain, somewhere near the states he/she seeks for them.

This is the essence of a feedback control system and is completely familiar to nearly everyone. *Except* that they look at the pieces of the loop without putting it together as a system! How many people understand the operation of a thermostatic control system in feedback-system terms? They use such systems as "on/off" switches. And they are usually satisfied. They just haven't learned the viewpoint and somewhat specialized language that PCTers find useful. I find them useful too, and I believe most others will, as well, if they take the trouble to learn at least the key parts of PCT.

Bill, you "agree with Rick Marken about decisions: they represent conflicts." To me, "represent" signifies equivalence; that is, when A

"represents" B, they are interchangeable. To me, "decisions" are quite different from "conflicts." In an earlier post, I pointed out the need for alternatives in order for a "decision" to be needed. I agree, of course, that when "all of the goals can be satisfied at once.... no decisions are needed." And none are made. Yes, conflicts imply needs for decisions—and decisions that have been made imply preceding conflicts resolved. Note, incidentally, that some conflicts are very incidental and are settled easily and quickly, so that the operation of the DME can easily be overlooked.

Bill, you say: "At the level I call 'programs,'... processes occur... a network of choice-points.... choice of a branch... determined by a rule... The term 'choice' seems to imply a decision, but in fact there are no decisions at this level either. The conditions... completely determine the path... Only when there is ambiguity or when the rule is self-contradictory... is anything like a decision required. If you have an algorithm for making decisions, you don't have to decide anything!" Yes, of course. But surely these networks, algorithms, etc. come from somewhere and are retained somewhere for application if needed. I call this "somewhere" the memory, created by the Recording Function. The choice-points you speak of, I would term past decisions remembered and applied in the present: recordings of decisions made at some previous times and retained in effect unless considered for review.

You say: "If we eliminate programs... from decision-making... we must reorganize or... suffer... conflict." Well, yes—that's what I'm talking about. Whenever decisions have already been established, whether by accident or careful study, no decision is needed, and the DME's attention is directed elsewhere. Of course, the pre-existing decision might be found wanting and need revision. Then the DME would reexamine the situation.

Regarding your reorganizing system, I point out the DME has the ability to be arbitrary. Indeed, it does have the "flavor of reorganization." After all, the organization to be reorganized consists of a network of previously made decisions. Note that the DME has access to all recordings and can project—"imagine"—anticipated outcomes that can be used as the basis for decision, ultimately using intrinsic reference levels as criteria. If previous selections have not worked out, the DME can arbitrarily select an alternative or alternatives.

Bill, you say: "This is typical of the reorganizing system as I perceive it." I did not set out to examine either the reorganizing system or the process of reorganization. I was intrigued by observing that I am very frequently making "small" (?) decisions.

There seems to be some question as to exactly how much, or "what," is included within the DME and/or the reorganizing system. I recall, of course, the Negentropy System. And I agree—it implies too much.

Indeed, to me, the same is true of the "reorganizing system." The term seems to imply some kind of complex, highly structured entity operating in mysterious ways to resolve conflicts—especially those involving intrinsic systems. I like your phrase "a built-in aspect of the organism that functions from the beginning of life." "Life" might be too broad—there are many forms of life that seem (to me) to lack some of the structures necessary for meaningful decision making. (Specifically, some form of Recording Function, and perhaps other items must be available for operation of a rudimentary DME.) I am tempted to suggest the need for some form of central nervous system as a minimum prerequisite.

I think our concepts in the areas of attention, awareness, and consciousness are quite similar. The biggest difference seems to relate to your treatment of the reorganizing system, which seems to include memory, imagination, anticipation, and, possibly, selection (among alternatives).

Bill says: "I think that this proposal is related to your concept of 'modes.' However, I do not see these modes of consciousness as being modes of just one level, your DME." These statements indicate that I have failed to communicate the locations of "modes" within the hierarchy and of the DME. The DME, as I see it, is not a "level" of the hierarchy. It is a separate entity, operating *on the* hierarchy, which is itself located within the memory. Bill, you speak of awareness moving from one place in the hierarchy to another. It is the DME that is "aware."

"Modes" are suggested as a convenient way to regard the contents of Sixth Order. They help to identify and organize the contents of Sixth Order. Defining Sixth Order as "control of interpersonal relationships" opens the door to a great variety of possibilities. With communication clearly an important aspect of interpersonal relationships, it occurred to me to observe the content of everyday conversations (in terms of Orders of control). I was surprised and pleased to find the concept of Modes very helpful in clarifying my ideas about higher levels. The content and structure of Sixth Order and higher orders need further refinement and discussion.

Bill says: "Higher reference levels are no longer fixed, except at the highest level. At intermediate levels, lower-order reference signals are varied as needed to provide a higher-level system with the perceptions it needs..." This is pretty much a restatement of the hierarchy concept. This works nicely as long as the structure is fixed, and the results are acceptable. Thus, these portions of the structure can run on "automatic," and the attention of the DME can be elsewhere.

And he says: "Reorganization is now needed only when the learned systems are not capable of maintaining intrinsic variables at their reference levels." They include "many 'intellectual' functions such as classifying... all of which are learned—the reorganizing system does not have to carry out any rational processes." *Exactly!* And none of them are "put into the DME." Where are they? To me, they are in the memory—indeed, they, the learned systems, are much of the content of the memory. To a large degree, what you have called the reorganizing system, I prefer to call the DME, with access to the memory and minimal content and capabilities. The DME can perceive recordings, can compare them, can project their implications—by using methods, rules, algorithms, etc. taken from other recordings. The rationality of the DME is determined by the content of its available recordings. In some situations, existing recordings might include "rules" that interfere with the availability of needed recordings and, perhaps, the DME's attention-directing capability.

One of my primary suggestions for modifying the levels of the hierarchy involves temporal variables. I find the treatment of time as an underlying independent variable necessary throughout the structure. At lower levels, it is implicit, because Third Order events, actions, etc. occur "over time," but that is taken for granted. As I examined sequences for parameters for their control, I found that temporal variables became evident. Bill suggests "transitions" as an alternative. This is a logically appealing category. However, such an abstract category does not suggest specific controllable variables to me.

In human activities, I find many temporal variables of importance. Tempo is certainly one, but so also are such items as rhythm, acceleration, deceleration, pauses, and delays. The lowest level where these variables become important is for manual skills (my suggested Fifth Order). Throwing a ball not only requires a certain sequence of configurations of selected muscles, etc., but the timing must be correct! Indeed, change the timing, and you get a different result! Fifth Order, skills, requires control of temporal variables in addition to sequences, etc.

Bill, you refer to Gary Cziko and "temporal variables in which *only* the ordering was important—in language." In communication of nonpersonal matters, this (the *"only")* might be true. When personality, attitude, intent, motivation, etc. are part of the communication, spoken language requires control of temporal variables. Consider temporal variations in emphasis, tempo, rhythm, pitch, loudness, and enunciation (the mechanics of speech) for spoken language.

It seems to me that there are several situations where control of temporal variables is critically important. For example, skills in speaking. When working with adult stutterers (with McFarland), it became apparent to me that some forms of stuttering result from inadequate control of the time relations among the vocal systems.

Time scale is a fundamental concept—I don't know just where it belongs. But consider the changes in interpretation and analysis when the time scale moves from microseconds, to milliseconds, to tenths of seconds, to minutes, to hours, to days, to decades, etc. Consider the fastest-acting neural control system on a millisecond scale. Sequences of events are easily traced around the loop in terms of a series of straight-through operations. Change the time scale to fractions of a second, and, behold! It works as a unit—a feedback control system! *Both views are right!* When larger-scale, slower interactions are examined, games, perhaps, or economics, one can follow the signals as they follow their pathways and interact to become an operating control system. Problems that arise tend to be solved by changing the parameters of the system components. Or by changing the connections. The instructions are changed, personnel are replaced, etc. Most everyday problems are being handled pretty well (?) already. Over the last 20 (?) years, management theory began talking about "management by objectives" (often misapplied). More recently, there has been the "team" concept (also not well-understood). These ideas are not well-developed at this time, and the utility of the feedback-control-system concept is not yet clearly perceived. But makeshift alternatives are being used. It appears that one of the reasons that control-system theory is not generally applied is that longer time-scale needs are being satisfied fairly well without explicit control-system analysis.

At the dose of his post, Bill refers again to "experimental test." I'm not at all sure that is the relevant consideration. Theoretical structures tend to be accepted or rejected not only on the basis of formal test, but also on the basis of convenience and applicability. After all, the main advantage of the heliocentric theory of the universe is the simplicity of the computations. What a mess when earth-centered! But, correctly done, the results are indistinguishable.

When a decision is needed (a choice between/among alternatives), the DME examines ("imagines") related experiences from memory. It considers conditions (remembered) that might limit the selection(s). Anticipated results (projected through imagination) are compared with the objectives for acceptability. The DME could combine selected procedures sequentially or use an average (weighted, perhaps) of the imagined procedures. These imagined procedures are used by the DME as structured inputs to the corresponding levels of the hierarchy. Under ordinary conditions, this might take only a small fraction of a second. But if the situation is complex (and time permits), extensive investigation and study might be used before finally selecting the procedure. The whole process is so familiar and quick that it is easily overlooked.

This very general summary becomes more meaningful when applied to real people in real situations. I saw a figure-skating contest (pairs)

last night. *Very* complex activities—mainly muscle skills, but commentators reported some of the personal interactions that can play a part. I was struck by the situation when a *disturbance* occurred, a fall to the ice. This is a very complex situation: the planned sequence, with its timing requirements, has been suddenly interrupted. This appears to require extensive reworking of the many systems involved. However, the response—compensating movements—was within a fraction of a second! Clearly, the skaters had available, almost instantly, an alternative procedure. It was designed both to avoid injury and to continue the program. Most of these skaters had 10 or more years of practice. If you have ever tried to ice skate, you know that much of early experience involves learning how to fall without bruises. Thus the experienced skaters have a large supply of alternative memories that can be quickly applied when needed. Notice, while this involves much "repetition," this is *not* "reinforcement," rather it is acquiring a repertoire of alternative variations of performance.

In terms of Orders of the hierarchy, such a contest certainly involves interpersonal relationships (my suggested Sixth Order) and, in various degrees, all lower orders. In performance, the selected relationships are played out. But in discussing the contest, communication skills are used. Here, words are used to represent perceived variables at several levels. The ice, the skates, and the arena are (more or less) objects that can be considered among the Second Modes of Sixth Order. The movements, with their timing, would be Third and Fourth Modes. The combination into skilled performance could be Fifth Order. Overall, there are the personal interactions of the skaters in a framework of competition. Here we have Sixth Mode of Sixth Order. This analysis can be carried further and applied to other activities.

If we want our ideas to be used, we must show where they help solve other people's problems, that is, help them achieve the goals they are already working on. What are their motives, their higher-order objectives? We should show where Perceptual Control Theory fits into and contributes to their ideas. We should compliment them on their knowledge and insights.

Behavior perceived as attack results in defense, retreat, or return of attack. Such conflicts might be fun for the winner, but they often result in losses on both sides. Examples abound.

Experimenters must begin with some kind of theory as a guide to experimental design in any field. Each has several alternatives: try to validate some theory (not necessarily his or her own), try to invalidate some theory (likewise), try to determine interesting parameters (perhaps a recipe for a candy), or just do something for "the heck of it."

Why are experimental studies undertaken? Perhaps a student needs a thesis topic. It will be subject to assorted approvals. To maximize probability of approval, it should fall within the range of currently acceptable ideas in that area.

Perceptual Control Theory is in competition with various other theories, some of which (in economics? in sociology?) are rudimentary indeed. But those other theorists tend to be interested in finding ways to support and defend their ideas. PCT people show the same behavior. What is needed is to change from “conflict” to “cooperation,” or at least neutrality. (There are several ways to resolve a conflict.) Some attention could be directed toward strategies of interaction, using communication skills to “make friends and influence people” (Dale Carnegie, of course). Theories and ideas gain acceptance by being useful, not by winning arguments. Regarding PCT vs. other theories, PCT can describe and analyze the behavior of their opponents, while the others can only describe PCT behavior by misunderstanding, overlooking, discarding, etc. various common observations. Bill has listed some of them. These remarks are very condensed, but I think my viewpoint is clear—at least to PCTers.

Bill Powers: Bob, in talking about your “modes,” you appear to take an external view of someone else’s organization. That is, you seem to be looking for levels that will apply to “psychological” aspects of a person, to explain the how and why of that person’s behavior. I’m taking a different viewpoint my definitions of levels are meant to describe how the world appears from the standpoint of the person, regardless of the context. When I speak of “system concepts,” I’m referring not just to things like a self or a personality or a character, but to *all* system concepts. To a physicist, for example, there exists something called physics, a discipline. This is, of course, a perception. The entity called physics, I have proposed, is a concept built from a set of principles and generalizations, which both provide the material within which the entity physics is perceived, and which, as goals, are specified by the goals we have for physics—that is, for what kind of entity we want it to be. The principles and generalizations, in turn, are built out of a set of rational, logical, reasoned mental processes that I call, generically, “programs.” In a set of programs, we can discern general principles; at the same time, the principles we wish to maintain in force determine what programs we will select to use.

My intention in proposing these levels of perception was to provide a framework within which we might understand all human experiences, no matter what they are about. If the subject matter is one person’s experience of other individuals, then what I call “system concepts” would correspond to what you term “personality,” and perhaps what I call “principles” would correspond to your “character,” and my “programs” to something like “habits” or “abstract skills” or “intelligence.”

These are ways of perceiving other people. But these general classes of perception and control include more than our experiences of other people. As I said, they include all experiences of all kinds. To a manager, the system concept called “my company” is as much an entity as “my children.” To a patriot, “my country” is a real living entity. To a sociologist, “society” is a system concept with as much reality as “self.” And to a chemist, chemistry is an entity with characteristics that depend on principles that are implemented as programs, without any organisms in the picture.

So what I am most interested in are the general classes of experience, not specific contexts in which we might give them more specialized names. The concepts of “character” and “personality” are inventions, but they are examples of fundamental classes of perception shared by the educated and the uneducated alike, and constant across cultures (I sincerely hope).

Bob Clark: Firing long-range weapons provides an example of the importance of the “speed of feedback.” When guns were first used on targets that were beyond visual range, results were poor. Soon “spotters” were introduced to report the results. Thus, the gun became more accurate. This combination can be regarded as a negative-feedback control system, even though the return signal is relatively slow compared to the speed of the projectile. It does not permit control of each shot, but provides improved control of the overall performance of the gun.

Without spotters, feedback was slow indeed; hours to days were needed to get reports. Adding spotters reduced the delay, providing much faster feedback. Self-guided weapons are now available: cruise missiles, smart bombs. These work better yet, with much faster corrections. With these capabilities, they correct for aiming errors, possible movement of the target, and varying winds.

Analysis is influenced by the time scale selected. When times of the order of seconds are of interest (approximately the time needed for the projectile to arrive), there is no control without self-guidance. Here, open-loop analysis applies. Events are followed around the loop without treating the system as a whole. When events are examined in terms of the time for firing the gun several times (several minutes), closed-loop analysis applies to each firing of the gun as the assigned target is followed. Assuming the necessary components are present, either closed- or open-loop analysis may be suitable, according to the time scale of interest.

A primary question for any control system is: “What is the perceptual variable being controlled?” In this case, it is the point of impact of the projectile. This variable is a combination of several perceptual variables used to specify location in terms that can be communicated to the gun crew. The observer's conscious attention is required in combining and communicating this information.

This system, assembled for the purpose of controlling the impact of a projectile, can be used as a general example of feedback systems. These observations may help in the analysis of other systems where the separate operations are unclear. Each of the parts of a control system can be identified: the feedback function is the spotter (plus communication equipment); the output function is the gun (plus the powder, projectile, aiming devices, and crew); the reference signal is the target (provided by higher command, the "Decision Making Entity"); and the comparator is the human (or a specialized device) determining the size and direction of the error provided to the crew to adjust the aim of the gun. For a time scale fast enough to observe these events as they occur, analysis can emphasize any one of the components. A mathematical equivalent of each of their separate operations can be written. For a time scale so slow that the system has come to equilibrium, analysis concerns the operation of the entire system as a unit. This is equivalent to solving the equations for the controlled variable in terms of the reference signal (and system parameters). The result is the familiar form used to describe the operation of a closed-loop feedback system.

The preceding discussion has been in rather mechanical, abstract terms. Regarding the people operating the system, each one is primarily concerned with his/her own part in the detailed sequence of events, rather than with the combined operations as a feedback system. Each person uses the skills needed for the immediate purposes. He or she selects and applies them as he or she understands their function in the larger organization. He or she also coordinates them with his or her individual internal conditions and needs.

The commanding officer, using a time scale suitable to his/her needs, regards each combination of gun, crew, and spotter as one of the parts of his/her output function. To him/her, each "rifle squad" is a simple straight-through system: he/she assigns the target, and the system performs. This can be considered as a stimulus-response system with its performance improved by adding a negative-feedback loop. This treatment, however, omits the events in between the "stimulus" and the "response." For some purposes, it is adequate.

The above is an example of a control system with two levels. Selecting suitable response times helps separate and identify the different levels. By adding another level of command, we have a three-level system. For the gun crew's spotter, the time scale would be of the order of minute, the time to fire a few shells. The commanding officer is concerned with the operation of his/her several guns. His/her time scale would be from minutes to hours, and, in turn, the higher commander works with larger-scale tactics/strategy and even longer time scales; to him or

her, the individual gun and crew with its assigned target is simply a tool to be used. He or she is concerned with larger-scale results.

Consider, in passing, what happens when the chain of command is bypassed and higher-order corrections are introduced too early!

Memory, expressed in several forms, is essential to the operation of this system. Some of the data are in the form of maps and instructions. Some are in the form of the aiming and firing mechanism of the gun. Some are in the form of remembered procedures and instructions. Some are in the form of remembered orders "from above." And so on.

In fact, the entire set of concepts, ideas, procedures, and skills are all located within the memories of the participating individuals. Each must have available, as a minimum, those portions of the operation that apply to him or her. Perhaps this could be simulated with high-speed computers and software, but the operating components must all be included in some form. Although the mechanical requirements are relatively modest, the memory capacity and programming to provide for automatic selection among many alternative actions is mind-boggling!

Each participant must direct his attention to the assigned task, while "simultaneously" "paying attention" to several other variables, especially those in his or her immediate environment. This requires frequent shifting of attention among several perceived variables.

Bill Powers: Bob: a very nice analysis, with lots of interesting observations. One thing your examples about "synthetic" control systems show is how crude control actually is when an organization tries to imitate individual control systems. But even an organization wouldn't think of computing how to aim the gun and firing it without looking to see where the shell landed.

One minor quibble. You say that "the reference signal is the target (provided by higher command, the 'Decision Making Entity"); and the comparator is the human (or a specialized device) determining the size and direction of the error provided to the crew to adjust the aim of the gun." When the commander says, "Put a warning shot just in front of them," the aiming point is not the target, but a point that bears a specified relationship to the target position. So it's the relationship between the impact point and the target that is the reference signal, and it exists only in someone's head prior to the shot.

An added observation: in order to adjust the gun position over repeated shots, the error must be turned into a new gun position. In order to get the final error as small as possible, you need a high loop gain. But if you have a high loop gain, an error of +50 yards would lead to a large correction, and the next error might be -500 yards. The solution is to use a slowing factor, such that only a constant fraction of the computed correction is actually applied on any one trial. In that way, you can have high loop gain and accuracy, without instability of control. The Same principle applies in spinal control loops with transport lags.

Dag Forssell: I was fully trained in the Swedish Army Artillery. A behaviorist might listen to Bob Clark and hear him say that this is a chain of cause-effect happenings. We in PCT notice the *multiple* iterations required to arrive at the target and can see the similarity with the iterative calculations of Rick Marken's spreadsheet model. We *can* see that the difference is quantitative, not qualitative, since we see the error signals at work, pulling in some (hopefully correct) direction, and we know that the process works well even without perfectly planned and executed output functions.

Bob Clark Bill, I am surprised by your reaction to one of my remarks. You say that I "seem to be looking for levels that will apply to 'psychological' aspects of a person, to explain the how and why of that person's behavior." This suggests that I *begin* by selecting "psychological" aspects, *then* search for lower-order systems (variables?) that might fit. To the contrary, I begin with the lower-order variables. Thus, I look for perceptual variables that use combinations of selected skills (including their related lower-order variables). With a rather large assortment of these perceptual variables, the question is one of assigning useful labels. Labels are needed to facilitate their selection and application, both for use as sources of sets of reference signals and for communication. Labels are preferred that will be generally understood and thus communicate to more people.

I am basing my analysis on your very important observation that *behavior is the control of perception,* and that perceivable variables are the heart of the structure. I might have overlooked some important aspects of the situation—I am sure you will point out where my suggestions can be improved.

Bill Powers: There's a subtle difference between "sequential" and "lagged" control. Bob Clark gave an example of truly sequential control: lob a shell, wait for the spotter to see where it lands, wait for the spotter to send the message back to the gun site, lob another shell, etc.

Lagged control is like aiming a fire hose. The water shoots through the air and lands somewhere. The fireman is watching where the water lands and corrects his/her aim according to the error between perceived and intended landing spots. There is water continuously flowing and continuously landing, and the fireman is continuously monitoring the landing spot. There is always water leaving the nozzle at the same time

that water is landing on the fire, at the same time that the fireman's eyes are seeing the water landing, at the same time that the fireman's muscles are altering the aim of the nozzle. The processes in various parts of the loop are all going on at the same time, literally simultaneously—even if it takes two or three seconds for any one drop of water to fly through the air and land on the fire, and a hundred nanoseconds for the image of the water landing on the fire to reach the fireman's eyes, and 50 to 200 milliseconds for the image to be converted into a perceptual signal, and an error signal, and a new muscle tension.

The second case is the most common in human behavior, although there are valid examples of the first (corresponding by e-mail, for example).

Many analysts of human behavior have confused sequential control with lagged control. They assume that while a stimulus is occurring, everything else in the control loop is on hold until the stimulus finishes its pattern. Then, with the stimulus input finished, the response commences, goes through its pattern, and stops. At that point the effect of the response alters the stimulus conditions, with neither stimulus nor response occurring. Finally, the next stimulus occurs and the sequence begins again.

Even inside the nervous system, this same erroneous image seems to be used. A neuron fires, sending an impulse along a fiber to its end, where the impulse triggers off the next impulse in line. The maximum number of input-output events per second therefore seems to depend on the time it takes for an impulse to travel through the nervous system to a muscle.

In reality, there can be anywhere up to 10 or so impulses traveling along the same nerve fiber at the same time (length of path, say 0.5 meters, divided by speed of travel, say 50 meters per second, times impulse frequency, 1000 per second or more). The maximum number of input or output pulses per second is set by the maximum impulse rate, regardless of transit time through the nervous system. If you count redundant paths carrying similar information, the maximum rates are even higher than that.

This confusion is the result of trying to describe a dosed-loop process in words. Using words, we can say only one thing at a time. We can't be talking about input processes while we're also talking about output processes and the processes in between, or the effects going on in the external part of the loop. So language forces us to describe first the input, then the comparison, then the output, then the effect on the environment, then the effect on the input again, as if this were a sequence of mutually exclusive events. If one lets words dictate thought, the mental image of the process will have the same sequential nature, leading to incorrect analyses and failed predictions.

Bob Clark: "Anticipation" has been used without being really tied to PCT very well. A common example: the time when you got on an elevator, pressed the "up" button, and it went down. This is quite upsetting the first time it happens, because your remembered experiences lead you to expect—"anticipate"—it to follow the button's label. There are other common experiences of many sorts. (Such as going up—or down—stairs and finding one step, more or less, than was expected.) The point here is that "anticipation" and its related concepts are common occurrences.

What does anticipation consist of? It begins with the existence of a situation where there is a goal to be achieved. The Decision Making Entity examines the memory for ways to reach that goal. There could be an established procedure—a set of related reference signals—that needs only to be put into operation. Absent such an established method, the DME "looks" for an alternative that appears to result in reaching that goal. It (the DME) selects a promising procedure and uses that remembered set of reference signals. The DME might be using a previously successful procedure, or it might be extrapolating from remembered events. Either way, future events are expected, that is, "anticipated."

Anticipation dearly plays a significant part in a decision to "go to Paris." And the DME finds in available memories (including maps, travel agents, etc.) the procedures needed. These procedures are then used to provide suitable reference levels as inputs for the systems needed. In this situation, various skills are needed: communications to assorted people, handling money and tickets, passports, etc., etc. Variables of configuration, sequence, and time must be included. And all of these involve suitable control of the lower-order muscle skills. Bill, I think this is consistent with your view, but you have stated it in such abstract terms that some of this might be overlooked. It is very helpful to have the concepts of temporal variables, skills, etc. available in addition to that of configurations.

It seems to me that the concept of "Intentionality" recognizes that people make decisions (action by something I call the Decision Making Entity), selecting future events/situations to be achieved. E.g., I got in the car with the "intent" of going to the dentist. I "anticipated" little or no traffic and expected the car to perform as it has in the past. I remember the route and the conventions regarding other cars. To me, "Intention" is a Sixth Order concept—one uses available skills to accomplish higher-order purposes. Is this a problem?

Bill, I think your "minor quibble" about gun-aiming is more serious. The new "aiming point" is the "new target" for the gun crew. The target for the crew is no more, no less than that ordered by the commander. To specify it in terms of the preceding target might be a convenient shorthand way to communicate the position of the new target.

It seems to me you are following events around the loop, resembling open-loop analysis. Using a time scale including several shots, appropriate to the view of the commander, high loop gain should improve the resulting accuracy. Examining the series of events, we begin with the first shot. It misses by some amount, and the location of the impact is reported by the spotter. If the spotter is very sensitive, this location could be reported in feet or inches, although yards might be sufficient. The aim is then adjusted by the crew to whatever accuracy the equipment permits. High gain means that the aim is corrected very precisely. However, the second shot could be off considerably if, for example, there is a gust of wind, the target moved, or whatever. But high loop gain would still tend to minimize the error, instead of creating an over-correction. An over-correction might occur if the gun controls were not properly calibrated. As I understand it, a bracketing procedure is often used to calibrate the gun controls.

Indeed, the "bracketing" concept is useful in any situation (exploration, experimentation) lacking accurate or reliable data.

My statement that "Analysis is influenced by the time scale selected," would have been more clear as "Whether open-loop or closed-loop analysis is appropriate depends on the time scale selected."

Closed-loop analysis is appropriate for a time scale in which the firing of the gun is completed before the higher-order system (the commander's system) can respond. The loop gain has little effect on this analysis, because the loop serves as part of the commander's output function. The gain of the loop determines the accuracy with which the output signal follows the reference signal. Loop gain is determined by combining the sensitivity of the spotter with the sensitivity of the gun aiming equipment.

In the open-loop analysis, the concept of "high loop gain" does not apply. There is no "loop" to have a "gain." It particularly does not apply to the gunner alone. The gunner adjusts the aiming equipment according to the correction called for by the spotter. If the report is "100 meters too far," the gunner makes the corresponding correction (perhaps aiming two degrees lower); the spotter reports again, etc.

Which view is more useful depends on the purpose of the analysis. The commander's view, with its longer time scale, uses closed-loop analysis; the spotter's view uses open-loop analysis.

Dag comments: "A behaviorist might listen to Bob Clark and hear him say that this is a chain of cause-effect happenings. We in PCT notice the *multiple* iterations required to arrive at the target and can see the similarity with the iterative calculations of Rick Marken's spreadsheet model." I am not familiar with the Marken spreadsheet, but I can infer the general nature of the demonstration. The iterations are, of course, steps in the correction process. When observed with a longer time scale, these iterations disappear; at a shorter scale, they become more obvious. Purely a matter of viewpoint and choice of time scale for observation.

Bill, your example of the fire hose for "lagged" control seems to work very well. But I don't think the fire chief cares which form of control it is as long as the water lands where *he or she* specified. The chief uses a time scale of perhaps minutes, vs. the seconds needed for the water to flow.

The existence of conflict depends not so much on the nature of the perceptual signals as it does on the relative time scales. Thus, the "gun crew plus spotter" is controlling the point of impact of the shell, and so is their commander in assigning the target. If the commander observes excessive spread in the pattern, he or she might make changes in the lower-order system. He or she might, for example, adjust the position of the spotter to improve his or her sensitivity. Both systems are concerned with the same perceptual signal, but their output systems operate differently.

As suggested, "conflict" occurs when the time scales overlap. If the spotter is repeatedly moved to a new position before the operations from the preceding position have been completed, a loss of accuracy (perhaps temporary) results. Some forms of stuttering provide another illustration. If the individual attempts to correct the formation of his phonemes too soon, i.e., before completing a word or phrase, stuttering is unavoidable. Many other examples are readily found.

Bill, in your words: "I'm taking a different viewpoint: my definitions of levels are meant to describe how the world appears from the standpoint of the person, regardless of the context. When I speak of 'system concepts,' I'm referring not just to things like a self or a personality or a character, but to *all* system concepts. To a physicist, for example, there exists something called physics, a discipline. This is, of course, a perception. The entity called physics, I have proposed, is a concept built from a set of principles and generalizations, which both provide the material within which the entity physics is perceived, and which, as goals, are specified by the goals we have for physics—that is, for what kind of entity we want it to be. The principles and generalizations, in turn, are built out of a set of rational, logical, reasoned mental processes that I call, generically, 'programs.' In a set of programs we can discern general principles; at the same time, the principles we wish to maintain in force determine what programs we will select to use."

"Programs we will select"—who, or what, does the selecting? The DME?

Your selection of these higher-level structures reflects your extensive knowledge, together with the application of a high degree of logical skill and reasoning. However, what about those who are not as knowledgeable? How do they manage? What are the categories, etc. that they form and live by? When they interact with other people, what are the concepts they use? How can we talk to them without some common language?

Bill, I am troubled by your move from your Fifth Order, control of sequence, to discussion of "concepts." Are these concepts derived from combinations of lower-order perceptual variables? If so, how? And which? Does the operation of these concepts include setting reference levels for Fifth Order and/or lower-order perceptions? How, and by what is this done?

In the Glossary of *Behavior: The Control of Perception*, I find: 'Perception: A perceptual signal (inside a system) that is a continuous analogue of a state of affairs outside the system." Finding no special definition of "concept" in that Glossary, I consult my dictionary: "concept, n. 1. a general notion or idea; conception. 2. an idea of something formed by mentally combining all its characteristics or particulars; a construct." I think that's essentially what you mean. What are the perceptual components of "concepts"? It seems to me that this term is too broad and vague a category to be assigned as an Order of control in the hierarchy.

Also, for "entity," as in "entity called physics," above. Not in the Glossary. Dictionary: "entity, n. 1. something that has a real existence; thing. 2. being or existence, esp. when considered as distinct, independent, or self-contained." This is how I use "entity" in "Decision Making Entity."

Your view of "physics" seems to differ from mine. To me, a physicist, it is not "a" concept, rather it is a specialized language, including its own special words, syntax, etc. It is an assemblage of definitions, observations, methods, procedures, formulas, derivations, etc., etc. I find these in various locations in my memory—given suitable situations, they are available to select for use, or whatever. In one way or another, any of the lower-order perceptual variables might be pertinent. But it does not seem to me to serve as a "concept."

Concerning "what kind of entity we want it [physics] to be," I don't have any particular "goals" for "physics." It is "set of tools," very useful for certain purposes, but irrelevant for others.

My proposal is to assign control of temporal variables to Fourth Order, placing sequence at Third. Sequences have temporal aspects which are perceivable and controllable. Combinations of sequences with temporal variables, also perceivable and controllable, form skills. These provide new sets of perceivable and controllable variables. Skills can be selected: "Shall we dance the waltz, or the tango?"

You say: "My intention in proposing these levels of perception was to provide a framework within which we might understand all human experiences, no matter what they are about. If the subject matter is one person's experience of other individuals, then what I call 'system concepts' would correspond to what you term 'personality,' and perhaps what I call 'principles' would correspond to your 'character,' and my 'programs' to something like 'habits' or 'abstract skills' or 'intelligence.'" The "correspondence" you suggest appears to be limited to a similarity in position in the sequence of levels in the hierarchy.

To me, "personality" refers to a group of perceptual variables with names that are convenient because they are commonly "understood" by ordinary people. They relate to short-term interactions and include such perceptual variables as "friendly," "helpful," "dominant," etc. What you call "system concepts" draws pretty much a blank, except among those with unusual information and experience. Logical, yes, but the connection with perceptual variables is not clear to me.

To me, "character" refers to another group of perceptual variables. These variables also have names that are "understood" by ordinary people. They relate to identifiable, therefore perceptual, underlying forms of behavior displayed in repeated interactions. Examples include such concepts as "honest," "reliable," "thorough," "careless,"—they are not necessarily favorable. What you call at this point "principles," in the sense you seem to intend, also draws pretty much a blank, except among those with special knowledge as above. Logical, again yes, but what is the nature of the "perceptual variables" from which they are derived, or *for* which they might provide reference signals?

Similar comments apply to your "programs." "Habits," "abstract skills," and "intelligence" I would treat quite differently. To me, these raise important questions not included *in* my present comments.

You emphasize: "These [referring to my proposed terminology] are ways of perceiving other people." Yes, but they are also ways of perceiving yourself. We agree that one cannot observe (perceive) one's own acts during the performance of those acts. However, this does not prevent their perception by examination of recent (perhaps very recent) memories of those same acts.

Bill Powers: Bob, I wasn't accusing you of beginning with psychological constructs and then filling in lower-level systems. My point is different.

Sometime between 1960, when we parted company, and 1973, when *Behavior: The Control of Perception* was published, a change in my thinking about the levels seems to have occurred. Or maybe, being on my own, my direction of thought became clearer. This all seems to be clearer now that you're describing your hierarchical concepts once again.

At any rate, the "pre" idea was much like yours, that we were attempting to characterize human beings by identifying levels of control with various aspects of human functioning. Somewhere in that 13 years, I realized that this was not the right problem.

As I now think about it, the problem in understanding human nature is not so much to understand human beings as to understand the world that human beings experience. In this world, I include not only the three-dimensional world around us, complete with living color, stereo sound, smellivision, and so forth, but also the "inner" world of imagination, memory, thought, reasoning, understanding—the whole world of inner commentary on sensory experience. In short, the world of experience includes everything experiencable, whether we think of it as being "inside" or "outside."

This world, to the best of my knowledge, originates in signals emitted into the nervous system by sensory receptors. That observation seems fundamental to me; to deny it would be to wreck the entire structure of physical theory, which I do not propose to do just yet. There is no way for the state of the world outside the nervous system to be registered in the brain without first appearing as a set of raw unanalyzed sensory signals. Nothing by way of information about the outside universe can get into the brain in any other way.

This means that the world we experience must consist of sensory signals and other signals derived from them. The "other signals derived from them" include the totality of what we can experience, from the taste of chocolate to Fermat's Last Theorem, as well as our experienced "interest" in that Theorem, if any, and any "thoughts" we might have about it. Nothing is exempt.

When I say "it's all perception," this is what I mean. We live inside a nervous system, and all we know is what goes on inside that nervous system. Even our idea of the existence of the nervous system exists as a set of neural signals, perceptions. The physical world outside us is a network of hypotheses existing in neural networks in the brain. Part of this neural hypothesis is a conjecture to the effect that there is an objective physical world outside our sensors. Sciences like physics and chemistry are very well worked out neural hypotheses. At bottom, they rest on sensory experience and all that the brain can make of such experiences. Our very attribution of physical theory to objective phenomena is itself a phenomenon in the brain.

This changes the problem. Now the problem is to classify all of experience, not just experiences of other people. We might perceive another person driving a screw into a piece of wood as showing a "skill" type of control, but this leaves unexplained the screwdriver, the screw, the piece of wood, and the relations among them. Those are also perceptions, and they are being controlled. The term "skill" refers mainly to something about the person's organization, but to explain how a skill like that is carried out, we have to explain the screw, screwdriver, wood, and relationship as well. The perceptual organization needed to represent these four things explains their existence for the actor; the actor's behavior is explained, in PCT, as control of these perceptions. Whether we characterize that control as constituting a "skill" is more or less beside the point. If we can explain the behavior in terms of controlling perceptions of wood, screw, and screwdriver individually, and in terms of adjusting those controlled perceptions to maintain control of a particular space-time relationship among them, we have explained "skill," too. But we have also explained how any person interacts with the world, whether the immediate world contains other living systems or not.

What I attempted to do with my definitions of levels was to represent the way the world seems to appear to us—meaning to myself as a representative human organism. This was very much an idiosyncratic first try, and it has undergone revisions as I have attempted to refine the descriptions. The process involved was quite unscientific, in that I didn't take any polls or do any objective experiments. I simply looked and listened and felt and tried to understand what was going on from the standpoint that I was an observer watching the outputs of neural data-processing functions. "What am I taking for granted?" I asked over and over. What is it that I'm doing or experiencing that is so familiar and so self-evident that I don't even recognize it as a perception? What part of my experiences am I setting aside as having some special status, or treating as the background of more important things, or brushing out of the way so I can look at something more interesting?

The "relationship" level was a latecomer to the hierarchy. I had spent a lot of time looking for relationships between one perception and others, and between action and perception, but it took years for me to realize that relationship *itself is* a perception. The same is true for all of the levels added or modified since 1960. I had spoken for years about the "principles of control," without realizing that principles can't exist unless we perceive them, and to perceive them we necessarily have to have principle-perceiving functions. Similarly for "physics." What is physics, that I can know it exists? It's a perception, of course. If I couldn't perceive such a thing, it wouldn't exist for me. So what sort of thing is it? I have proposed calling such things "system concepts," for lack of any better term. And what other sorts of experiences are of that same sort? There are many, once you realize that this *is* a sort of perception.

I think that the key to understanding how I think of the levels is to get into a mode of observation in which, as they say in Washington nowadays, "everything is on the table." No thought, no concept, no background perception can be let go because it "doesn't count." Everything noticeable counts. Everything noticeable is evidence about

what at least one brain is doing. If you accept the basic premise that the experienced world begins as a set of unanalyzed sensory signals, the only conclusion is that everything noticeable is activity in a brain, and hence it has to have a place made for it in a model of a brain.

I don't think that I've characterized the higher levels of perception very well. The most I hope to get across by the terms I use is the approach, the idea of calling into question everything we normally take for granted, all of the operations and perceptions that we use in thinking about and acting on something *else. I* don't think we'll arrive at a consensus on the levels until more people go through this very personal sort of exploration and report their findings.

On a different topic, Bob says: "The new 'aiming point' is the 'new target' for the gun crew. The target for the crew is no more, no less than that ordered by the commander." Yes, there are two levels of control involved here. Considering only the commander's level, the target always remains the same: the position where the shell is intended to land. The error is the amount by which the gun crew misses the target. The commander must alter the target position given to the gun crew *slowly,* however, to avoid treating dispersion in the pattern of shots as a systematic error.

Bob also says: "It seems to me you are following events around the loop, resembling open-loop analysis." This is indeed difficult to convey accurately. Loop gain is in fact the product of all amplification factors encountered in one trip around the closed loop, so calculating it seems like following events around the loop. To get high loop gain when there are transport delays in the loop, one must also use dynamic slowing of error corrections, a low-pass filter. With the filter in place, the behavior of the system at low frequencies is just as though no transport lag existed. So even though all real systems do entail such lags, they can be neglected! A difficult point to get across.

And: "Whether open-loop or closed-loop analysis is appropriate depends on the time scale selected." I don't think this is quite right. If one does an analysis on a short time scale where delays are visible, but neglects dynamic effects, a control system with a loop gain more than -1 will be incorrectly predicted to be unstable. The existence of large negative loop gains can be explained in a sequential analysis only if the proper low-pass filtering is taken into account—and it is usually not taken into account in open-loop analyses.

Consider a control system in which the controlled quantity is equal to the output of the system, the input function has a gain of 1, and the output function has a gain of 100. If there are lags in this system, as there are in all real systems, you would predict on that basis alone that the system would go into violent overshoots increasing without limit by a factor of 100 on every iteration. But now add a slowing factor that follows the rule: on each iteration, calculate the new output, and then let the actual output change by 1% of distance from the previous amount to the new calculated amount. This is a low-pass filter that does not alter the final steady state. The system will suddenly become stable; in fact, it will bring the error down to 1% in a single iteration! The effective long-term loop gain is still 100, so errors will be kept small over the long run.

If you try to eliminate the overshoots in this sequential system by just lowering the output gain to less than 1, the result will be stability, but the error remaining at equilibrium will be 50% of the value of the reference signal on the average. So you get stability, but almost no control. The high-gain system with the low-pass filter will counteract errors slightly more slowly, but will eliminate 99% of their effects. The low-gain system without filtering will counteract disturbances instantly, but will cancel only half of their long-term effect.

So there is a difference between closed-loop and open-loop analysis that is independent of the time-scale.

And you say: "I don't think the fire chief cares which form of control it is as long as the water lands where *he or she* specified. The chief uses a time scale of perhaps minutes, vs. the seconds needed for the water to flow." My point was that all components of a closed-loop system of this sort are operating literally simultaneously; they don't take turns acting, with no action between. This is how the nervous system works; sensors are generating signals at literally the same time that actuators are producing forces.

And finally: "As suggested, 'conflict' occurs when the time scales overlap. If the spotter is repeatedly moved to a new position before the operations from the preceding position have been completed, a loss of accuracy (perhaps temporary) results." With proper design, the system would work better if the spotter were moved immediately, rather than waiting for the previous results to come in. This would be the right strategy if the calculations were being continuously averaged over several shots, as would be necessary to distinguish random from systematic errors.

On another topic, you say: "What about those who are not as knowledgeable? How do they manage? What are the categories, etc. that they form and live by?" I see your point and agree that it has to be considered. My levels are intended to describe categories of experience that all people (and even animals) employ without any training or knowledge. All people perceive and control relationships, by my account. They also perceive and control categories, sequences, logical functions, etc., not by thinking about it but simply by having the world presented to them in such terms by the *basic* equipment of their own brains. I don't know how to put it better than that.

You also say: "Bill, I am troubled by your move from your Fifth Order, control of sequence, to discussion of 'concepts.' Are these concepts derived from combinations of lower-order perceptual variables? If so, how? And which?" The levels as of now (February 1993) are (1) intensity, (2) sensation, (3) configuration, (4) transition, (5) event, (6) relationship, (7) category, (8) sequence, (9) program, (10) principle, and (11) system concept. Each one, when analyzed into components that are not just smaller groups of the same level, proves to be a function of perceptions of the next lower level (or lower still). So a system concept like physics is drawn from perceptions of many physical principles, while principles are drawn from perceptions of many specific logical/mathematical operations, and so on down the list.

As to *how* a perceptual function of one level combines lower-level perceptions, I have no idea. The nature of the functions must be very complex at the higher levels, or at least of a kind that we can't analyze now. The apparent dependencies were arrived at from analysis of experience, much as we can see that configurations are composed of sets of sensations. Also it was helpful to ask how we would go about maintaining a perception of any given level against disturbances—how, for example, we would maintain the principle of honesty. To perceive ourselves as honest, we set reference signals for certain programs of action and thought which we call reasoning or analysis or procedures. None of this is very firm; I'm just reporting how it seems to me after as close an inspection as I can carry out. Other people's opinions are obviously needed.

I chose the term "system concept" with the emphasis on "system," not "concept." In my view, "concept" falls within the range of meaning of "perception," because it's something we can experience as occurring or existing. I could have said "system perception." It just means the sense of an organized entity of some sort being present, the kind that is composed of principles, generalizations, heuristics, characteristics, whatever you want to call them. Perceiving a specific person whom you know well leads to this sort of system concept or perception—the impression of a particular person, a personality, a system. Shoot, how am I suppose to be more specific about an idea that's not very clear to begin with?

You go on to say: "To me, a physicist, it [physics] is not 'a' concept, rather it is a specialized language, including its own special words, syntax, etc. It is an assemblage of definitions, observations, methods, procedures, formulas, derivations, etc., etc." Yes, that's what I mean by a system concept. The very fact that you can, without enumerating, refer to all of its components as some sort of bringing-together into an "assemblage" of a variety of more specific elements shows that you have formed a conception of physics as a unified system of ideas, definitions, observations, methods, procedures, etc., with the "etc." indicating that the picture includes much that is not enumerated. "Physics" is clearly a system concept quite different from "religion" or "family." Enumerating the lower-level details of these other system concepts would entail quite a different list.

When you say, "I am a physicist," the "I" being indicated is associated with the system concept of physics. For the moment, the center of awareness is operating from that position. But when you say "I am a father," the system concept is the one we refer to as "family," and the "I" now takes on new characteristics associated with a different system concept.

Or at least that makes a good story.

As to other differences, let's just go along with them for now. I'm feeling a bit overloaded.

Bob Clark: Bill, I was preparing for a final edit of another post on levels when I received your latest post. I am pleased, but not surprised, to find our primary views of "the world" have remained identical over the years: "This world, to the best of my knowledge, originates in signals emitted into the nervous system by sensory receptors." And: "This means that the world *we* experience must consist of sensory signals and other signals derived from them. The 'other signals derived from them' include the totality of what we can experience, from the taste of chocolate to Fermat's Last Theorem, as well as our experienced 'interest' in that Theorem, if any, and any 'thoughts' we might have about it. Nothing is exempt." Also: "When I say 'it's all perception,' this is what I mean. We live inside a nervous system, and all we know is what goes on inside that nervous system."

Given this viewpoint, with which I completely agree, there are several pertinent problems.

You report that your "pre" idea was "attempting to characterize human beings by identifying levels of control with various aspects of human functioning." That does not quite fit my recollection, but we probably need not resolve the matter at this time.

My present views have developed irregularly over the years. They have been modified since you and I were in contact in 1987, and further developed since I met Greg Williams in 1988. Some of the ideas I have been presenting recently are still being revised. I certainly expect further changes as discussions proceed—just as I think you also expect.

In your most recent post, you have restated your current view: "Now the problem is to classify all of experience, not just experiences of other people. We might perceive another person driving a screw into a piece of wood as showing a 'skill' type of control, but this leaves unexplained the screwdriver, the screw, the piece of wood, and the relations

among them." (An aside: in my view, "skill" is not a "type of control," rather it is a combination of perceptual variables that includes perception of objects (screwdriver, etc.), the one using the tool, the location of the several objects, and the sequence of events and interactions required in order to "drive a screw into a piece of wood." This "combination of perceptual variables" includes several less-complex skills, such as reaching for the screw, placing it in the required position, etc. This entire combination could be referred to as "driving a screw, etc.," which is one among many muscle skills that can be used to accomplish higher-order purposes. Thus, "skill" is a category of perceptual variables, selected for purposes related to interactions with other people and distinguished from lower-order variables by combining them (sequences of muscle tensions combined with temporal variables) to form the specific skill selected. Perhaps that is not an "explanation," but I think it is "understandable," and I hope that it communicates something of my view of Fifth Order.) You use two familiar, frequently used words: "understanding" and "explanation." Exactly what does each "really" mean? I find my dictionary of little help here—let me try to define them: "Explanation" seems to consist, at a minimum, of being classified, that is, placed in a category. That category might or might not pre-exist, but to be useful, it probably should contain more than one element.

Is a dog "explained" by having its breed specified? Or by naming its species? Or by its genealogy? How about its physiology, or neural systems? Of course not. Neither is "control of a perceptual variable" "explained" by pointing out that its actions resemble those of a negative-feedback system.

Instead of "explaining" some thing, activity, system, or whatnot, I prefer "description" of parts and their connections with each other and with other items. "Interactions" among the parts and with other items describe its "behavior." I am pretty sure that this is what you mean.

"Understanding" is the goal of every teacher for his or her students. For me, too. However, it seems to me that there are two aspects to this concept: internal and external. The "internal" aspect is displayed by simply asking, "Do you understand this matter?" If "Yes" is the reply, this signifies that there is no perceived recognition of inconsistency within another's internal array of information (perhaps after modification to include the new material). The "external" aspect is more complicated, being displayed by asking the other party to "solve" a problem that requires "proper" use of the material to obtain "the" solution. If the result is "acceptable," it indicates (does not "prove") that the comparable parts of each party's systems are in agreement. This is desirable, of course, because further discussion is facilitated, possibly leading to revision (perhaps by both participants).

It is interesting that I have had the experience of saying "Yes" to the question, but finding that the external test reveals some degree of "misunderstanding." Indeed, I think that most people have had this experience in one form or another. "Consistency" is demonstrated by this procedure, but not necessarily consistency with other parts of either party's systems.

Since our views of the lower levels are rather similar (with the possible exception of my Fourth Order, temporal variables), we move to higher levels.

Here I seek controllable, perceivable variables that are formed by combining lower-order variables. It occurred to me that muscle skills can be regarded as sequences combined with temporal variables. There are many such perceivable combinations. Some are relatively "simple," like walking, pressing fingers on buttons, pulling rubber bands, etc. And some are very complex skills, like vocalizing, running, throwing, dancing, acrobatics. Thus muscle skills, a group of perceivable, controllable combinations, can be assigned to Fifth Order, "skills." Such muscle skills are readily perceived not only in others, but also in oneself. Many are learned, some probably have genetic origins. In the process of learning how and when to use them, variations of many sorts are explored. Such experiments and their results are recorded (as "memories") as they occur. Thus, they remain generally available for later use.

What comes next? What would be the nature of Sixth Order activities composed of controllable, perceivable variables based on combinations of lower-level variables, especially skills of Fifth Order? As I was seeking to distinguish Fifth Order from Fourth Order, there was a tendency to consider interactions between/among individuals. Thus, with Fifth Order assigned to skills, Sixth Order could include all activities using combinations of skills for purposes requiring control of interpersonal interactions. Examples include games, competition, cooperation, government, dubs, businesses, and entertainment. In addition, language, mathematics, philosophy, systems, principles, and programs are included here. Here we find all theories, whether of the natural world, the world of imagination, the world of behavior, etc., including Perceptual Control Theory.

People generally have some sort of structured views of the nature of their surroundings and how to achieve their objectives. Their methods might be based on gross misunderstanding, superstition, or whatever, but they are sufficient for most people most of the time.

Communication, complex combinations of many muscle skills, taking many forms, is used throughout interpersonal interactions for many purposes. Should this be considered another level? In examining that possibility, it occurred to me to pay attention to everyday conversations among my friends and associates. Much conversation pertains

to Zero Order systems—health, sensations of temperature, and physiological events. There was discussion of combinations of sensations perceived as "objects." In turn, sequences forming postures, movements, etc. were of interest. These various combinations were used for ordinary, customary purposes of communication.

As "topics of communication," these might be called "Modes" of Sixth Order, corresponding to Orders of control, without themselves being control systems. Topics relating to skills would be Fifth Mode of Sixth Order. Those relating to communication and other interpersonal variables would be Sixth Mode of Sixth Order. The Modes do not function as control systems, but they assist in analyzing the structure and performance of the systems.

Continuing these observations, one finds comments about personalities and characters of individuals. What does this mean in terms of perceivable variables? The dictionary answers these questions rather well: 'Personality: 1. The visible aspect of one's character, as it impresses others: 'He has a pleasing personality.' This looks as though it could belong to Sixth Mode of Sixth Order, but it seems to me to go a bit further. Thus, we have people who are actors, behaving to portray varying personalities, emotions, etc. They appear to be controlling their behavior to produce certain interpretations by those around them. Being "pleasing," "friendly," "courteous," "hateful," whatever, can be controlled, even if contrary to the performer's own internal feelings. Thus "personality variables" can be regarded as controllable, perceivable variables in the performer's own repertoire. Interestingly, because combinations of skills are needed to display these variables, the time scale needed to perceive these variables is moderately long vs. the time needed for demonstrating lower Modes. "Character: 3. moral or ethical quality, 4. qualities of honesty, courage, or the like; integrity." Other definitions seem too inclusive or specialized. I think this does pretty well. Here, there is another increase in the time scale. While personality can sometimes be demonstrated in a matter of minutes, character requires observation of several incidents distributed over a much longer period.

These topics, "personality" and "character," are sufficiently different from each other and the other Modes of Sixth Order that they could be treated as Seventh and Eighth Modes of Sixth Order. Their importance in forming "images" of other people also suggests assigning them to Seventh and Eighth Modes of Sixth Order. This assignment would imply the existence of Seventh and Eighth Order Control Systems, based on corresponding perceptual variables.

This discussion suggests that something like "self image" could be considered Ninth Mode of Sixth Order, with corresponding Ninth Order control system. This treats personality and character as important components of self image, in addition to all other perceptions of whatever composes one's "self."

Where and how the DME, "Decision Making Entity," would relate to this structure is postponed for the present.

I am not very confident that the above distinctions among personality, character, and self image are appropriate, but they might be useful for discussion.

Conceived, I think, as a truly general theory of behavior, PCT should apply not only to observations of the behavior of other people, but also to ourselves, both individually and in the process of constructing a theory of behavior. "Personality" and "character" certainly can be used for describing other people.

On examining my memories of my own behavior, I find that I can generally perceive even these high-order variables in my own remembered behavior. Perhaps more important, I find that, if I care to, I can generally change my behavior. This might take more time than I like, but my perceived and changed behavior has become more nearly what I sought.

Further revisions are certainly needed. Perhaps most important, PCT should be applied to problems of general interest.

Martin Taylor: On reading Bob Clark's set of levels and comparing it with that of Bill Powers, I am for the umpteenth time reminded of the great difference between the internal view and the analyst's view of a hierarchy. Maybe I am being unfair, but Bob's sounds to me like the view one would see from the outside, rather than a description or model of what goes on inside an organism, whereas Bill's seems addressed to the mechanism inside the organism (again seen by an outside analyst).

Bill's levels deal with different kinds of perceptual input functions (PIFs). They speak, from the analyst's viewpoint, about what the organism *might* be controlling, and they have been developed by an organism that has attempted to consciously perceive what is normally unconsciously controlled. It is an empathetic view. Each level exists because there is a requirement for a different kind of perception, and the differences among the levels are (if I understand correctly) only in the perceptual input functions characteristic of the different levels (I can imagine that the output functions also differ, but I don't remember that being talked about).

Bob's levels strike me as speaking to what a social contact might perceive of a person; no single elementary control system (ECS) would act at a "skill" level, unless I greatly misunderstand what is meant. An external observer can see skill, and the performer, *looking from another viewpoint,* can assess his or her own skill, but no skill-level control

system can be extracted from a hierarchy. Maybe Bob can describe a skill-level elementary control system and prove me wrong. But I can't at the moment imagine "skill" as a level of control in the way that I can imagine "sequence" or "program."

Bob, you say: "As I was seeking to distinguish Fifth Order from Fourth Order, there was a tendency to consider interactions between/among individuals. Thus, with Fifth Order assigned to skills, Sixth Order could include all activities using combinations of skills for purposes requiring control of interpersonal interactions. Examples include games, competition, cooperation, government, clubs, businesses, and entertainment. In addition, language, mathematics, philosophy, systems, principles, and programs are included here." All of this is external, isn't it? You are talking about the applications for which Sixth Order systems would be used, not what Sixth Order systems do, or how they are constructed. Perhaps what you are saying is that Sixth Order ECSs individually contain language models, games models, cooperation models, etc. that they use in forming their perceptual functions. Such models are, indeed, possible. Symbolic artificial intelligence depends on them. But do they belong as intrinsic components of individual ECSs?

I think I have become more sensitive recently to the importance of separating the external (analyst or observer) viewpoint from the internal viewpoint. Many of the issues raised in recent postings seem to hinge on a failure to note, and sometimes on a tendency to mix, the two viewpoints. The organism can control what it can perceive, and it cannot perceive its feedback paths, other people's perceptions or references, or its own outputs. But the analyst can perceive feedback paths and the outputs of other organisms and can develop implausible theories that *require* the organism to perceive them. S-R theory cannot work if it requires the organism to control R, for example. The analyst can see that under relatively undisturbed conditions, there is a moderately consistent relationship in an experiment between S and R, as the analyst perceives them, and makes the unjustified claim that the subject produces R as a result of perceiving some transform of S. But the fact that the analyst can perceive both doesn't mean the subject can.

Many posters to the net, myself included, fall into the trap of writing about something the analyst can see as if it were something the analyzed organism can see, and asserting or assuming that the analyzed organism uses that property in some way. I don't know how to avoid this problem it is built into our language. Seeing that the problem exists is one way to avoid being caught by it. Sometimes.

Rick Marken: Excellent post, Martin! I think that the difference between Clark's and Powers' levels might be based on more than the internal/external distinction, but your discussion of that distinction was brilliant. I agree with you that it is probably the essence of the difference between the PCT and the conventional perspective on behavior.

Bill Powers: Bob says: "You use two familiar, frequently used words: 'understanding' and 'explanation.' Exactly what does each 'really' mean? I find my dictionary of little help here—let me try to define them: 'Explanation' seems to consist, at a minimum, of being classified, that is, placed in a category." The problem with this sort of definition is that all you get is a claim that the thing to be explained is like (or at least classified with) something else, which generally is also unexplained. Then Bob says: "Instead of 'explaining' some thing, activity, system, or whatnot, I prefer 'description' of parts and their connections with each other and with other items." I like this better. To explain a phenomenon is to describe its operation at a lower level. So models are explanations of the phenomena that they reproduce. Then: "The 'external' aspect [of "understanding") is more complicated, being displayed by asking the other party to 'solve' a problem that requires 'proper' use of the material to obtain 'the' solution. If the result is 'acceptable,' it indicates (does not 'prove') that the comparable parts of each party's systems are in agreement." Yes, the question when someone says, "I understand what you mean" is just what the other person's understanding is. This is the basic problem of communication.

Bob, you say that "'skill' is not a 'type of control,' rather it is a combination of perceptual variables..." This might be a difference between our approaches that I hadn't recognized. My levels are supposed to be types of controlled perceptual variables and, by implication, the systems that control them. When I label one level "programs," I don't mean just a level where programs are executed. I mean a level where we perceive *what program is being carried out,* and continually correct errors if we perceive a deviation from the correct program. An example would be watching people play cards. After a while, watching the play proceed, you recognize the rules in effect, and say, "Ah, they're playing five-card stud." Then, if someone violates a rule of five-card stud, you can perceive the error and (unwisely perhaps) point it out to the players to get them to conform to the rules. A rule is a form of program. To say "combination of perceptual variables" doesn't tell us much unless you say what kind of combination you're talking about.

You say: "Thus, 'skill' is a category of perceptual variables..." I agree with that: it is a perception at the level of categories in my definitions of levels. The category level is where we use one perception (here the noise or series of marks, "skill") to refer to a collection of perceptions of lower order.

Bob Clark: Bill and Martin, instead of making specific comments on your recent posts, I am offering comments of a more general nature.

Martin, you have focused on a general concept: "viewpoint." In view of your remarks, I am trying to summarize my (present) orientation in the following. This turns out to be much more difficult than I expected—and probably will change with additional review.

My general view. Quoting Bill: "We live inside a nervous system, and all we know is what goes on inside that nervous system." As I noted in an earlier post, that is also my viewpoint.

Categories. When I investigate what I have available ("inside that [my] nervous system"), I find several easily identified categories. Many other categories can be used as desired. I find the following categories particularly convenient and useful:

1. "Decision Making Entity" (DME; "Center of Awareness"). This is the entity that "uses" viewpoints. "I" is not used because it tends to include too much. This entity can direct its attention to any of the neural signals entering the central nervous system. It can shift its attention rapidly from one signal (or group of signals) to another. It also can select which of the available signals has its attention at any given time. It responds to "built-in" reference levels by selectively "paying more attention" to some signals than to others.

2. "Recording Function"; "Memory"; "Conscious". This is the entity that forms records of signals to which attention is directed. Attention can shift fast enough that it appears that all signals are recorded. Mere "exposure" to perceivable events seems to be insufficient for remembering. Conscious attention, i.e., perception, appears necessary. Teachers, parents, supervisors, etc. are invariably concerned that their students "pay attention."

3. "Perceptual Signals"; "Attention". These are the signals to which the DME's attention may be directed. From time to time, the DME selects them from the available signals. These form two groups: a. "Sensory signals" reporting the current condition of all physiological systems with neural connections to the central nervous system. They can form various combinations, resulting in production of additional, derived, sensory signals. b. "Imaginary signals" are recorded sensory signals and other recorded signals as selected by the DME. The imaginary signals include all perceptual signals derived from recordings. Generally, they are organized in some manner by the DME for convenience and accessibility. Such organization will distinguish between those coming from "external sources" and those coming from "internal sources." When selected by the DME for examination, they resemble audio-visual-sensory recordings. They normally run from past time events toward the present, and the DME can extrapolate them to future time. Likewise, memories can be combined in various ways, both

sequentially and simultaneously. In this respect, they resemble editing of videotapes.

4. “Output Signals.” These signals are recorded in the memory together with the corresponding perceptual signals. After review, the DME determines the “desired” effects on the perceptual signals. The DME then applies the remembered perceptual signals to the corresponding output systems. They act as “reference signals” for the systems connected to them. Effects are determined by the nature of the systems to which they are connected. The DME cannot directly perceive these signals (they are not “incoming”), but their effects are determined by observing corresponding perceptual signals.

5. “Comparator Function.” The DME makes its selections on the basis of comparison of the “desired” effects with the anticipated results offered by alternative sets of imaginary signals in relation to current sensory signals (and their combinations).

Viewpoints regarding the hierarchy. Martin, I have already been thinking about pointing out alternative views of the basic feedback-control system. However, you have focused on a more general concept: “viewpoint.” When I apply that concept to a minimal system, I find five identifiable viewpoints. Perhaps others can be found. Different viewpoints might call for different classifications and definitions of the hierarchical levels/Orders.

1. The “User’s” view. The User’s DME selects the desired condition (activity, etc.) of his/her own system, as it relates to its surroundings and applies the indicated reference signals. The User observes the resulting activity, etc. for possible deviation from intended performance. If deviations are observed, corrections are applied as indicated. The corrections are selected from memory, including anticipation, analysis, and theory (as the User understands them). This process continues as long as results are acceptable. If the results are not within limits, changes might be needed in the remembered structures. Although the concept of a hierarchy is not essential for the usual User, it can be very helpful when there is difficulty in finding adequate results.

2. The “Engineer’s” view. This view is “objective,” in that the Engineer treats the subject as external to himself or herself, omitting the part(s) he or she plays in this activity. The Engineer studies the details of the various elements of the system(s) and their interconnections. Each element is evaluated in terms of the relation(s) between its input(s) and its output(s), expresses them in logical/mathematical terms, and analyzes the results. If this is unacceptable, modifications of one or more elements and/or interconnections are examined for possible alternatives. The Engineer supplies standards of performance selected from his/her memory by his/her DME. In this process, the Engineer’s DME controls the activity. Although the concept of a hierarchy is not essential for the

usual Engineer (many are quite successful without it), it can be very helpful in more complex and multi-dimensional situations.

3. The “Outsider’s” view. The Outsider, that is, his/her DME, is observing the activities of another “living-behaving” entity. His/her information about that entity is derived exclusively from his/her own input systems—sensory, as modified and interpreted by his own established internal systems. He/she uses his/her knowledge to construct a description of the internal structure of the other entity. All of this activity, together with the conclusions, is stored in his/her memory and continues to be available for future application, modification, etc. These activities might include discussions, etc. with other Outsiders. Although the concept of a hierarchy is not essential for the usual Outsider/Observer, it can be very helpful in analysis and interpretation of results.

4. The “Experimenter’s” view. This view is also “objective,” in that the Experimenter treats the subject as external to himself or herself. He or she assumes that the subject’s reference levels are determined by the Experimenter’s instructions combined with the subject’s pre-existing decisions. The Experimenter selects and applies some action to the subject’s externally accessible inputs. The results are interpreted in terms of whatever behavioral theory he/she wants to apply, Although the concept of a hierarchy is not essential for some experimental purposes, it can be very helpful both in experimental design and interpretation.

5. The ‘Theorist’s” view. The Theorist pays attention to all of the views listed above, as well as any others that can be proposed. He/she resembles the Experimenter in searching for confirmation or denial of proposed theoretical and/or analytical ideas. The User’s and Outsider’s views provide additional data for evaluation of proposals. The Engineer’s view provides guidelines as to the logical and technical limitations that are intrinsic to the external surroundings. Although the concept of a hierarchy is not essential for some theoretical purposes, it offers the most inclusive and effective theoretical framework I know of.

Two views of hierarchical levels/Orders. These are both Theorist’s views:

1. Bill Powers’ view: “My levels are intended to describe categories of experience that all people (and even animals) employ without any training or knowledge.” Bill is concerned with “categories of experience.”

2. In my own approach, I have focused on the perceptual signals as they combine to form the hierarchy. “Hierarchy” is defined in *Behavior: The Control of Perception,* page 78: “This model consists of a hierarchical structure of feedback control organizations in which higher-order systems perceive and control an environment composed of lower-order systems; only first-order systems interact directly with the external world. The entire hierarchy is organized around a single concept: con-

trot by means of adjusting reference-signals for lower-order systems." I am concerned with categories of perceptual signals as they combine to form a hierarchy of perceptual signals.

Martin Taylor: Bob's "Engineer's" viewpoint is fine, but in using it, the engineer must try to empathize with the many viewpoints that occur at all places within the system. If point A is a perceptual signal that has as part of its input a sensory signal B, the engineer cannot assume that every variation in B is reflected exactly in A. The question must be "what does A see of the variation in B" before the engineer can properly assess what will happen at A. None of Bob's viewpoints seem to me to be of the class that I might call "internal."

From the outside view, there is a complex in the world that seems to be what the "subject" is controlling. It is the experimenter's view of the putative controlled environmental variable (CEV). The theorist outsider can also "see" the subject's perceptual signal that is the actual controlled variable. As far as the subject is concerned, that signal *is* the CEV. It is all that the ECS in question can know about the state of the world.

There are various kinds of "outsiders," as Bob has pointed out. One of them is the DME, which views all sorts of signals in the hierarchy. All outsiders use their own perceptions, rather than the one actually being controlled by the observed ECS. It is from the outsider's viewpoint that we can see a dichotomy between the CEV in the world and the perceptual signal. The subject cannot see it.

The outsider, who might be using very precise measuring instruments, can see that there are discrepancies between the state of the putative CEV and the state of the derived perceptual signal, even if the total perceptual input function is correctly interpreted. These discrepancies have to do with the resolution of the perceptual system. The subject might not be able to detect that any individual discrepancy exists, but he or she might be able to detect the possibility that discrepancy exists, by virtue of the success of control. (This is much the same in principle as the way astronomers judge the numbers of meteor craters on the moon that are smaller than they can see, or the way ecologists judge the number of species not yet identified.)

The perceptual signal, in this way of looking at things, does not define the CEV. It defines the operations on the sensed world that create the CEV, but the CEV is a structure in the world, not in the mind. It is a conceptual structure that mirrors the mind, and it might not be detectable to anyone else than the mind that created it, but, nevertheless, it is in the world, not in the mind. For example, a CEV might be "the distance between my fingertip and my nose." Forgetting the irregularities of skin and the like, there is a perceived value for that CEV—the

perceptual signal that corresponds to it. If I hold up my finger, I might perceive that distance as stable (or nearly so, with a slow drift), but I know from other information that if I could only see it, there is a rapid oscillation in the distance. Someone with a laser interferometer could probably measure fluctuations that are not in my perceptual signal. But I would say that they are in the CEV that the perceptual input function determines. So, the CEV is not defined by the perceptual signal; it is represented by the perceptual signal. It is defined by the perceptual input function.

There's a hidden issue here, one that relates to reorganization. There is no CEV that corresponds to the function that causes the actions of the subject to control an intrinsic variable. Reorganization controls the control operations, but it does not work on any perceptual signal in the usual sense: a perceptual signal based on a function of sensory input variables. Reorganization works, but it works only because the behavior of the world (unperceived) is factually stable over periods longer than the time it takes to reorganize. That factual stability can be inferred from the success of the reorganization. It cannot be perceived (I'm tempted to say "in principle," but I don't know if I could argue that). An outsider with a perceptual function that operated over a long time scale (I include memory here) could perceive the stability that permits reorganization to happen. Likewise, with a normal perceptual signal and its corresponding CEV, an outsider could perceive discrepancies between the CEV and the perceptual signal that represents it, even though the user of the perceptual signal cannot. But, as with reorganization, the user of the perceptual signal might possibly infer that there is a factual discrepancy.

I realize that the word "factual" in the above paragraph raises its own issues about boss reality and the like. I assume that all such issues are resolved against the solipsist position.

Bob Clark: Perhaps the following will clarify my earlier post.

The Engineer's goal seems to be the construction, at least in principle, of an assembly of hardware (or equivalent computer-cum-software) that performs the same way that a human (or, perhaps, a simpler organism) does.

Some Engineers approach this in terms of levers, gears, pulleys, etc. arranged so that inputs ("disturbances"?) at certain locations result in movements at other locations. By adding suitable "leading" terms (time derivatives) and "lagging" terms (time integrals), these systems can be made quite effective for specified applications.

The PCT Engineer, if that is a suitable term, bases his/her design on the properties of negative-feedback control systems. These are combined into a hierarchical structure, HPCT, assembled and modified to operate according to his/her desires. The Engineer proceeds by selecting from his/her inventory of memories, including physical and other principles, in order to bring his/her proposed structure into correspondence with his/her view of human behavior.

The design might include "recording and playback" capability, as well as ability to "reorganize" itself. In principle, these are both included in HPCT.

The following remarks bring up another subject, one that can lead, I think, to some very interesting and helpful results.

Over the last 10 days, I have tried to write this material from several different viewpoints. Each is pertinent and interesting, but it tends to become too long and complicated for a reasonably short post. This viewpoint appears to offer a framework that can be used to explore additional important (useful) subjects.

The Decision Making Entity (DME) can be considered from several viewpoints. Each is interesting, but the Theorist's is the most general, and it might be the most useful. This viewpoint is defined here by paraphrasing and quoting from *Behavior: The Control of Perception* (page 18).

The HPC Theorist proposes to construct a "model of the brain's internal organization" where "observed behavior is deduced... from the way in which these internal entities interact with each other and the external world." These entities have been chosen not only to "behave properly," but also to fit anatomical hints about the nervous system, physical models of the organism and its environment, subjective experience, and elementary mathematical logic.

1. Primary concepts: greatly condensed summaries of *Behavior: The Control of Perception.*

A. "Behavior is the control of perception"; "perceptual variables."

B. The negative-feedback control system and its intrinsic properties.

C. The hierarchical structure of negative-feedback control systems.

D. Problem-solving programs: fixed instructions with choice-points.

E. Intrinsic variables (genetically determined).

F. Reorganization: change in the properties or number of components.

G. Memory: recording and playback switches.

These concepts, with their analysis and development, cover a remarkably large range of human (and other) activities. However, this structure is largely fixed in form, changing only by the addition of new recordings or reorganization. Problem-solving programs, including associated choice points, are composed of recordings. They are derived from combinations of recordings and/or reorganizations. New programs result only from new/rearranged recordings and reorganization. This results in limited flexibility, leading to several problems.

2. Possible problems.

A. Minor changes in behavior might be needed because of inadequate or "incorrect" problem-solving programs. Reorganization is unnecessary and not initiated.

B. Minor changes in behavior might be needed because of inadequate or "incorrect" recordings. Reorganization is unnecessary and not initiated.

C. An operator is needed to control the recording-playback switches.

D. A source of reference levels is needed at the top of the hierarchy.

E. Arbitrary action is observed in the absence of intrinsic error.

F. Initiative is observed but not explained in present PCT.

G. Anticipation of unexpected events is observed but not explained in present PCT.

H. Errors, accidents, and misdeeds are observed but assignment of responsibility is not provided in present PCT.

I. Subjective reports ("User's view") of the processes of selecting among alternatives are not described in present PCT.

J. An "Observer's view" of subjects' unexpected actions is not described in present PCT.

A Decision Making Entity (DME) is proposed as a partial solution to these problems. The concept seems to be generally taken for granted and accepted by many people—including most (if not all) of those on CSGnet. Such acceptance is demonstrated by the frequent use of the first person singular. "DME" is proposed as a name for this concept when personal associations are removed, leaving nothing but the process of selecting from among alternatives for action. It offers a straightforward way to solve some of the above problems, and possibly others, by the addition of a single element with its associated capabilities and characteristics. This concept is consistent with several others discussed in *Behavior: The Control of Perception,* and it helps clarify the operations and relations within HPCT, as summarized above.

3. Operation of the DME: summary.

A. Reacts to attention-getting events.

B. Searches for relevant memories (by association and/or content).

C. Compares their anticipated results.

D. Selects those preferred on the basis of selected guidelines.

E. Puts them into effect by using them as reference levels for selected Orders within the hierarchy.

These and other topics can be discussed separately. Enough for now.

Martin Taylor: Bob says: "A Decision Making Entity (DME) is proposed as a partial solution to these problems. The concept seems to be generally taken for granted and accepted by many people—including most (if not all) of those on CSGnet. Such acceptance is demonstrated by the frequent use of the first person singular." Count me among the nonacceptors. Your DME sounds very much like the old homunculus who sits behind the sensors and effectors, manipulating. How does it work? Does it have its own little hierarchy?

Bob Clark: Martin, from my ("Observer's") viewpoint, your latest post is equivalent to two decisions: 1. "Count me among the nonacceptors" is equivalent to your having decided that the first person singular does not refer to a Decision Maker. 2. The posting of your decision to the net is equivalent to a second decision.

Who, or where, is the "me" included in your post and involved in creating it? Please explain your alternative(s), with or without using PCT. Remember, I am assuming a situation where both the established "nodes" or "choice-points" are, for any reason, unable to provide a "decision," and intrinsic error is neither present nor anticipated.

My dictionary gives "homunculus, n. 1. a diminutive human; midget. 2. a fully-formed, miniature human body believed, according to some medical theories of the 16th and 17th centuries, to be contained in the spermatozoon." I don't think you intend the term "homunculus" to be taken literally per the definition above. More important, in my posts I have tried to restrict the capabilities of the DME to those without which it could not perform its defining function: making decisions. Can any of these be omitted? Should any others be added?

How does the DME work? See my earlier posts. Here is another attempt to describe the essential characteristics of a Decision Making Process—a DME in operation.

It is assumed that no built-in automatic branch-point is available, no intrinsic error currently exists, and no intrinsic error is anticipated. I have tried to limit this description to those items without which decisions cannot be made. Thus, the proposed items are these:

Current perceptions. The DME selects the signals to which it directs its attention. They are selected from among the incoming neural signals available. These signals are available for use as feedback signals if needed.

Current objectives (reference levels, etc.), if not already in operation, are selected from recordings of past decisions, events, etc.

Past perceptions: recordings (memories). The DME finds recordings both by named addresses and by similarities of content. They could result from simple "recognition" ("reminders"), or (more or less) extensive searches for relevant material.

The recordings found are examined for relevance and possible application ("feasibility") to current perceptions (perceived situation).

5. The recordings are further examined, by imagination, for anticipated future effects as they relate to current, relevant reference levels.

The entire hierarchy is available to serve as the output function for the DME. In ordinary situations, only limited, selected portions will be needed.

On the basis of the above examinations, etc., the DME selects and activates a recording. The DME's selection can be arbitrary. The recording selected can consist of revised and/or combined recordings.

If this is a "homunculus," so be it.

For the Perceptual Control Theory of behavior to be complete, it seems essential to me that "decision making activities" are included somehow.

In addition, I think that these elements are consistent with most, if not all, of the ideas either stated or implied in Behavior: The Control of Perception.

Finally, to repeat: please let me know your (Martin's) procedures for making ordinary decisions, and what part (?) of you (Martin) does it.

Martin Taylor: Bob asks: 'Who, or where, is the 'me' included in your post and involved in creating it? Please explain your alternative(s), with or without using PCT." Two questions and an assumption. I recognize the existence of consciousness in me, and I extend you the courtesy of assuming it exists in you. I have no explanation of it, other than the simple presumption that its content must be based on signals in the hierarchy, and that it is not itself such a signal. Consciousness is a multi-dimensional experience. "Me" is an element of consciousness.

The assumption: that this is a "User's" viewpoint. What I mean by a "User's" viewpoint is that you can take account of only the signals accessible at that point. The User's viewpoint of an ECS is not that of a person within whom the ECS operates. It is consideration of what is accessible at some point within the ECS often the perceptual signal, but possibly one of the other signals. Your DME does not have a User's viewpoint of the action hierarchy. It has access to signals from all over the hierarchy.

What is a "decision" *within* the control hierarchy? It must happen at the program level or above (assuming Bill Powers' set of levels). Below the program level, there might well be multiple means to achieve any particular perceptual signal value, but the variation of means must be caused by differences in the reactivity of the world. The increase of difficulty (I sometimes say "impedance") of one lower-level control might mean that a higher-level perception is brought under a control by an entirely different set of actions. This is not "decision" as I understand it. It is a natural consequence of there being a non-linear system with more (in this situation) degrees of freedom for output than there are perceptual degrees of freedom being controlled at a high level.

Something nearer "decision" can occur within the hierarchy below the program level, when possible actions are played through imagination loops in various ECSs. I suspect that this happens all the time, and is not switched. The effectiveness or otherwise of this imaginary control might affect the real gains of different ECSs, resulting in different real patterns of action when the imaginary control is actualized. Again, there is the appearance of decision without any actual decision.

At the program level, "decision" is intrinsic to the level. It is the nature of the program level to select among sequence reference levels, and there, decisions have an explicit place within the hierarchy. I would think that they would be accessible also at higher levels. But that's pretty high in the hierarchy.

If you are talking about a Powers type of hierarchy, you must be talking about the program level or above, because below this level there are no choice points. The PIFs do not permit them.

How does one "anticipate" intrinsic error? One can't even perceive it when it does occur, according to Powers. I don't think it is relevant to the issue of the DME.

You say: "I don't think you intend the term 'homunculus' to be taken literally per the definition above." No, of course not. One of the reasons that behaviorist psychology became popular in the early years of this century was that people saw that most of the 19th-century psychological theories were recursive. To explain what a human did, they in effect passed the results of sensory processing to a "little man in the head" who decided which levers to pull and push to make the muscles work. All of the issues of the psychology of the human were incorporated within the LMITH, and he was usually called "the homunculus."

You also say: "The DME selects the signals to which it directs its attention. They are selected from among the incoming neural signals available. These signals are available for use as feedback signals if needed." On what basis is this selection made? What is the perception that the DME is controlling by means of varying its choice of neural signals? "For use as feedback signals" in what control loop?

And: "The entire hierarchy is available to serve as the output function for the DME. In ordinary situations, only limited, selected portions will be needed." So the hierarchy is the environment on which the DME operates, exactly as does the Powers reorganizing system? Your seven characteristics certainly seem to indicate this. But how does the DME itself operate? Is it controlling anything? If so, what can it be controlling but its own perceptions? And if it is controlling its own perceptions, do not the same considerations apply to it as to the main hierarchy: it is a hierarchy of perceptual control systems, needing a sub-DME to make decisions on its behalf, such as what signals in the hierarchy to attend to?

"If this is a 'homunculus,' so be it." Well, it still sounds like one, in that it solves an acknowledged problem within the control hierarchy by replicating the problem at a new level. The recursion, as with the original psychological conception of the homunculus, is potentially infinite.

"For the Perceptual Control Theory of behavior to be complete, it seems essential to me that 'decision making activities' are included somehow." Yes, but why must they be outside the control hierarchy? Isn't the program level adequate? Remember that in the Powers system, perceptual input functions may accept any neural signal as input, though in our diagrams and analyses we usually consider only the perceptual signals of the next lower level of control.

"Finally, to repeat: please let me know your (Martin's) procedures for making ordinary decisions, and what part (?) of you (Martin) does it." If I knew that, I would join the ranks of those making pronouncements about the truth of the world, and I might be rich in the bargain.

Look, my problem with the DME as an entity isn't a matter of faith that everything can be solved within the main hierarchy (though I like to think that true, and it is one reason I continue to think of local reorganization instead of postulating a separate reorganizing system). My problem with the DME is that it seems to do the same kind of job within the main hierarchy that the main hierarchy does in the outer world. That means that the DME must need its own DME, which needs its own DME, which.... In other words, introducing the DME does not seem to solve the problem it addresses. If I misunderstand what the DME is supposed to be, then I'm quite happy to retract all I have said. But I have indeed read your postings, and refrained from comment for lack of time. I simply didn't want silence to be taken as acceptance when you made that an issue.

Bob Clark: Martin, thanks very much for your prompt and thoughtful reply to my last post.

Before turning to your specific remarks, I'd like to state "where I'm coming from." It seems to me that the present theory is incomplete in certain respects, and that it would be much more useful if ways could be found to improve it. Here I point out two places where it is incomplete:

1. Behavior: The Control of Perception, Chapter 5, Memory, pages 220 ff., and Figure 15.3 showing two position switches. Here, the four possible combinations of the switches are described in terms of "modes." After discussing these modes, we find on page 224: "Note how skillfully I bypassed the question of *what* flips the memory switch. 'One' flips it! I plead guilty to obfuscation—the model obviously lacks some details which I am not now prepared to supply."

2. Chapter 13, Higher Levels, pages 173 ff.: "I must now account for choice of particular system concepts as ninth-order reference levels, and I can't." Also: "So I must say for the time being that this is my model of behavioral organization, as far as it concerns the ongoing performance of a competent adult human being. I must leave questions unanswered, hoping that others will find this approach interesting enough to expand upon and modify." Then: "Another possible—even probable—source of ninth-order reference levels is *memory.*" Further: "The solution that I prefer for this problem involves a discussion of learning of a particular type, and so will be presented later." This "learning of a particular type" is, of course, the reorganizing system, genetically determined and operating outside the hierarchy with intrinsic error providing the driving force. See Chapter 14, Learning. (Incidentally, this system was originally proposed as the "Negentropy System," with essentially the same properties as the present "reorganizing system." It was proposed in order to account for observed changes in the operation of the systems composing the hierarchy. It was regarded as operating "outside" of the hierarchy—without definition of "outside."

The "one" in the first item above is regarded as existing, somehow, somewhere. I am suggesting a more meaningful name that will help identify the "items" needed to accomplish the indicated results. Perhaps Decision Making Entity (DME) is not the best name for this, but some equivalent seems to me unavoidable. I have previously listed seven items which seem to me necessary for the DME to perform its switching function effectively.

The DME is proposed in order to account for certain observable events called, perhaps loosely, "decisions." Many of these are readily accounted for in terms of the existing hierarchical structure, including preset "choice-points," as discussed in *Behavior: The Control of Perception,* Chapter 14, Learning, pages 177 ff.: 'Programs are fixed lists of instructions (reference levels for lower-order systems in human beings) with choice point in the lists. Both memory and present-time inputs are important elements.... the same list of operations remains in use, and... the subprograms may retain their same organization. All that changes is the path followed through the network of contingencies, *all possible paths being determined when the writing of the program is finished."* (Italics added by me.)

Although a mature adult might have adequate programs to cover all possible situations, it seems unlikely. It seems especially unlikely for an infant, where a major part of its learning consists in learning such programs.

Operation of the reorganizing system might account for those "decisions" when an intrinsic error exists. "Decisions" made in the absence of intrinsic error require other operations.

Martin, you say: "Consciousness is a multi-dimensional experience." I'm afraid I don't know what this means. To me, "consciousness" refers to the condition of the perceptual systems. If they are in working order, the individual is "conscious." I think this is consistent with the following *(Behavior: The Control of Perception,* page 200): "Consciousness consists of perception (presence of neural currents in a perceptual pathway) *and* awareness (reception by the reorganizing system of duplicates of those signals, which are all alike wherever they come from)."

Martin, you also say that "'me' is an element of consciousness." Is "me," then, a subgroup of perceptual signals assigned the label "me"? In the same sense as the "tree in the yard"? Is it always passive? Is it sometimes active? In what manner, subject to what conditions, if any?

You say: "What I mean by a 'User's' viewpoint is that you can take account of only the signals accessible at that point." "Point" in the hierarchy, or "point" in time, or both?

"It is consideration of what is accessible at some point within the ECS, often the perceptual signal, but possibly one of the other signals." If "it" refers to the "User's" viewpoint, I don't understand this statement either.

"Your DME does not have a User's viewpoint of the action hierarchy." Since I don't understand your definition of User's viewpoint, I cannot comment.

"It has access to signals from all over the hierarchy." Yes, this is what I said.

You ask: "What is a 'decision' *within* the control hierarchy?" Your discussion here seems to consist largely of a description of the ordinary operation of the hierarchy, using its existing choice-points at whatever levels might be required. I am concerned with situations in which problem-solving programs are, perhaps, incomplete or otherwise unable to provide needed solutions. But, at the same time, no intrinsic error exists.

"If you are talking about a Powers type of hierarchy, you must be talking about the program level or above, because below this level there are no choice points!' So I am suggesting a situation where there is no suitable "program" available, with or without pre-existing choice-points. Perhaps this is rare, although at early stages of development there might be rather few effective "programs." Here is where a DME might produce action before any intrinsic error develops.

"How does one 'anticipate' intrinsic error? One can't even perceive it when it does occur, according to Powers." Martin, do you agree that it is possible to "anticipate" some possible future events? Do you plan your posts before sending them? In giving a talk, do you plan for possible questions or interruptions? Is a toothache an intrinsic error? Do you remember having one? Or any other intrinsic error? Do you go to the dentist to avoid a future toothache? Do you take action to avoid repeating a situation involving an intrinsic error?

Also: "What is the perception the DME is controlling by means of varying its choice of neural signals?" It is controlling its perception of the overall situation as it relates to an unexpected event. This includes its examination of those memories that seem to be related.

And: "... in what control loop?" In whatever control loop has a problem, but lacks a pre-existing problem-solving program.

After referring to my sixth point, you offer: "So the hierarchy is the environment on which the DME operates, exactly as does the Powers reorganizing system? Your seven characteristics certainly seem to indicate this But how does the DME itself operate? Is it controlling anything? If so, what can it be controlling but its own perceptions?" Very pertinent and important questions.

First, the DME has a strong resemblance to the 'Powers reorganizing system." It differs in that it only operates with respect to those perceptions that have its attention, whether selected arbitrarily or in response to some attention-getting event. This is in contrast to the reorganizing system, which is in contact with all intrinsic signals all of the time.

Second, as long as the established hierarchy has no problems, the DME need not be active, but it is capable of arbitrary action, perhaps "curiosity."

Third, its primary source of material to use as inputs to the hierarchy is the contents of the memory. There is little discussion, in *Behavior: The Control of Perception* or otherwise, of the contents of the memory. But, after all, without the memory there really is no hierarchy beyond the genetically determined neural pathways. Note, again, that the reorganizing system output seems to be arbitrary, if not entirely random, making no use of contents of the memory.

Fourth, after either the reorganizing system or the DME acts, the DME can review (in imagination) those actions as they were recorded in the memory. They can then be described in terms of a "problem-solving program," complete with choice-points. Indeed, this additional program becomes available for future use. It seems to me that this is pretty much the way these problem-solving programs come into being.

Thus, the DME needs no hierarchy of its own. When there is a problem, it provides assistance to the hierarchy based on selection from the contents of the memory.

In our early work, Bill and I were both greatly concerned with possible circularity or other recursive properties of the developing theory.

The DME requires at least the beginnings of a recording function and the formation of memories. These are included, as noted above and in Behavior: The Control of Perception Chapter 15, Memory, within the present theoretical structure.

Martin, you say that your "problem with the DME is that it seems to do the same kind of job within the main hierarchy that the main hierarchy does in the outer world." If I thought so, I, too, would reject the DME concept. Instead, I see it as playing a critical part in the development of the hierarchy.

I hope that I have shown you where and why I think the DME concept offers a useful extension of the original theory.

Thanks again for your interest, questions, and ideas.

I've been essentially out of touch with PCT developments and discussions for some 25 years! Migawd, Bill published his book in '73, and I've been away from Chicago since 1968! (There was a brief correspondence with Bill in 1987.) Indeed, I only learned of Bill's book from Greg Williams in late 1988! Greg was also kind enough to provide me with a copy of Robertson's book. These contacts have led me to join the CSG and the net. I am still not familiar with the several viewpoints represented by the members of the Group.

However, during those years I have been applying the ideas Bill and I initiated, adding my observations, and developing my viewpoints, while engaged in other activities. I have had opportunities to work with and study a variety of organizations (including business, manufacturing, lobbying, political, government, sales, tax exempt, etc.). I have been in a position to participate in and observe their operation, planning, viewpoints, concepts, etc. I have also played a major part in the formation and operation of several organizations.

I find that having the underlying concepts of a hierarchical array of feedback control systems readily at hand has been very useful in all of these activities, enabling me to participate at whatever level interested me.

Currently, I am exploring possibilities for involvement with local business, school, and government activities. I don't yet know how it will work out, but it should be interesting! It is intriguing to observe how ordinary people doing ordinary things recognize and use the concepts underlying Pa without any need for formal technical, scientific ways to communicate them.

Propagation of these theories is not my primary purpose—rather, I hope to find ways to assist people in achieving their goals. I expect that they will gradually find that certain PCT terms and concepts are helpful.

Bill Powers: Bob, I no longer think of intrinsic error as limited to purely physiological variables. For example, the presence of chronic significant error signals in any control systems of the brain is itself an indication of something amiss, and would fit the basic definition of an intrinsic error. It's also possible that the scope of the reorganizing system has evolved along with the structure of the brain that permits us to develop higher levels of control systems. So I don't object too much to your concept of the DME, which apparently operates in terms of criteria considerably more advanced than physiological states ("appropriateness," for example). Perhaps your DME is simply a more evolved version of the primitive "Negentropy System" with which we began almost four decades ago.

I do have one argument with your DME, which is that it seems to have many capabilities that I would rather see as aspects of the learned hierarchy. In my development of ideas about levels in the hierarchy, I tried to isolate types of perceptions that at least in principle could be controlled by learned control systems. Anything of that nature clearly doesn't belong in the system responsible for shaping organization, because what is learned is not present at first, yet the process of reorganization has to work from the beginning. I see too much that is systematic and algorithmic in your descriptions of the DME and what it does. If those were stripped away and assigned to the learned hierarchy instead, I think our concepts would come much closer together.

Bob Clark: Bill, I think we are, in fact, very close indeed. Dividing ideas into their components—and naming them—can be very helpful. Your remarks relating the old "Negentropy System" to your evolving concept of the "reorganizing system" suggest a need for some form of "intermediate system." Such a concept can be placed on a continuum, with a "pure intrinsic error-driven system" at one end and a "random, arbitrary curiosity system" at the other, leaving your evolving reorganizing system to combine with a "structured memory-using" system—the DME—in between.

The "pure intrinsic error-driven system" operates at a level of desperation, having been driven beyond organized efforts, no longer able to access existing memories. At the same time, the Recording Function continues to produce records that become available for later access.

The "curiosity system," on the other hand, operates when the individual is awake, alert, and without any immediate actions called for. Perhaps this is close to a state of "boredom," or perhaps is a "standby" condition, waiting for something to happen. The Recording Function would also (of course) continue to produce records that become available for later access.

The intermediate system (DME) would be characterized by its use of memories as the source of ways to achieve the goals (provided by higher levels), that are currently dissatisfied. As previously described, it would use memories as a guide for the selection of promising actions in seeking the goals in question. This would include anticipation (via imagination) and application of learned problem-solving programs

where they seem useful. Memory would also be a source, perhaps resulting from application of problem-solving programs, of reference levels for application throughout the hierarchy.

Of course, this entire process could be no better than the assortment of memories available to the individual at the time they are reviewed. Since this entire process is recorded together with continuing current perceptions, the result can be considered a form of "reorganization," at least of the learned systems.

You comment, Bill: "I see too much that is systematic and algorithmic in your descriptions of the DME and what it does. If those were stripped away and assigned to the learned hierarchy instead, I think our concepts would come much closer together." I sympathize with your view here. These "systematic and algorithmic" aspects are partly due to my difficulty in describing my concept of the operation of the DME without using language and concepts typical of the learned hierarchy. I have tried to distinguish between the "pure" DME and a description of the logical requirements for it to perform as defined. The suggestion of a "continuum" might be helpful, with "purely learned" reorganizing systems supplemented on occasion by action of the DME. This results in the role of the DME being a bit "mixed," in the extremes, with those of the "pure reorganizing system" and the "curiosity system." Perhaps we can devise better ways to describe and distinguish among these concepts.

After all, these verbal systems, as well as theories in general, exist only in our memories (and memory supplements called "books," "periodicals," etc.)!

I hope we can move ahead with this, Bill, because there are several more areas for discussion.

You might be interested in the developments as I become more involved with the local city, Forest Park. By selecting suitable time scales, all aspects of Hierarchical Perceptual Control Systems become apparent! This includes intrinsic error, learned systems, reorganizing systems/DME, and curiosity! Fascinating!

Forest Park, Ohio: Population about 20,000, about 30 miles north of downtown Cincinnati. About 600 businesses, about 75 of which are members of the Forest Park Business Association. A few years ago, I helped revise the by-laws of the FPBA—I was a member of the FPBA Board at the time. Forest Park's government uses a City Council-Mayor with a City Manager. I have a copy of their Charter, which impresses me very favorably. (At one time or another, I have been involved in writing/revising various by-laws as well as working with the results, so I have some basis for evaluation.)

It didn't take long to identify each level of an HPCT system as it operates. In addition, the major orientations of several individuals were observable as they cooperated and interacted in supporting their mutual interests. These observations are also helpful identifying labels for the several levels.

It is fascinating to observe the way the participants think of (read: "perceive") their own actions and interactions. By and large, their thinking is PCT-type thinking—they have goals, personal- and community-, which they are working toward achieving. It is very straightforward, with very little S-R contamination.

I find Control Systems Group members repeatedly concerned with getting some kind of recognition from the "scientific community." This is a losing game—the "scientists" hold the cards. They select the independent variables, the dependent variables, the time scale—and arrange to have any relevant reference levels held constant. Any deviation from these rules guarantees rejection.

Instead, how about working with applications of HPCT? I've been doing it, informally, for about 30 years. This can be done without having to teach anyone the special lingo used in HPCT—most of the key concepts needed in HPCT are already available in everyday language. One only needs to look for them. The applications could be presented (possibly for publication?) in a form that shows how these methods work and how the ordinary language of application can be expressed in generalized abstract theoretical terms. After all, a great many "practical" applications were used in many fields long before "modern" theoretical methods were developed.

Applications that HPCT Might Explore

Learning in contrast to teaching
Conflict resolution (internal, interpersonal, intergroup, etc.)
Social systems
Economic systems
Management principles
Government (at all levels)
Argument vs. persuasion
Marketing and sales
Decision making by groups
Motives
Emotions
Cooperation
Personality
Development (of individuals and species)
Genetically determined neural systems
Planning
Character

Anticipation
Memory playback
Freedom
Responsibility
Consciousness

Quite a list—presented here "off the top of my head." Most of these have concerned me at one time or another, and they seem to be to be highly relevant to HPCT treatment.

In several places in *Behavior: The Control of Perception,* Bill notes some uncertainty and expects later additions and revisions. I find the "established" designations of higher levels very logical—but not very useful in attempting to work with "real people." Sometimes the PCTers don't look very "real" to me.

Bill Powers: Bob, putting the concepts of PCT into ordinary language as you suggest is a fine idea, and I endorse it. There are sticky spots in doing this, however: those where PCT and common sense part company. Many people speak of emotion, for example, as if it's something that the outside world does to them, and with which they must then try to cope. It's not easy to present a compelling case in ordinary language for the idea that emotion is part of voluntary action and is the product of the person's own attempts to seek goals.

The higher-level definitions of behavior in HPCT aren't meaningful until you translate the terms into real experiences. For example, in your interactions with government types, you have probably seen that many of them state "facts" about human nature—what "people" are like, what to expect of them, and so on. These are system concepts, as I think of them. You probably also hear many people stating generalizations; not specific programs for actions, but *principles* of action. In government, they are often called policies, where the program-level stuff consists of rules, regulations, or laws stated in if-then terms, designed to suit an overall policy at the principle level such as equal treatment, fair pay for adequate work, loyalty, and so forth.

The concept of PCT itself is a system concept. It is composed of principles like control of input and resistance to disturbance, which describe no particular control system but are meant to apply to all control systems however they are designed. At the program level, control becomes a mathematical-logical model containing specific quantitative relationships, no one operation being a control system or accomplishing control of input in itself. The "emergence" of control from the quantitative relationships among parts of control systems is evident only to a higher level of perception, the principle level at which we perceive the principles of control. And from these emergent principles, once we can perceive enough of them, there emerges the concept—the yet-higher-level perception—of an autonomous self-organizing hierarchy that constitutes human nature itself: a system concept.

Once you start translating from the too-formal terms of HPCT, you can begin to see the phenomena to which these terms were intended to point. Everyone has principles and system concepts. All you have to do to believe this is to sit in a blue-collar lunchroom day after day and listen in on the conversations. Listen to the people talking about union problems, about work rules, about unfair treatment given to one person or another. Ask their advice on how to get along in the company, and you will be drenched in principles. Ask how they think the company should be organized, and you'll get clear statements of system concepts, not to mention lots of descriptions of errors at that level. (All of which, I must admit, makes me wonder about the relationship of language to these levels—how can such things be described? Some aspects of language must surely operate at the system-concept level, too or higher.)

Behind the simple terms in my proposed levels, there are phenomena that I think are quite real and observable, at all of the levels.

Bob Clark: Bill, my suggestion regarding "ordinary language" was for PCTers to select from and use that language. Listening to others can be very useful, revealing much about the concepts, ideas, theories, observations, etc. that are, in fact, in use by "real people." To me, your "simple terms" are far from simple, and I agree that there are such "real and observable" "phenomena." I find it much more useful to work with these phenomena, rather than your abstract (and reasonable) classifications. To do this, I look for words/concepts with more immediate relations to the levels of the hierarchy.

Bill Powers: Bob, your recommendations about ordinary language are very much to the point for communicating control theory under everyday circumstances. This is really the "end-around play" that Dag Forssell proposes—forget about the Establishment of psychology and go directly to the people. Control theory, however, contains concepts which are not already in the vocabularies of many audiences, and the existing words usually mean something that has to be overcome before the wanted meaning can be communicated. By trying to make PCT concepts seem *too* familiar, in the hope of getting a friendly reaction from the audience, one can end up convincing them only that there's nothing new in it.

The Control Systems Group is a membership organization which supports the understanding of cybernetic control systems in organisms and their environments: *living control systems.* Academicians, clinicians, and other professionals in several disciplines, including biology, psychology, social work, economics, education, engineering, and philosophy, are members of the Group. Annual meetings have been held since 1985. CSG publications include a newsletter and a series of books, as well as this journal. The CSG Business Office is located at 73 Ridge Pl., CR 510, Durango, CO 81301; the phone number is (303)247-7986.

The CSG logo shows the generic structure of cybernetic control systems. A Comparator (C) computes the difference between a reference signal (represented by the arrow coming from above) and the output signal from Sensory (S) computation. The resulting difference signal is the input to the Gain generator (G). Disturbances (represented by the black box) alter the Gain generator output on the way to Sensory computation, where the negative-feedback loop is closed.

Back cover

Closed Loop

Threads from CSGnet

This reproduction of *Closed Loop* was created by Dag Forssell in 2001. Addresses and phone numbers have not been updated. Most are obsolete.

Posted at www.pctresources.com

Proofread as of March 8, 2001

Fall 1993

Volume 3

Number 4

Front cover

Closed Loop

Threads from CSGnet

Fall 1993 Volume 3 Number 4

Edited by Greg Williams, 460 Black Lick Rd., Gravel Switch, KY 40328

CONTENTS

Members of the Control Systems Group receive *Closed Loop* quarterly. For more information, contact Ed Ford, 10209 F. 56th St., Scottsdale, AZ 85253; phone (602) 991-1860.

CSGnet, the electronic mail network for individuals interested in control theory as applied to living systems, is a lively forum for sharing ideas, asking questions, and learning more about the theory, its implications, and its problems. The "threads" in each *Closed Loop,* stitched together from some of the net's many conversations, exemplify the rich interchanges among netters. Some issues of *Closed Loop* also feature research reports by netters, in hopes of initiating new conversations.

There are no sign-up or connect-time charges for participation on CSGnet. The Internet address is CSGL@VMD.CSO.UIUC.EDU while CSGL@UIUCVMD is the Bitnet address. Messages sent to CSGnet via these addresses are automatically forwarded to over 120 participants on five continents, as well as to hundreds of NetNews (Usenet) sites where CSGnet can be found as the newsgroup bit.fistserv.csg-l. CSGnet also can be accessed via CompuServe, AT&T Mail, MCI Mail, or any other computer communication service with a gateway to Internet or Bitnet. For more information about subscribing to CSGnet, contact Gary Cziko, the network manager, at G-CZIKO@UIUC.EDU, phone him at (217) 333-8527, or send a FAX to (217) 244-7620.

Two Views of Control-System Models

Hans Blom: What is the fastest way to get a spaceship to Mars? The solution is well-known, although impractical: apply full thrust until you are at the exact midpoint of the trip, turn your ship around and apply full thrust again, braking until you arrive at Mars with zero speed. This is an example of what is called "bang-bang" or "minimum time" control, a control paradigm quite different from the "stabilizing control" that is usually discussed on CSGnet. Features of bang-bang control are these: (1) outputs are either zero or at their maximal limits, (2) the only important parameters are the times at which outputs go from zero to maxima, or from maxima to zero, (3) in general, it is quite difficult to find optimal values for those times, and (4) for long periods of time (between the decision points), it might seem to an outside observer that control is absent, because nothing changes—because there is no modulation of the outputs.

This is discussed in "The Neural Control of Limb Movement," by William S. Levine and Gerald E. Loeb, in the December 1992 issue of *IEEE Control Systems*. Does the organism use bang-bang control? No. "The experimental data... show a substantial deviation from the optimal control model." Why is that? Partly in order to protect the organism: "the feedback from the joint sensors, while certainly present, would be too late to prevent injury if a human jumper tried to perform a mathematically optimal [i.e., top-performance] jump." And partly because "it is important for both biologists and control engineers to remember that the control systems that have been invented to date are almost certainly a meager subset of all possible types of control and even of all control methods used in biological systems. Thus, the study of biological systems should not be confined to testing whether their performance is compatible with control schemes invented to date but must include detailed examination of their inner workings to discover new types of control."

Some type of stabilizing control is needed in all cases where full-time control relative to a setpoint cannot be relinquished even for a moment. But stabilizing control is incompatible with top performance, such as in sports. In high jumping, only the maximum height of the jump is important, not the full trajectory. In the Mars rocket, the output resources are used at 100% capacity during 100% of the time; the only control decision is to find the exact point in space-time of the turnaround. Mathematically, due to the nonlinearity of the problem, finding this point is generally intractable and therefore usually a matter of trial and error (search) or creative insight. In humans, finding the optimal decision points requires a considerable period of tuning and fine-tuning (training).

The authors pose more questions than they provide solutions: "much more work needs to be done before the above suggestions can be called a theory." Yet, in my opinion, this paper provides some insights into why stabilizing control, which works so well in ordinary circumstances, breaks down when maximum performance is required.

Bill Powers: Hans, human control systems are pretty close to the design limits set by the materials used. It's possible, for example, for an arm muscle to pull itself loose from its attachments to the bones, if feedback is lost and an energetic movement is attempted. Even with an intact set of control systems, tendons and muscles can be ripped loose if an emergency situation results in sending abnormally large reference signals to the spinal motor neurons.

The "substantial deviation from the optimal control model" that Levine and Loeb mention might not be a deviation from what is optimal for the whole human being using the control system. Control models of an arm usually propose the application of torques at each joint, but in the human system there are no motors at the joints. Instead, there are nonlinear muscles attached in clever ways that produce many kinds of torques, some through clever linkages (as in the two bones of the forearm), and some by having the muscle wrap around the joint in a strange way (like the biceps).

Even the muscles work differently from the servo motors that engineers use. They don't apply forces directly, but by shortening the contractile elements in the muscle to alter the resting length of the series spring component. In principle, a movement could be carried out by suddenly shortening all of the contractile elements in a muscle and storing energy in the spring components, then letting the spring components execute most of the movement without any further expenditure of muscle energy until time for deceleration. Actual movements work somewhat in this way. This is something like the solution for maximum rocket efficiency given a finite fuel supply. In fact, the human system is far more efficient than any robot so far invented; it moves 100 to 200 pounds of weight around all day expending only two or three kilocalories of energy and using less than 0.1 horsepower of total muscle output power. And the fuel supply has to support not just the muscles, but the brain and the general metabolic requirements.

The reason a human being can't perform a mathematically optimal jump is simply the rocket problem: you would need to produce an impulse of muscle force of zero duration and infinite amplitude. That would hardly be a feasible solution for a servomechanism, either.

The "feedback too slow" argument turns up even here, doesn't it? Actually, the speed of feedback in a human control system is just right—to explain the behavior we see.

And you also say: "But stabilizing control is incompatible with top performance, such as in sports. In high jumping, only the maximum height of the jump is important, not the full trajectory." Human beings hardly ever control the "full trajectory." They control the variables that matter to them. Rodney Brooks has the right idea here: don't plan trajectories, avoid obstacles. It isn't necessary to know where obstacles will be, if the system has sensors that can detect proximity to an obstacle.

"Stabilizing control" is something of a misnomer, suggesting that all that a control system does is to keep something constant. More generally, it makes the perceptual signal track the reference signal. This means that a control system for producing a directed force (as in throwing a ball or launching a high jump) can make the sensed acceleration have the right magnitude and direction right up to the moment of release. When we learn how these perceptions must change in order to have a desired result remotely or later, we vary the reference signals to repeat the experienced thrust as nearly as possible, and we get pretty dose. Of course, if we got too dose, people would stop doing such things—or they'd set the bar higher, or put the target farther away, until errors in control once again made the game interesting.

I think that when normal human movements such as walking are finally modeled fully, we will find that the system uses as little energy as possible, letting momentum and spring effects carry most of the movement through, with muscle contraction being used primarily to trim the result into a useful form. When we walk, we choose a pattern of walking to control that is as dose to the zero-energy pattern as possible, given the higher-level goals of actually getting somewhere in a reasonable time. Only when we have some reason to get there faster, as in running a rare, do the control systems try to produce patterns that cost a lot of energy. And even then, the patterns finally chosen are pretty efficient—after all, the fuel supply and distribution have to suffice to get the body to the finish line.

Hans Blom: Bill, you say: "Human beings hardly ever control the 'full trajectory."' If that is the case, "new types of control," which do not try to maintain minimum error between reference values and perceptions at all times, might provide superior performance in some cases. Or greater ease. When I fly to New York, I (attempt to) control my destination, but in the plane I have to trust the pilot. Part of my trajectory will be, as far as I am concerned, ballistic.

What makes control in organisms so difficult to study is the simultaneity of a great many different ongoing goals, whose importance might, moreover, fluctuate from moment to moment due to influences beyond our control and usually beyond our knowledge. Only in the simplest of experiments one variable can be considered to be controlled, if at all. "Keep your finger pointed at the knot." But the subject also has to control the upright position of his or her body and otherwise keep all sensory channels open, if only to hear you say, "You can stop now."

Still, a high jumper wants to jump as high as possible, period. An objective measure is provided to test that performance. All else is unimportant (within limits, see below). What more can you ask for? There is no prescribed trajectory to be followed; a new world record often is an unprecedented experience for the jumper.

Human control systems normally function well within design limits. We have very little experience with operation near those limits: pain effectively causes us to stay away from them. But pain is carried by slow nerve fibers; in emergencies, the experience of pain can arrive too late to prevent harm. Is a case where "tendons and muscles can be ripped loose" really an indication of "an intact set of control systems"? I consider that to be pathology, a control system gone haywire, operating beyond its design limits. I would maintain that one of the most important of an organism's objectives is, at all times, not to seriously damage itself. But that cannot be formalized by control in the usual sense of the word, that is, a perception following a reference signal. The control system is operating under constraints, i.e., it must stay away front certain experiences with a high probability of success. Short-term goals are rarely important enough to jeopardize long-term goals, which need an intact organism.

You say: "The reason a human being can't perform a mathematically optimal jump is simply the rocket problem: you would need to produce an impulse of muscle force of zero duration and infinite amplitude. That would hardly be a feasible solution for a servomechanism, either." Impulses are not required, step functions will do nicely. After all, a trainer just wants to study the peak performance that a real individual is capable of, given his / her motor equipment, and search for whatever means there are to teach him / her to fire his / her nerves in such a way that this peak performance is reached.

Also, Levine and Loeb do not say that feedback is too slow; bang-bang control requires very accurate timing. They say that when the need for performance becomes extreme, protection mechanisms are required to prevent muscles and tendons from being torn loose. Feedback from those protective sensors would probably be too slow if training did not slowly familiarize the high jumper with the sensations that they provide. (Much of psychotherapy seems to serve the same function: trying

to get the client "in contact" with his / her feelings without being overwhelmed by them.) This is much like walking as dose to the abyss as you dare without risking the damage that a fall would cause. The fall would provide you with feedback, of course, but you wouldn't want that feedback, would you? (In psychotherapy, one of the frequent goals is to show the client that much of his / her "fear of falling" is imaginary, and that the abyss is much farther away than he / she thinks. This, too, is a difficult and often fearful type of exercise.)

You say: "Human beings hardly ever control the 'full trajectory.' They control the variables that matter to them." Yes. And bodily (and mental) integrity matters a great deal.

You also say: "'Stabilizing control' is something of a misnomer, suggesting that all that a control system does is to keep something constant. More generally, it makes the perceptual signal track the reference signal." Exactly how would you know that the jumper follows a reference signal when for the very first time he / she jumps higher than he / she ever did before? How does the reference signal get established in the first place? I do not allow the answer that it is an "imagined" reference signal; that would be impossible to either prove or refute, and it would therefore be unscientific (following Popper). I do allow the answer that the reference signal is discovered "by accident," through trial-and-error learning. But that would mean that the very first time there was no reference that could be followed, i.e., that not all behavior (here: peak performance) is control of perception.

Perception is not the only human capability that we depend on to control our behavior. Sometimes memory will do: a child will stay away from a hot stove after having been burned by it only once. Sometimes "knowledge," such as from a newspaper, will do: stay away from Chernobyl for a while. In neither case do you control for an exact distance from the feared location, you just want to keep at least a minimum distance away from it.

Maybe we have a different conception of what perception is. For me, perception is everything that my senses register and what can be derived from that. You might include memory as some type of "observation" through "inner senses." Is that what you mean?

That leaves the discrepancy of wanting something and not wanting something. More philosophically, I think that this distinction explains what gives us freedom. *There is* not one optimal location that is dictated by a match between our inner drives (reference levels) and our perceptions of the outside world. I do not dispute that we have reference levels and that we use our perceptions to get us dose to them. I just want to add something like "negative reference levels," things to stay away from. Freedom is a name for ranges in n-dimensional objective space where you can move about "at will," because the objective function is flat. It is as if you try to find the highest peak in a mountain range, and once you get there, you discover a wide, high-altitude tableland.

An example: you get conflict when the heater is set to 22 degrees Celsius and the air conditioner to 20 degrees. You get a region of "freedom" if the heater is set to 20 degrees and the air conditioner to 22 degrees.

As a control systems designer, I do not create control systems in the hope that they function correctly; hope has no place in the model. I do not rely on things going right only usually. I specify an objective function that I know will lead to a correct design. And if I cannot guarantee correctness, I will at least strive for optimality in some sense, such as longest mean time between failures or longest time before first failure. Would evolution be sloppier, given its billions of years of experimentation?

I assume that evolution, through a harsh billion-year-long struggle for survival, could have come up with some pretty clever solutions to the control problems that have arisen. *E. coli* has a funny (partly random) but clever control law that results in what is called a biased random walk This "primitive" control law serves it quite well; E. coli is far more numerous than *Homo sapiens*. Higher organisms have other (better?) control laws, some of which we seem to have more or less uncovered (control of voluntary muscles in humans) and which resemble linear quadratic control, at least as long as muscles function well within their force limits. Linear quadratic control works well in stabilization, i.e., stand-still and slow movements. In other cases, we know that there are better control laws. An example of that is when peak performance is required and the forces that muscles can deliver come to their limits. In that case, the nonlinearities of the actuators cannot be neglected any more, and linear quadratic control becomes sub-optimal. Intuitively, I agree with Bill Powers when he supposes that there is only one control law that governs the control of muscles. Linear quadratic control is, in my opinion, its more readily understandable "special case," just like Newtonian physics is a more readily understandable special case of general relativity.

People are very good (but often highly nonlinear) controllers. Moreover, it is my perception that people have a whole range of control schemes and frequently even apply the appropriate one at the appropriate time. This is a continual source of amazement (and envy) for control engineers who generally do much worse.

I know by now what perceptual control theorists mean by the mantra "organisms control perception." As so often with jargon, it is an abbreviation for a whole philosophy and only understandable for those who have gotten to know that philosophy. It is right, from a certain perspective. From another perspective, organisms control their outputs. I find it hard, in a control loop, to see one apart from the other. But, of course, sometimes you concentrate on the one, sometimes on the other. Very often, the output is controlled as well, for example in cases where different actions are possible (steak or salmon?), all leading to similar perceptions (great food!). Then you actively have to choose between outputs ("I would like...").

Perception is controlled by actions; actions are controlled by perception Remember the loop!

I agree with Bill Powers' "it's all perception" in the sense that perceptions (of the outside world and of our inner physical and mental mechanisms) are the only sources of information available to us. But perceptions are built upon and result in higher-level things that I would not call perceptions any more. Beliefs, superstitions, the "facts" of our lives. All those together constitute what I call a model (of the world, ourselves included). A model is, technically, always a simplification, and always has a purpose. That it is a simplification is due to the facts that we have experienced only a limited set of perceptions, and that our processing of those perceptions must be done by a mere three pounds of flesh. Models are never unique; it is always possible to translate one model into another, equivalent one. Sometimes a simple, approximate model works well enough, sometimes only a very complex and very accurate one will do, depending upon the goal that it serves. The highest purpose of the biological model is, in my opinion, best described by Dawkins: transmission of genes. Everything serves that supreme goal. The evolutionary process has weeded out every organism that did not serve its purpose well enough. A high degree of optimization has taken place during billions of years, and in that sense all currently existing organisms can surely be called well-designed control systems. Control systems, because they need to achieve a goal. There are numerous ways to achieve that goal. Viruses, bacteria, cats, and humans do it differently, thus far equally successfully. All other goals are sub-goals, designed through evolution to serve the one supreme goal. The sub-goals of each organism are uniquely related to its potential for actions, i.e., its body. A virus needs very few perceptions to achieve its goal; it mainly relies on the forces of nature ("free" energy) to work for it. A human, on the other hand, cannot survive without a great many perceptions.

In short, I think that the PCT perspective is extremely valuable when you study human behavior. A different perspective might be better for me, because I study very simple things like control systems. Let's by all means keep exchanging perspectives! Sometimes it seems less limiting to have two different perspectives on the same reality at the same time. Could that be why binocular vision proved to be successful?

Control engineers have a broader conception of control than you seem to have. Control does not necessarily imply feedback. In fact, engineers prefer non-feedback systems if at all possible, because they cannot possibly have stability problems. Regrettably, non-feedback control is possible only if the system to be controlled in invariable and not significantly subject to disturbances.

I think that by now I understand what perceptual control theory is about. I have followed and enjoyed the discussions for more than a year now, mostly quietly. Once in a while I grab the chance to vent some of my ideas, which are more or less related, hoping for a useful reply—usually not in vain. Reconciliation is not what I look for; I find that friction—clashing points of view—generates much more creative energy. Engineers and psychologists are not dose neighbors. They speak different languages, have a different culture, and work on different problems, although it is fascinating to discover similarities. I believe that engineers can learn as much from psychologists as the other way around. Doesn't this net show it?

Bill Powers: Hans, you say that "'new types of control', which do not try to maintain minimum error between reference values and perceptions at all times, might provide superior performance in some cases. Or greater ease. When I fly to New York, I (attempt to) control my destination, but in the plane I have to trust the pilot. Part of my trajectory will be, as far as I am concerned, ballistic." I think you're going about this backward. When we study human behavior, we aren't comparing it with some "optimal" or "best" way of controlling. We're just trying to understand what people are actually controlling under various circumstances. In some regards, people control things very well indeed, by clever means that surpass what any engineer knows how to build. In other ways, people control stupidly and poorly, and suffer the consequences.

More to the point, people use the means available to achieve whatever degree of control is possible. When I buy a ticket on an airplane, show up for the flight, and strap myself in, I have done all that is possible to get myself to the destination by that means of transport. So that's all of the control I have; if the plane is hijacked to another destination, that disturbance is beyond my ability to resist. All I can do is wait until the plane lands and I can get off it, and then start controlling again for getting to the destination by some other means. It could easily be that I would have arrived at the destination sooner, even without the hijacking, by taking a bus. But I didn't think of that. People are not optimal controllers; they just do the best they can.

You say: "What makes control in organisms so difficult to study is the simultaneity of a great many different ongoing goals, whose importance

might, moreover, fluctuate from moment to moment due to influences beyond our control and usually beyond our knowledge." The hierarchical model helps here, because higher-level goals change more slowly than lower-level goals. Many of the fluctuations in conditions are just disturbances, which lower-level systems automatically compensate for by adjusting lower-level goals. Much of the apparently chaotic nature of behavior becomes more understandable when we ask about higher-level goals. We can then understand many external events as disturbances, and we can see how the changes in detailed behavior oppose their effects. This reveals regularity where formerly we couldn't see any. I think that most behavior is actually quite regular, once we understand what's being controlled at many levels.

You're right about the fact that more variables are under control than we can measure in any one experiment. But it's interesting that without much trouble we can get those other variables to remain constant enough to get good repeatable data.

You say: "Still, a high jumper wants to jump as high as possible, period. An objective measure is provided to test that performance. All else is unimportant..." The highest-level goal is to win the contest, not to jump as high as possible. There is strategy involved, as well as just trying to produce maximum effort. Some jumpers will pass at a certain height, saving their strength for later. they don't try to jump at all. Also, if you assume that every time you see a high jumper, the objective is to jump as high as possible, you will usually be wrong; most of the time, the high jumper is just trying to go high enough to clear the bar. On many other occasions, the jumper might not be concerned at all with controlling for height. The jumper might be working on the approach or the takeoff, or the form at the peak of the trajectory, or the flip that raises the legs at the critical instant, and not be worrying at all about maximum height. You can't tell what a person is doing just by looking at what the person is doing. The Test for the Controlled Variable helps you to understand what is actually being controlled (as opposed to what you logically assume is being controlled).

You say: "Is a case where 'tendons and muscles can be ripped looser really an indication of 'an intact set of control systems'? I consider that to be pathology, a control system gone haywire, operating beyond its design limits." Certainly it is. If pathology is involved, it is a higher-level system that is misusing its lower-level control systems. Is it pathological for a father to lift a Volkswagen off his child, suffering born muscles and ligaments (and a lot of pain) as a result? When a person shoots himself in the head, all of the control systems for grasping the gun, aiming it, and pulling the trigger are working perfectly well until the last moment; all that's haywire is the higher-level system that has chosen this outcome. And even that choice might not be pathological, if the person is facing torture or the pain and humiliation of a vicious disease by staying alive.

And: "I would maintain that one of the most important of an organism's objectives is, at all times, not to seriously damage itself." Normally, perhaps. Not always.

You say that "bodily (and mental) integrity matters a great deal." I disagree. This is like saying that organisms control for "survival." Organisms control specific variables relative to specific adjustable reference levels. An outcome of doing so might be that the organism "survives" or preserves "physical and mental integrity," but that is not a concern of the organism. It's an opinion of a third-party observer. I don't think that there is any reference signal specifying survival or integrity. Organisms don't survive or preserve their integrity, anyway. They all die.

You ask: "Exactly how would you know that the jumper follows a reference signal when for the very first time he / she jumps higher than he / she ever did before? How does the reference signal get established in the first place?" The trajectory is a side-effect of controlling variables that the jumper can control. It is not itself a controlled variable. Once the jumper has left the ground, there is no action that can alter the trajectory of the center of gravity. There are, of course, many variables that can be controlled during the trajectory, such as the relative configuration of the parts of the body. These can make quite a difference in whether the bar falls or not, but they have no effect over the path followed by the center of gravity. One of the tricks of high jumping is to control the body's configuration so that the center of gravity passes under the bar while the body itself passes over it. That process is under continuous control all during the trajectory.

I think that competitors control what they can control: the approach, the takeoff, and the body configurations. The outcome depends on how well they are able to control those variables.

The peak height of the trajectory, perceived over dozens or hundreds of occasions, might be a controlled variable if there are things the jumper can do to affect this average peak height. The associated control system would be very slow, and would operate by adjusting many lower-order reference signals for such things as practice time, amount of effort, adjustments of form, and so forth. During any one jump, of course, this averaged perception can't be controlled. But over time, the jumper can gradually raise the reference signal for height jumped, as long as this is consistent with maintaining the necessary elements of the jump in the right forms. On the initial jump of a competition, no jumper strives for maximum height. The reference height is set comfortably above the bar, but no higher than necessary.

I think you would have a dearer picture of the PCT approach if you

kept the hierarchy in mind. The first time anything is accomplished, there can be no reference signal derived from experience of accomplishing it. At worst, one can have reference signals only for the lower-order components of perceived behavior that are to be put together in a new way. There are many possible ways for that to happen, including instruction followed by imagining the meaning of the instructions. At best, you've studied movies of someone else doing it and have some concept of the coordinations required.

On the first attempt, one seldom achieves perfect control. But the first attempt provides a perception of doing the control action, and from that experience, more realistic reference signals can be selected. Also, the new control system's parameters are probably not set to the best possible values; reorganizing them takes many trials, too.

To speak of "the" reference signal being "discovered" doesn't sound right to me. A reference signal is variable; it can be set to high or low levels. In any complex behavior, reference signals must be varied during the behavior if high-level perceptions are to be controlled at their given reference levels. Even when a behavior is well-practiced, the reference signals can be set to different states within the possible range. As I said, a jumper doesn't set a reference signal for the maximum possible jump early in the competition; you don't see champion pole vaulters clearing a 15-foot bar by five feet. I don't think that "maxima" have anything to do with it, anyway. The jumper simply sets a target height that is enough above the bar to clear it. When the bar is set too high, the target is still set above the bar, but now the jumper can't produce lower-level control actions sufficient to dear the bar, and fails.

If a jumper really set a reference signal for "maximum height'' (say, one kilometer), there would be an enormous error signal, and the output function would saturate, destroying control. To achieve maximum performance, one should set the reference signal just slightly above the level that the maximum possible efforts can achieve.

You say: "Maybe we have a different conception of what perception is. For me, perception is everything that my senses register and what can be derived from that. You might include memory as some type of 'observation' through 'inner senses'. Is that what you mean?" That all sounds OK to me. Perception is what we know of the world and ourselves. It exists physically as signals in a brain.

And: "I do not dispute that we have reference levels and that we use our perceptions to get us close to them. I just want to add something like 'negative reference levels,' things to stay away from." There are many reference settings that result in staying away from something. The simplest kind is a reference setting of zero. If you set your reference level for the perception of a loose tiger to zero, then any perception of a loose tiger constitutes an error, and you will act to reduce the perception of the tiger to zero by moving it away or yourself away from it.

And: "Freedom is a name for ranges in n-dimensional objective spare where you can move about 'at will', because the objective function is flat." You get the same result from an inverse-square function If you keep the perception of the tiger at zero, you still have all of the other degrees of freedom of movement, the only restriction (which you set yourself) being that the perception of the tiger should not depart significantly from zero. Actually, by the way, you would probably not set the reference signal to zero, but to some small nonzero amount. If there's a tiger on the loose, you want to see a very small image of a tiger, but you definitely want to see some image of the tiger. It would not be wise to lose track of where it is.

You say: "As a control systems designer, I do not create control systems in the hope that they function correctly; hope has no place in the model." Well, you hope that somebody doesn't pull the power plug, or that the motor doesn't burn out a bearing, or that the environment doesn't become so nonlinear that your design becomes unstable, and so on. Every system, however carefully designed, has failure modes, doesn't it?

In fact, designed control systems live in an environment that's almost totally predictable, so you can be pretty sure that nothing disastrously unexpected will happen. But human beings roam free through an undisciplined environment that is far more complex than any of them can understand. That environment is also full of disturbances that can't be predicted (weather, for example) or even be sensed before they occur. Most of our "predictions" are statistical in nature; sometimes they work, and sometimes they don't So there's no way that living systems could evolve to anticipate every circumstance or act correctly every time.

There's another factor that the designer has considerable control over. the forms of the analytical functions involved in the design. Most control systems are deliberately designed with linear components for the simple reason that we can't solve the equations with nonlinear functions—not because nature doesn't present us with nonlinear situations. In most real control problems, if you actually use the mathematical forms that fit the behavior of the environment most accurately, you find that you can't solve the equations and can't complete the design without trial and error. So we all use approximations; we fit a quadratic to the curve, instead of using a power of 2.113, which would fit better.

The human control systems have to work with the components that are given. They can't approximate.

My job is actually easier than yours. I'm not trying to optimize anything—just to match the behavior of a model with that of a real human subject. I'm just trying to produce a model that controls as well as people

do, not to produce engineering miracles.

Of course, real control engineers know a lot more than I do about the design of complex control systems, and some day they will take PCT much further than I possibly could. My job is not to compete with them or tell them their business. It's to get them to look at control in novel ways, ways that are not part of the customary approach—and not to improve the control systems they design, but to help us understand the behavior of organisms, most of which are not control engineers, either.

Pure reason isn't going to identify the actual variable under control by a given person in a given circumstance. A guess about what someone is controlling for could be quite right, or quite wrong. The only real way to find out is to apply a disturbance to the proposed controlled variable and see whether it's resisted in the way a control system would resist it. An even better way is to match a model to the behavior and find the parameters that give the best fit, and that predict future behavior in detail. This is why we refer to the Test for the Controlled Variable—because it provides a formal way of determining what is in fact being controlled, as opposed to what seems reasonable. People are not always reasonable. They don't all control for the same things in the same way. Sometimes they seem positively determined to do things the hard way. All we can do as theoreticians and experimenters is to find out what's really going on in a given person

You say: "I know by now what perceptual control theorists mean by the mantra "organisms control perception." As so often with jargon, it is an abbreviation for a whole philosophy and only understandable for those who have gotten to know that philosophy. It is right, from a certain perspective. From another perspective, organisms control their outputs." This isn't really jargon or "in" talk, but it is a problem with word usage. When I think of the "output" of a system, I mean the physical effect on the environment that is due to the actions of the behaving system *alone.* In the human system, this means muscle tensions, because that's that last place in the chain of outgoing effects where environmental disturbances can't get into the process and alter the consequences. Measuring the consequences any farther from the nervous system can give a false impression of what the nervous system is actually doing.

In a servo system, with this understanding of "output," I would not call the output of a motor the shaft position or speed, but the torque applied to the armature of the motor (at low speeds, anyway). Only that torque can be varied by the active system without regard to what the environment is doing. Only the torque output gives an accurate indication of the electrical output of the control system. The shaft position or speed will depend on the torque *and* on external loads and disturbances, so can't be used to indicate the output activities of the control system by itself (especially if the loads and disturbances aren't predictable).

So this is more a matter of labeling than ideology. I'm sure you would agree that a servomechanism doesn't control the torque applied to the armature of its motor, but only some consequence of that torque measured farther downstream in the causal chain. As disturbances come and go, the servo system varies its output torque, but it doesn't try to maintain any particular torque (unless torque itself is being sensed and controlled, which isn't the most common case). The torque has to be free to vary if disturbances of position or speed are to be counteracted.

The "control of perception" part is also a matter of labeling. I think you'll agree that in order to control an effect of a system's actuator output (to distinguish it from "outputs" farther along the chain), that effect must be monitored by a sensor and accurately represented as a signal. The more accurate the representation, the more accurate the control can be.

Furthermore, if the sensor characteristics change, the signal will still be brought to a match with the reference signal, but the variable it represents will no longer be maintained in the same condition. If the temperature-sensing element of a thermostat goes out of calibration, the thermostat will still think it is controlling the same temperature and will keep its movable contact nearly at the same position as before, but the room temperature will be controlled at a different level.

The only aspect of a control loop that is under reliable control, therefore, is the sensor signal. The external counterpart of that signal remains under reliable control only as long as the sensor keeps its calibration accurately. So, if we had to pin down any one aspect of the loop to be "the" controlled aspect of the situation, we would have to choose the sensor signal. Sensor signal = perceptual signal; hence, control of perception.

I think that my way of defining output and control is the least ambiguous. After all, if you define output at a place where disturbances can have an effect, you can't reason backward to the power or force output of the control system just from knowing the state of the variable called "output," because disturbances are contributing an unknown amount to the state of that variable. It seems strange to me to define output in such a way that by knowing the output you can't deduce what the control system is putting out. I don't object to looser usages for the sake of convenience, but when we want to avoid misunderstandings, I think my usage is the least ambiguous.

You say: "Perception is controlled by actions; actions are controlled by perception. Remember the loop!" Let's not confuse "control" with "affect." Control entails bringing a variable to a specified state and keeping it there. Perceptions don't bring actions to specified states and keep them there. It's the variations in the actions that bring perceptions to specified states, despite disturbances that bend to change their states. If you add a disturbance to the actuator output of a control system, the control system will alter its own output effects, not keep them the same.

In ordinary environments, the loop is asymmetrical. There is power gain going through the organism, power loss going through the environment. The part of the loop with the power gain does all of the controlling.

Hans Blom: Bill, you say: "When we study human behavior, we aren't comparing it with some 'optimal' or 'best' way of controlling. We're just trying to understand what people are actually controlling under various circumstances. In some regards, people control things very well indeed, by clever means that surpass what any engineer knows how to build. In other ways, people control stupidly and poorly, and suffer the consequences." That is not my impression. In my opinion, in the billions of years of experimentation through evolution, people (and organisms in general) have found superb ways to realize their goals. If we think that they are stupid, then we are in error. we just have not properly identified their (many!) goals. This is in line with your remark: "Much of the apparently chaotic nature of behavior becomes more understandable when we ask about higher-level goals." In my world view, an organism's behavior is perfectly in line with its top-level goals. Reaching idiosyncratic goals can, of course, be hindered by the laws of nature and of society. Every organism is always at its own local optimum. Of course, we might not agree with its definition of optimum and think that it is just plain stupid. We might even have convinced the organism of that "fact." I realize that this is a personal world view that can in no way be proven. Nevertheless, it is one of my basic life rules, until a better-working one appears. By the way, your use of "suffer the consequences" applies in any case. Behavior has unforeseeable short- and long-range side-effects, always. Our perception is limited, although training might improve things slightly.

You say: 'The highest-level goal is to win the contest, not to jump as high as possible." How do you know? The rules of the game are usually considered to be as follows: when I invent a hypothetical situation, I know what goes on in that situation, because I invented it. You go against the rules here. I say, in effect, "assume that X," and you reply "no, I cannot assume X, I assume Y." You do not play according to what I think the rules are. When I think of a reason, I can only come up with the suggestion that high jumping looks different to you than to me. Your high jumper wants to win the contest. My high jumper really wants to jump as high as possible; he is not interested in winning the contest since he already knows that he is by far the best of those he meets today. No, he is setting his sights much higher. he is training for the next Olympics. He has to compete not with his direct competitors this day, he has to compete with the figures in the world records book that he studies every day. But not even that is enough. He knows that a world record holds only for six years on average. He wants to do better than that and hold the record for many years to come. He will just give this jump his very best effort.

Are these extra perceptions helpful in seeing the situation differently? You could have been right. Your understanding might have explained somebody else's behavior. But in different persons identically looking actions can result from completely different motives. A few lines later you do seem to take that position: "You can't tell what a person is doing just by looking at what the person is doing." And later again: "Pure reason isn't going to identify the actual variable under control by a given person in a given circumstance. A guess about what someone is controlling for could be quite right, or quite wrong." Yes.

You say: "My job is actually easier than yours. I'm not trying to optimize anything—just to match the behavior of a model with that of a real human subject." I have to be precise here: our jobs are very similar. You are trying to optimize something: you are trying to find an optional match between a model and a real human subject.

You say: "Of course, real control engineers know a lot more than I do about the design of complex control systems..." Maybe, maybe not. Anyway, that extra knowledge might not account for much when it comes down to designing good control systems. After all, there is not much good theory around to travel by. "Feeling" and "intuition" are required as substitutes for knowledge. I don't think you lack those.

The question of "control" versus "affect" seems to have to do with either intended versus unintended or full versus partial correlation. In either case, it has to do with our limited predictive powers. The first raises the question of what it means to "intend" or to have "goals." The second raises the problem that actions will always have effects in addition to those "intended. " My point is that the human perceptual and conceptual systems are so beautifully designed that they even extract information from very "noisy" perceptions. Control must always be limited; the world is just too complex for our three pounds of brains to model it and our 50 pounds or so of muscles to subdue it.

In engineering, we take great liberty in defining inputs, outputs, and systems. I can take for an input anything that I can manipulate, and for an output anything that I can measure. A system is anything in between. One person's choice might differ from another one's.

Bill Powers: Hans, I don't think many evolutionists would agree with

your statement that "in the billions of years of experimentation through evolution, people (and organisms in general) have found superb ways to realize their goals." Evolution doesn't optimize anything; it just weeds out unworkable organisms. What's left is just barely good enough to survive—for a while.

I would have to agree with your implication that organisms control as well as they can. That's a matter of definition. But in looking at the state of our world, I am not greatly impressed with the way people control for social harmony, economic viability, or maintenance of an environment fit to live in.

You say: "In my world view, an organism's behavior is perfectly in line with its top-level goals." I think you're defining top-level goals from outside of the organism. When I speak of goal-seeking, I'm not normally dunking of "goals" like maintaining the life-support system and combating invasive microorganisms, or even "surviving"—the unlearned goals that I assume to drive reorganization. I'm thinking more in terms of the learned goals, things like being a good person, making a decent living, and so forth. I don't think that people are particularly adept at constructing systems of goals that hang together, are consistent with each other. Most of the people in the world live in poverty, hunger, and illness. I don't see how you can claim that they are optimal control systems.

In offering alternatives to the highest-level goal that you suggested (jumping as high as possible), I wasn't denying that some people might actually have the goal of jumping as high as possible. I was only pointing out that other goals are equally plausible, and, in my experience, more common (particularly when you ask what the *immediate* goal is). In explaining to me that in different persons identical actions might come from different motives, you're simply echoing my point.

You say: "You *are* trying to optimize something: you are trying to find an optimal match between a model and a real human subject." You're a pretty slippery customer. What you say is true: I'm controlling for the best fit between the model and the real behavior. Achieving this requires the same sort of trial and error that tuning a radio or focusing a lens requires, because the amount of error doesn't tell you which way to move, and there's no a priori way to specify the magnitude of the effect at the minimum (or maximum). This sort of control does happen. It's not very common. And it's not very tight.

Same subject: "My point is that the human perceptual and conceptual systems are so beautifully designed that they even extract information from very 'noisy' perceptions." They do that only as well as the statistics and the accuracy-time tradeoff permit. I don't worry much about extracting signal from noise; most of the behaviors we observe work at signal levels where noise can be neglected.

Then you say: “Control must always be limited; the world is just boo complex for our three pounds of brains to model it and our 50 pounds or so of muscles to subdue it.” Well, I won’t be nasty and remind you of how wonderful our evolved control systems are supposed to be. What’s really wrong with your statement is the implication that it’s hard to find instances of good control. Control is, to be sure, limited—but it’s hard to find examples of behavior in which control isn’t pretty good by anyone’s standards. “Limited” is one of those qualitative terms; the importance of the limits depends on quantitative definitions. Human motor behavior works with a bandwidth of only about 25 Hz—certainly too limited to enable us to balance on end a stick one inch long. On the other hand, this bandwidth seems to be just sufficient to handle most of the disturbances that actually occur on scales that matter to us. On those scales, the limitations are irrelevant.

You say: “In engineering, we take great liberty in defining inputs, outputs, and systems.” I think this is one of the reasons that engineers failed to come up with PCT. When you’re focused on producing some outcome in the environment, there’s no organizing principle for laying out the control system. You can put your stabilizing filters in the input function or add little loops anywhere you like that will do the job. The result is that there are no real principles of design in control engineering (that I know of). There are plenty of principles, but none having to do with how to design the functions of a control system in some systematic way. Basically, you kludge up a design that looks as if it will work, then buckle down to analyzing what you designed.

The PCT approach is to define the problem in terms of sensed variables: it is the sensed variable that will ultimately be controlled, so it should represent something specific in the environment to be controlled. The engineer can violate this principle, because the engineer knows what is to be controlled. But if the control system is in an organism, its perceptions have to be useful in a variety of higher-level systems, and they can’t have haphazard relationships to the outside world. This forces the modeler to propose a consistent set of definitions of input, output, system, and environment.

I think that a little more systematicity would also help control engineers, but that’s their business.

Hans Blom: Bill, you confuse “optimal” (an engineering word with an exact meaning) with “good” (a moral categorization) in both of these remarks: “I would have to agree with your implication that organisms control as well as they can. That’s a matter of definition. But in looking at the state of our world, I am not greatly impressed with the way people control for social harmony, economic viability, or maintenance of an environment fit to live in.” “Most of the people in the world live in

poverty, hunger, and illness. I don’t see how you can claim that they are optimal control systems.” The “optimal” of engineering means only that some system reaches its grand overall goal as closely as possible, by definition. Engineering is not concerned with the question of whether that goal is “good.” Engineers are, though. In my own personal, idiosyncratic world model, I tend to equate “optimal” with “good” (subjectively, for that person, given his / her opportunities, limitations, and life plan). Maybe that provoked your remarks.

You go on to say: “I don’t think that people are particularly adept at constructing systems of goals that hang together, are consistent with each other.” In optimal control theory, there is only one “supergoal” that can be controlled. There can be subgoals, however. It would be possible to declare the two (seemingly conflicting) goals “drive in the middle of the road” and also “drive one yard to the right of the middle.” But then you would have to combine them into one goal. This can be done, for instance, by stating that the first goal is twice as important as the second goal, or that the first goal is 100% important during the first leg of the journey and 0% thereafter. No conflicts here. Again, I think, “conflict” is a uniquely human word with a moral implication.

I had remarked: “My point is that the human perceptual and conceptual systems are so beautifully designed that they even extract information from very ‘noisy’ perceptions.” You commented: “They do that only as well as the statistics and the accuracy-time tradeoff permit I don’t worry much about extracting signal from noise; most of the behaviors we observe work at signal levels where noise can be neglected. “ This is certainly true in the domain of muscle control. But is it also true in the other domains which concern you like “being a good person; “ “making a decent living,” and so forth?

You say: “The PCT approach is to define the problem in terms of sensed variables: it is the sensed variable that will ultimately be controlled, so it should represent something specific in the environment to be controlled.” Modern control theory thinks differently. It is, of course, the sensed variables that are our only source of information about how our actions affect the objects that we want to control. But the control problem is not necessarily to bring some variables to some prescribed values and keep them there. That is, of curse, a legitimate field for study, but control theory is far broader. By the way, I think that your use of the notion “reference level” confuses some psychologists and their ilk into having to think about “homeostasis. “ Recognition of this confusion might make the PCT approach more acceptable to journal editors and referees.

Bill Powers: Hans, if you’re trying to wrap up an entire organism as a single hypercomplex control system, I suppose you would have to look for some grand overall system and a single overall purpose. That isn't the approach in hierarchical perceptual control theory. There might be many highest-level control systems acting in parallel, with relative independence. Of course there is an overall control system in the HPCT model, too, a reorganizing system, but it isn't concerned with learned behavior. Its reference levels and perceptual signals are built-in, and its mode of action is to reorganize the rest of the system. It isn't really a single entity, but a collection of control systems concerned with maintaining the life support systems, each one being concerned with a specific variable.

You say that I "confuse 'optimal' (an engineering word with an exact meaning) with 'good' (a moral categorization)... The 'optimal' of engineering means only that some system reaches its grand overall goal as closely as possible, by definition." I'm sort of between these meanings. If there are two control systems with incompatible goals inside the organism, clearly they are going to expend a lot of energy canceling each other's efforts. This is suboptimal under certain assumptions: that energy expenditure is probably a cost to the whole organism and that reduction of the control range resulting from conflict reduces the ability of both control systems to counteract disturbances. These losses of ability aren't "morally" bad, but the organism would be able to control over a wider range and for a longer time if they were not present. Of course, given the conflict, the control systems are in fact coming as dose as possible to reaching their goals. But with a suitable adjustment of the system organization, they could come a lot closer. A great deal of psychotherapy is aimed at helping people resolve conflicts; I suppose you could say that helping them is a moral choice, but it does have engineering overtones.

"In optimal control theory, there is only one 'supergoal' that can be controlled." Can you explain why this has to be true? What if there is more than one control system operating at the highest level of organization? Of course, you could make up some "supergoal" having to do with an optimal balance between these systems, but in that case the criterion of optimality would be in the eye of the beholder—there would be nothing in the system itself trying to achieve that optimality.

I think that one of the legacies of traditional psychology is a general impression that human behavior is complex and chaotic, with regularities appearing only as statistical averages, and with the future being a matter of rather shaky predictions. PCT, once you get used to seeing the things it calls to attention, shows a very different picture. Most behavior is highly regular and closely controlled; there is very little left to chance.

If this were not true, the world we experience would be very different. People would keep getting lost on the way to work; buildings and

houses would constantly be falling down, or fail to have doors or windows, or be located in inaccessible places. Cars, if they ran at all, would always be crashing into each other or wandering off across fields. Nobody would know how to grow crops, or harvest them, or transport the food to some regular destination, or how to cook the food or keep it from spoiling. Most of the things that we use, encounter, or rely upon wouldn't even exist.

What astounds me is the way in which psychologists could have looked at the endless regularities of human existence, mostly maintained by and completely products of human efforts, and failed to recognize them. It is terribly naive just to take the world the way you find it without asking how it could possibly be that way. Psychology has focused on unusual side-effects, on tiny irregularities, and has failed to see the massive regularity that characterizes all living systems and the environments they have shaped to fit their wants. The signal-to-noise ratio in most aspects if life is very, very high. That has not prevented scientists from concentrating on the noise and ignoring the signal.

"It is, of course, the sensed variables that are our only source of information about how our actions affect the objects that we want to control. But the control problem is not necessarily to bring some variables to some prescribed values and keep them there." No, I have never said it was. PCT leads to HPCT, in which higher levels of control act by varying the reference signals for lower systems. They do so as their way of controlling derived perceptions, more generalized perceptions. Those systems, in turn, have their reference levels adjusted by still higher systems, concerned with still more abstract perceptual variables. The only dissonance between this view and your ideas of optimal control has to do with your assertion that at some level there is a single highest control system with a single highest goal.

As to your criteria of optimality, they are completely discretionary. I don't see any reason to suppose that organisms have adopted such criteria or seek to realize them. You're talking about engineers building control systems, not the processes by which living control systems evolve. The engineer can, by choice, combine all lower goals into supergoals, but there is nothing that compels us to suppose that organisms do the same thing—except when they're trained as engineers.

All that the brain knows about the external world comes to it in the form of perceptual signals in the afferent neural pathways. There is no other way for that information to get into the brain. If the brain wants to control the position of a real glass on a real table, it's out of luck: it doesn't have any way to know about the real glass and the real table. It can, however, adjust its output signals so that a neural signal representing the glass can be manipulated to achieve a certain relationship with a neural signal representing the table. *That* the brain *can* do.

I should think that all of this would be self-evident to any engineer who has ever actually built a working control system. A real hardware control system can't interact directly with the physical plant it is controlling. All it can do is alter its electronic output signals and see what happens to the signals being generated by its sensors. That's all it knows about what is happening outside it. If the sensors jump out of calibration, the control system will happily continue controlling the miscalibrated perception, while the technician in charge rushes to hit the *Stop* button What is controlled is *only* what is perceived. One hopes that what is perceived has some relationship to what is, but that is something that has to be determined indirectly.

This simple concept which should cause no problems for any control engineer causes immense problems for conventional sciences of life. The reason is that these conventional sciences ignore the difference between what is perceived and what is—at least when they're trying to explain behavior. And not having any experience with real system design, it seems perfectly reasonable to such conventional scientists that a stimulus input from real objects in the environment should be able to cause motor outputs that steer the organism through a variable environment along a path to the cheese or the mate or whatever. What's the problem? You can see them doing it, so it must be easy.

If you're an engineer watching an organism behave, you will have a hard time making your mental model behave in this simple cause-effect way. You will notice that the eyes keep moving around, that the head moves and bobs up and down, that the steps are a little imprecise and slightly wobbly, that things in the environment are shifting around. Being a person who is charged with making systems actually work, you will wonder how the organism gets away with such imprecision of action—where are all the stimuli coming from that cause the corrections of the little mistakes and overshoots and hesitations? How does the environment know that it should stimulate the organism just in the right way to correct for a previous stumble? How does that little unevenness in the path send just the right stimulus up the spine to make just the right muscles change their tension to keep the leg from jamming into the ground or flailing in empty air on the next step? Any engineer who pays attention in a professional way to the claims of S-R theorists would soon walk away shaking his or her head. No way!

Unfortunately, engineers seem to abandon their normal professional attitudes when they start trying to explain behavior. They start listening to the psychologists and physiologists and neurologists who think that behavior can just be "generated," open-loop. Perhaps they're just being polite because they're on another scientist's turf. They say, "Oh, is that how it works? OK, you must know what you're talking about; I'll see if I can make that work." And, of course, they can make it work. Good engineers can make any damn fool idea work. They can build an arm that's as solid as the front end of a Mack truck, equip it with precision bearings and gears and stepper motors, compute the driving signals using 80-bit floating point arithmetic, and make the arm move exactly as wanted. The smart ones must surely realize that this is *nothing* like the way a human arm works. But the psychologists see what they've done, and nod wisely. It works just the way they expected.

PCT is all about the realization that human systems simply can't work that way. Their outputs are rubbery and imprecise; their neural computers are good to maybe 1% at best; they don't sense everything in the environment that might interfere with the action Yet they work precisely and well, for four score years and six. A person with his little 1% analog computers can get out of bed in the morning and perform one action after another all day long, each action starting where the last one left off, and 16 hours later end up exactly at the side of the same bed, with no cumulative errors at all. Only one kind of system can accomplish that sort of behavior: a negative-feedback control system.

Hans Blom: I enjoy reading / scanning CSGnet a lot; I have derived many eurekas from it (not in the sense of discovering new "truths," but in the sense of gaining new perspectives), and I have come to respect Bill Powers' view of reality. The following remarks are probably more meta-science than science. But many of these discussions are, aren't they?

In systems science, we have the notion that any model accomplishes a particular end. You develop a model with a certain goal in mind; the goals might be different for different modelers. Models can be viewed as theories: you want to summarize all findings within a limited scientific domain in a certain form, e.g., a block diagram. Models can be viewed as tools: you want to encapsulate all properties of a system that you deem important into a simplified form, so that you can control the important aspects of an otherwise too-complex reality. Models can be viewed as predictors or extrapolators: if something happened in the past, it might happen again in a similar way. In *all* cases, we have to understand that each and every model is a simplification of reality, in which we leave out those aspects that we deem unimportant. Therefore, each model is a personal choice: what is unimportant to you might be the most important thing in the world to another person. Or, as the saying goes amongst control engineers: one person's noise is another person's signal.

Of course, such a personal choice might be picked up by others and become part of culture—but only if those others agree with how you split the world into "important" and "unimportant." Sometimes, agreeing is easy: color does not contribute to a body's mass. In other cases, it's not that easy: do people have free will? You might protest that "free will" is a badly defined notion. That is true. But so are "color" and "mass"; no two people or measuring devices will perceive exactly the same color or mass. You might complain again and say that the mass that two well-calibrated scales measure when exposed to the same body is *practically* the same. But that depends upon the practice at hand. In real life, we frequently (always?) seem to have to deal with fuzzy notions. In many cases, this fuzziness does not matter, but in others it can matter a great deal.

We each have a personal, emotional investment in our models. They encapsulate what *we* think is important and leave out what we believe is unimportant. Models are personal creations, much like works of art, that we experience as the best that we can produce. On this net, Bill defends what he sees as important. Of course. But so does everyone else. Isn't that one of the central tenants of your theory, Bill?

This brings me to the issue that, in my opinion, is expressed too little in PCT philosophy. Control is about *control*. You focus on *perceptions* as the important things—and, as a concomitant, on which perceptions are controlled. I have a different ordering of things important. Prime is that we have *goals* (reference levels, as you call them); a control system is a device that allows us to reach or approach those goals in the best possible ways, given our biological and mental limitations. This is also the orthodox control engineering vision. You have a goal, so go design a system that makes it come true. Use the information that the available sensors provide in the best possible way, using any type of processing and data storage that is available or can be newly designed. Control engineers do it this way, and evolution as well, I think. In control engineering, theory has its part; it provides a number of well-proven (partial) solutions. Hunches, trial and error, too, have their parts. No new design is exactly the same as a previous one, alas.

Does this difference in focus matter, you might ask? Yes, I think so. In science, it seems as if we have left all "grand unified theories" behind—although physics is still searching. It seems as if there are no "first principles"; you can go deeper and deeper all the time, if you have the resources. First principles seem to be theories as well. And they are practically useless to explain the world in all its complexity. The formulas of quantum mechanics are barely able to "explain" the movement of *one* electron around *one* proton (the simplest atom that exists), but anything more complex is beyond its powers of synthesis. The synthesis problem is much older, of course: the classical three-body problem of classical mechanics does not allow precise long-term predictions. We are now mentally just coming to grips with these strange facts: that even if first principles are given, a synthesis based on those first principles might

be too complex computationally (and mentally) to derive higher-order laws and "explain" more complex systems. That's what chaos theory is all about. Ask any practical control engineer: the existing theories do not suffice when you design a new control system. Always, some extra creativity is required. It is not that those theories are useless; they are not sufficient. Ask any AI-type who works with expert systems: it is not the "reasoning process" that provides the performance of a knowledge-based system, but the knowledge incorporated into it; the more, the better. But then we start to encounter the „complexity problem": a system with a large number of basically independent "knowledge chunks" starts to show unpredictable and uncomprehensible behavior because of the unforeseen ways in which those chunks (sometimes) interact. The result is that the paper model cannot explain or predict anymore. You actually have to *build* it and *run* it to see how it behaves. Philosophers who study culture start to recognize *the same* thing: post-modernists say that the time of the "grand stories," of the ideologies, is over. It is the 'little stories," the personal, subjective accounts, that are the important things that build up the world (and, if generally accepted, might grow into "grand stories").

In my view, no model is wrong, unless it is internally inconsistent. Of course, *any* model is wrong in the sense that it must necessarily be incomplete. In another sense, a different model might be right *as well:* it just has a different purpose (focus) and is based on different notions of what is important. This is true for all models, even PCT models—unless you talk in abstractions that can neither be proven or disproven. It follows from the basic notion that *every model is an approximation*.

If you can accept that different models reflect different goals and therefore incorporate and / or explain different observations, what is a fact to one modeler can be noise to another. A concomitant of this is that a model is (approximately) valid only within some restricted domain. It might "explain" a certain set of observations, but it is without value, or simply wrong, outside its domain. Einstein's E = mc2 certainly does not relate someone's "psychic energy" to his or her body weight.

Don't underestimate statistics. Astronomical data that remain from the days of Kepler show small and large measurement errors. Newton's laws could never have been derived without discarding quite a lot of outliers and assuming that the theory need not *exactly* fit the measurements. Yet, Newton's laws have shown their value. But they, too, are approximations, as Einstein showed. And, undoubtedly, Einstein's relativity theory is an approximation as well.

Bill's hierarchical control model consists of a multitude of simple, functionally identical blocks. The model is an elegant simplification, but we know that the brain isn't quite that homogeneous, neither at the cell level nor at the level of configurations of cells (wiring). Bill, you can marvel at the beauty of your model (it *is* elegant!), yet acknowledge that even in its very basics it cannot possibly be correct.

But that often does not matter much. One system can be modelled in a great many different ways, yet these models can *functionally* show (approximately) the same behavior. This I consider a basic conflict in your model: on the one hand, you want your model to represent physiology as accurately as possible; on the other hand, you want it to show the same *functional* behavior as a human. We are, I think, still very far from the point where we can link the lowest levels (cells, synapses) with the highest levels. In my opinion, and based on the arguments that I presented above, establishing such a link might be impossible in theory, as well.

As has often been noted on this net, things that "actually exist in nature" will forever remain outside our grasp. The best thing we can do is build *models* of what is out there. You know this, Bill, yet it seems that you cannot really accept it. What we require of a model is (a) that it is internally consistent, and (b) that it is consistent with our observations of the "real world." The problem lies in the latter, where we encounter the limitations. We cannot take into account every observational detail. We have to select. And *how* we select depends upon both what we deem important and what we have as capabilities, i.e., we make a personal choice based on our personal goals but within our personal limitations when we build our model. I strive for what I want, building upon what I already know. This is true in mathematics, in control engineering, in life.

Model or theory building is basically a creative process, in which you suddenly have this eureka-feeling of "yes, that's it!" But then science expects you to "prove" your model or theory, and you suddenly find that the theory does not explain all of the data or does not explain with full accuracy. That is when we have to introduce notions like "noise" (small discrepancies that we choose to disregard), "outliers" (large discrepancies that we choose to disregard), "statistics" (can I get an impression of how well my new theory fits the observations despite the fact that I disregard so much?) and things like that. Finally, a theory might start to lead its own life and be taken more seriously than the data. Bill, I assume that you, too, take Newton's laws more seriously than the data that they were originally based on, and more seriously as well than a great deal of more recent measurements.

All of the notions that you use are high-level abstractions, much like "force," "pressure," and "temperature," which have no objective existence but are cultural notions, ways of looking at what surrounds us. In every case, philosophers will tell you, we could have arrived at different but equally valid notions. To use a simple example: you use feet and Fahrenheit, while I use meters and Celsius.

As Rick Marken can tell me so eloquently: "It's all perception." Translate this into: "It's all your own personal subjective theory / model of what's out there," and you are dose to what I want to say.

As you can see, Bill, my 'life model" is, in many ways, different from yours. Why? Our models are based on different data, on different perceptions of what is important, and on different goals. My model has been built up through my experiences that have gradually taught me (a) how to perceive (what to notice, what to disregard), (b) which goals to set (the things that I have come to consider important) and (c) how to act (through the goal-reaching skills that have worked for me).

Everybody has one goal in common, however personal that goal looks: to make the world more controllable / understandable. Every trick in the book—as well as every new one that you can think of—is used to reach that goal. One trick is to observe others and see how they control; maybe (who knows?) their methods will work for me, too. Let's be inclusive, not exclusive. Let's find the best tricks and use all of them combined in our personal repertoire. Please take this contribution in that vein. As you might have noticed, I take your "life model" seriously. It provides a much needed additional perceptive. Yet, allow me to think that I, given different perceptions, might have discovered a "life model" that might have some value as well, even if it does not coincide with yours. In works of art, I often find it difficult to say which painting or sculpture is "better" than another. I am slowly discovering that I have a similar problem with scientific theories.

Bill Powers: Hans, you say: "In systems science, we have the notion that any model accomplishes a particular end." Yes, in the sense that any model that actually works does *something.* But there are two kinds of ends-achievement going on in PCT modeling. One is the modeler's goal of constructing a model that behaves like the real system. The other is the model's goal of bringing some perceptual representation of its environment to a reference-state endogenous to the model. If you construct a food-seeking model that depends on balancing smell intensities in a bug's antennae, but get the sign of the perceptual computation wrong (a - b instead of b - a), the bug will seek a goal, all right, but it will be the goal of traveling away from the food. So the model, while achieving its own goal, will not achieve the modeler's goal. The modeler wants the bug to want to get near the food and so will reverse that sign, altering what the bug-model perceives to make the outcome the same as what the modeler wants.

And you say: "You develop a model with a certain goal in mind; the goals might be different for different modelers." What I see missing in systems science is the concept of systems that have their *own* goals. That is, the system is designed to accomplish what the modeler wants done, but the idea that the system itself might want something doesn't seem to be addressed. Am I wrong about that? I admit that artificial devices aren't asked very often what they want, nor does it matter, but when we're modeling the modeler, we have to put the goals into the model.

You say: "Models can be viewed as tools: you want to encapsulate all properties of a system that you deem important into a simplified form, so that you can control the important aspects of an otherwise too-complex reality." The question remains, who does the controlling toward whose ends? Your statement seems to imply that the model's behavior is there only to satisfy the modeller's goals. This says that the model is not a model of the modeler, but of some device to be used for achieving the modeler's purposes. How, then, do we model the modeler, whose goals aren't being given by some other person to suit that other person?

You also say: "Models can be viewed as predictors or extrapolators: if something happened in the past, it might happen again in a similar way." I can agree to this in a very broad sense, but I wonder if it's the same sense you mean. Models in PCT aren't designed to produce particular behaviors under circumstances that led to those behaviors in the past. The *components* of these models could be seen that way—if a comparator has always produced a certain error signal given a particular reference and perceptual signal, we expect it to go on behaving that way. This is what we mean when we describe each function box with a mathematical form. We observe or propose that this function has been performed by that box in the past, and we predict that it will continue to perform the function.

A control-system model can be designed, on the other hand, to produce behavior like that of the real system, quite accurately, in the presence of conditions that have never occurred before. We can measure the control parameters for simple pursuit tracking, for example, and predict how a teal person will perform in a new task with a new pattern of movements of the target, and *with a second disturbance applied directly to the cursor*, which was not present when the parameters were evaluated. Now the model is presented with new conditions (as is the human subject), and the model still behaves just like the subject. This is not exactly extrapolating from past performance, is it? At least it's a kind of extrapolation that is very different from just observing disturbances and the behaviors that follow them, and predicting that recurrence of the same disturbances will produce the same behavior.

"In *all* cases, we have to understand that each and every model is a simplification of reality, in which we leave out those aspects that we deem unimportant." Yes, indeed. The trick is to know when you're leaving out or simplifying something vital. You find this out when you match

the model's behavior to the real behavior, or when you change conditions in a way that brings the omitted parts into play. But this is the whole modeling game, isn't it? You get the model to work in as simple a form as possible, then change the conditions until the model stops behaving like the real system. The way in which it fails can sometimes be traced to simplifications or omissions, in which case you go back and use a more detailed model. Other times, the model fails completely, and you have to reconsider it from scratch. The PCT models we use in tracking experiments today represent a long history of wrong guesses, although they're still so simple that it might seem impossible that they were overlooked in the beginning.

"Therefore, each model is a personal choice: what is unimportant to you may be the most important thing in the world to another person." In principle, maybe. In practice, it doesn't feel that way. Some models just don't work no chatter how hard you try to make them work. I suppose you could invoke psychoanalysis and say that if a model fails, its inventor really didn't *want* it to work. But it's hard to believe that when you can see a model designed exactly as you wanted it to be designed that behaves in a way completely different from the real behavior you thought you were modeling. No matter how much you like the model, no matter how many of your private beliefs or prejudices it expresses, if it doesn't work, it doesn't work, and there's no way but self-delusion to make it seem to work.

While I don't think that any models are the last true words about how nature works, I think that some models are definitely better than others. This isn't self-evident if you just construct conceptual models and never test them experimentally. It isn't self-evident if the models are simply descriptions of observations (there are countless ways of describing the same observations). The relative worth of models can be seen only when they're expressed as working simulations that can generate behavior out of their own properties. When you've committed yourself to the point of constructing a working model, there is no way you can make the model work other than the way you designed it to work—and if the way it works doesn't resemble the way the real system you're modeling behaves, you've just shown that your model is wrong.

"Of course, such a personal choice may be picked up by others and become part of culture. But only if those others agree with how you split the world into 'important' and 'unimportant.'" This is a different subject: not which model is best, but what aspect of experience you want to model. In PCT we generally agree that we want to model ordinary behavior: what people do in daily life, at many levels. We're not trying to model chakras or satori or survival after death or ghosts or metabolism or lots of things like that. Just plain vanilla behavior. Generally we took at the same things that other theories have looked at: environmental events near organisms, actions and their consequences produced by the muscles of organisms, perceptions of various kinds, nervous systems and their possible functions. We aren't emphasizing or de-emphasizing any of these phenomena; we're just asking what makes them work the way they seem to work.

"You might protest that 'free will' is a badly defined notion. That is true. But so are 'color' and 'mass'; no two people or measuring devices will perceive exactly the same color or mass." That's a bit qualitative for a valid comparison. We can characterize color and mass well enough to reproduce them within a few parts per thousand and agree on perceptions of them within a few percent, but I defy anyone to reproduce "free will" in any way that can be quantified. No two people perceive color or mass *exactly* the same, but no two people perceive free will even *approximately* the same: many claim they don't even perceive it. Let's at least compare apples with round things.

"In real life, we frequently (always?) seem to have to deal with fuzzy notions. In many cases, this fuzziness does not matter, but in others it can matter a great deal." The quality of our lives is vitally affected by fuzzy notions we would be better off without, or at least with, but in sharper form. The point of science, in my mind, is to clarify fuzzy notions or to get rid of them if they are intractably blurred.

"We each have a personal, emotional investment in our models. They encapsulate what tae think is important and leave out what we believe is unimportant." I have some investment in a model of tracking behavior in which the model's simulated handle position follows a course through time that deviates from the handle position created by a person in the same experiment only three to five percent, RMS. I think it is important for the behavior of a model to be as close to the behavior it supposedly models as possible. I'd like it to be closer, but so far can't accomplish that Someone else might consider this sort of match unimportant, preferring ire cream or skiing. Someone else might think that tracking behavior isn't very interesting, considering the problems in Somalia. But anyone who thinks that models of overt physical behavior should reproduce and predict behavior accurately has this model to contend with.

I doubt that the behavior of this model has much to do with my personal emotional investments.

"Models are personal creations, much like works of art, that we experience as the best that we can produce." There's a bit more than that to models that I respect. A model should deal with data that's publicly observable by means on which we can agree and reproduce independently. The reasoning that leads to the model should be laid out in public view in sufficient detail that anyone who understands basic logic

and mathematics could recreate the model from scratch if necessary, and come up with the same model. The model should behave the same way in anyone's hands and should fit behavior correctly as evaluated by any user of the model. I don't think that very many of these considerations apply to works of art.

"On this net, Bill defends what he sees as important. Of course. But so does everyone else. Isn't that one of the central tenants of your theory?" Certainly, and I'm glad that you see the theory as correctly describing human behavior.

'This brings me to the issue that, in my opinion, is expressed too little in PCT philosophy. Control is about *control*. You focus on *perceptions* as the important things—and, as a concomitant, on which perceptions are controlled." It would be pretty hard to focus on perceptions as the important things without the rest of the control loop. Perceptions aren't just sort of vaguely "important." It just happens that when you try to find the variable in a control loop that is the most reliably controlled under the most changes of conditions, it proves to be the perceptual signal. We didn't pick perceptions as pivotal for private or silly reasons, or just because we're perception freaks. Perceptions are all that an organism can know about the world outside it. That means you, too. It follows that goals have to be defined in terms of perceptions. You can't compare an internal goal with an external unperceived object; the object must appear as a perception in the same place where the goal is before any comparison can take place. PCT is about goals, too, and about error signals and input functions and actuators and all of the parts of a control system.

"In my view, no model is wrong, unless it is internally inconsistent." I guess our views differ. I demand that a model behave *like* the world it is supposed to describe or explain. A model can be internally consistent yet totally at variance with experimental observations. What is "important" has nothing *to* do with this. If a model predicts something unimportant incorrectly, it is still wrong. Models that don't have *anything to* do with observation and that produce no predictions of behavior *to* be compared with observation don't even count as models in my world. There's no reason to take them seriously unless the math grabs you.

"Don't underestimate statistics. Astronomical data that remain from the days of Kepler show small and large measurement errors. Newton's laws could never have been derived without discarding quite a lot of outliers and assuming that the theory need not exactly fit the measurements." "Measurement error" is something very different, quantitatively, from "variance" in psychological observations. You can measure a rat's running speed in a maze with a measurement error of perhaps 0.1 percent, if you use instrumentation. But the supposed effects of stimulus conditions on that running speed will have a variance of hundreds to thousands of percent. Newton and Kepler were trying to formulate models of celestial mechanics that would predict the positions of planets within the existing measurement errors. If the kinds of statistical methods used in psychology had been brought to bear on this problem, celestial mechanics would consist of the firm statement that the planets are up there, not down here. It is very hard to underestimate the power of statistics as used in the behavioral sciences.

"Model or theory building is basically a creative process, in which you suddenly have this eureka-feeling of 'yes, that's it!' But then science expects you to 'prove' your model or theory, and you suddenly find that the theory does not explain all the data or does not explain with full accuracy. That is where we have to introduce notions like 'noise' (small discrepancies that we choose to disregard), 'outliers' (large discrepancies that we choose to disregard), 'statistics' (can I get an impression of how well my new theory fits the observations despite the fact that I disregard so much?) and things like that." This is a rather remarkable statement, in that in summarizes exactly what I think is wrong in the behavioral sciences. Concepts like noise, outliers, statistics, variance, and so forth were invoked by psychologists as a way of explaining why their theories of behavior didn't predict worth a damn. Instead of blaming the poor results on a mismatch of theory to the organism, they blamed it on the organism. In PCT, any time we get results like the *best* statistical results in conventional behavioral experiments, we look for what is wrong with the model. And we usually find it. Behavior, I strongly suspect with some smattering of data in support, is nowhere near as variable as it has seemed to psychologists viewing it through their theories.

"All of the notions that you use are high-level abstractions, much like 'force; 'pressure; and 'temperature; which have no objective existence but are cultural notions, ways of looking at what surrounds us. In every case, philosophers will tell you, we could have arrived at different but equally valid notions." True, but high-level abstractions are grounded in lower-level ones, down to the level normally accepted in science as "observational"—the level where you can report just how much. How much of *what is* determined theoretically, but the relationships among observations are predicted at a low level of abstraction: how far one trace on a chart deviates from another.

As to the philosophers, it's easy to say that you *could* arrive at a different but equally valid notion. Actually doing that is a bit harder. What I hope for is a model for which *nobody* can think of an *equally valid* alternative. The fact that one might hypothetically exist doesn't bother me much. I'm concerned with the model we do have today, not one that might show up later.

The Hierarchical Behavior of Perception

Richard S. Marken
(Life Learning Associates, 10459 Holman Ave., Los Angeles, CA 90024)

Abstract

This paper argues that the coincidental development of hierarchical models of perception and behavior is not a coincidence. Perception and behavior are two sides of the same phenomenon—control. A hierarchical-control-system model shows that evidence of hierarchical organization in behavior is also evidence of hierarchical organization in perception. Studies of the temporal limitations of behavior, for example, are shown to be consistent with studies of temporal limitations of perception. The control model shows that the perceptual limits are the basis of the behavioral limits; action systems that are capable of rapid response cannot produce controlled behavioral results faster than the rate at which these results can be perceived. Behavioral skill turns on the ability to control a hierarchy of perceptions, not actions.

Introduction

Psychologists have developed hierarchical models of both perception (e.g., Bryan & Harter, 1899; Palmer, 1977; Povel, 1981) and behavior (e.g., Albus, 1981; Arbib, 1972; Greeno & Simon, 1974; Lashley, 1951; Martin, 1972; Rosenbaum, 1987). This could be a coincidence, a case of similar models being applied to two very different kinds of phenomena. On the other hand, it could reflect the existence of a common basis for both perception and behavior. This paper argues for the latter possibility, suggesting that perception and behavior are two sides of the same phenomenon: control (Marker, 1988). Control is the means by which agents keep perceived aspects of their external environment in goal states (Powers, 1973). It is argued that the existence of hierarchical models of both perception and behavior is a result of looking at control from two different perspectives: that of the agent doing the controlling (the actor), and that of the agent watching control (the observer). Depending on the perspective, control can be seen as a perceptual or a behavioral phenomenon.

From the actor's perspective, control is a perceptual phenomenon. The actor is controlling his or her own perceptual experience, making it behave as desired. However, from the observer's perspective, control is a behavioral phenomenon The actor appears to be controlling variable aspects of his or her behavior in relation to the environment. For example, from the perspective of a typist (the actor), typing involves the control of a dynamically changing set of kinesthetic, auditory, and, perhaps, visual perceptions. If there were no perceptions, there would be no typing. However, from the perspective of someone watching the typist (the observer), perception is irrelevant; the typist appears to be controlling the movements of his or her fingers in relation to the keys on a keyboard.

These two views of control have one thing in common; in both cases, control is seen in the behavior of perception. For the actor, control is seen in the behavior of his or her own perceptions. For the observer, control is seen in the behavior of his or her own perceptions of the actor's actions. (The observer can see the means of control but can only infer the perceptual consequences as experienced by the actor). If control is hierarchical, then it can be described as the behavior of a hierarchy of perceptions. Hierarchical models of perception and behavior can then be seen as attempts to describe control from two different perspectives, those of the actor and observer, respectively. This paper presents evidence that hierarchical models of perception and behavior reflect the hierarchical structure of control.

A Perceptual Control Hierarchy

The concept of control as the behavior of perception can be understood in the context of a hierarchical-control-system model of behavioral organization (Powers, 1973, 1989). The model is shown in Figure 1. It consists of several levels of control systems (the figure shows six levels), with many control systems at each level (the figure shows 11). Each control system consists of an input transducer (I), a comparator (C), and an output transducer (O). The input transducer converts inputs from the environment or from systems lower in the hierarchy into a perceptual signal, p. The comparator computes the difference, e, between the perceptual signal and a reference signal, r. The output transducer amplifies and converts this difference into actions which affect the environment or become reference signals for lower-level systems.

The control systems at each level of the hierarchy control perceptions of different aspects of their sensory input, but all of the systems control perceptions in the same way: by producing actions that reduce the discrepancy between actual and intended perceptions. Intended percep-

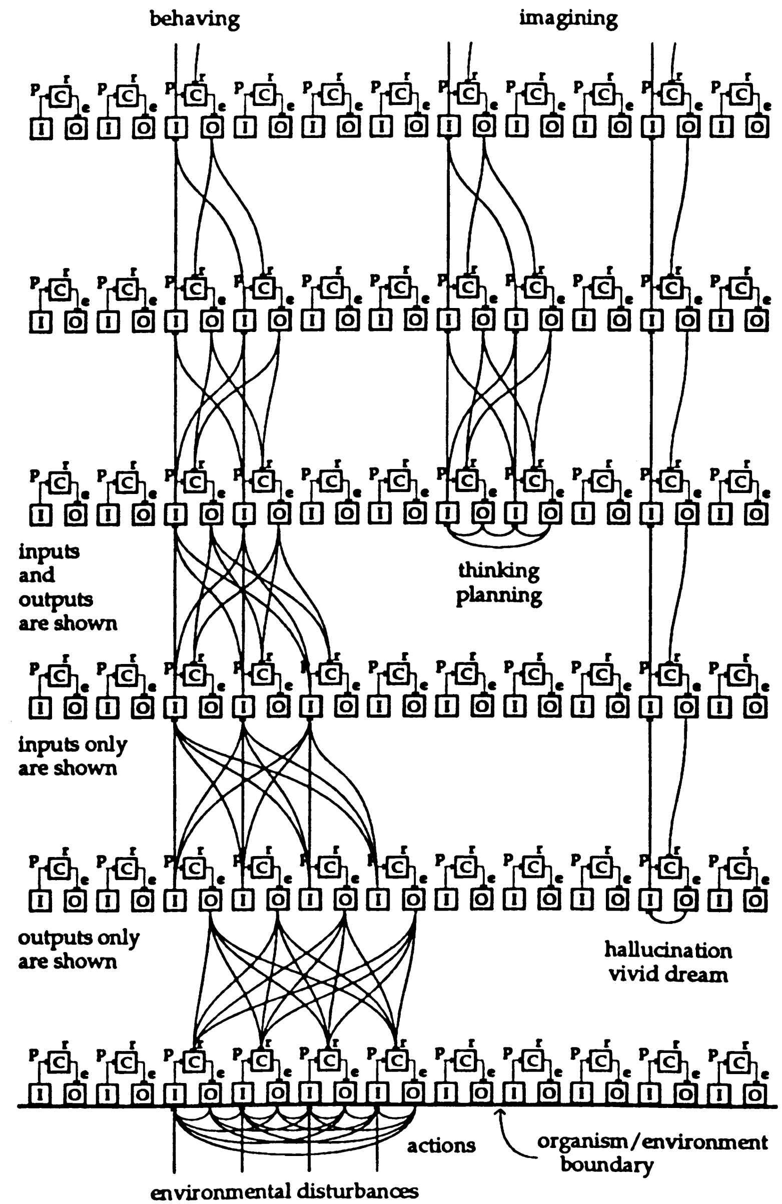

Figure 1. Perceptual control hierarchy (after Powers, 1989, p. 278).

tions are specified by the reference signals of the control systems. The actions of the control systems coax perceptual signals into a match with reference signals via direct or indirect effects on the external environment. The actions of the lowest-level control systems affect perceptions directly through the environment. The actions of higher-level control systems affect perceptions indirectly by adjusting the reference inputs to lower-level systems.

The hierarchy of control systems is a working model of purposeful behavior (Marker, 1986, 1990). The behavior of the hierarchy is purposeful, inasmuch as each control system in the hierarchy works against any opposing forces in order to produce intended results. Opposing forces come from disturbances created by the environment, as well as from interfering effects caused by the actions of other control systems. The existence of disturbances means that a control system cannot reliably produce an intended result by selecting a particular action. Actions must vary to compensate for varying disturbances. Control systems solve this problem by specifying what results are to be perceived, not how these results are to be achieved. Control systems control perceptions, not actions. When set up correctly, the control systems in the hierarchy vary their actions as necessary, compensating for unpredictable disturbances, in order to produce intended perceptions. Indeed, the term "control" refers to the process of producing intended perceptions in a disturbance-prone environment.

Levels of Perception

Powers (1990) has proposed that each level of the hierarchy of control systems controls a different class of perception. Moving up the hierarchy, these classes represent progressively more abstract aspects of sensory input. The lowest-level systems control perceptions that represent the intensity of sensory input. At the next level, the systems control sensations (such as colors), which are functions of several different intensities. Going up from sensations, there is control of configurations (combinations of sensations), transitions (temporal changes in configurations), events (sequences of changing configurations), relationships (logical, statistical, or causal co-variations among independent events), categories (class memberships), sequences (unique orderings of lower-order perceptions), programs (if-then contingencies among lower-level perceptions), principles (general rules perceptible in the behaviors of lower-level perceptions), and system concepts (particular sets of principles exemplified by the states of many lower-level perceptions; see Powers, 1989, pp. 190-208). These 11 classes of perception correspond to 11 levels of control systems in the hierarchical-control model. All control systems at a particular level of the hierarchy control the same class of perception, though each system controls a different exemplar of the class. Thus, all systems at the configuration level control configuration perceptions, but each system at that level controls a different configuration.

The rationale for hierarchical classes of perceptual control is based on the observation that certain types of perception depend on the existence of others. Higher-level perceptions depend on (and, thus, are functions of) lower-level perceptions. For example, the perception of a configuration, such as a face, depends on the existence of sensation and intensity perceptions. The fare is a function of these sensations and intensities. The lower-level perceptions are the independent variables in the function that computes the higher-level perception. Their status as independent variables is confirmed by the fact that lower-level perceptions can exist in the absence of the higher-level perceptions, but not vice versa. Sensation and intensity perceptions can exist without the perception of a fare (or any other configuration, for that matter), but there is no fare without perceptions of sensation and intensity.

The Behavior of Perceptions

From the point of view of the hierarchical-control model, "behaving" is a process of controlling perceptual experience. Any reasonably complex behavior involves the control of several levels of perception simultaneously. For example, when typing the word "hello," one controlled perception is the sequence of letters "h," "e; " l, " "l," and "o." The perception of this sequence is controlled by producing a sequence of keypress-event perceptions. Each keypress event is controlled by producing a particular set of transitions between finger-configuration perceptions. Each finger configuration is controlled by a different set of force sensations, which are themselves controlled by producing different combinations of intensities of tensions in a set of muscles.

The perceptions involved in typing "hello" are all being controlled simultaneously. Transitions between finger configurations are being controlled while the force sensations that produce the configuration perceptions are being controlled. However, the typist is usually not aware of the behavior of all these levels of perception. People ordinarily attend to the behavior of their perceptions at a high level of abstraction, ignoring the details. We attend to the fact that we are driving down the road and ignore the changing muscle tensions, arm configurations, and steering wheel movements that produce this result. Paying attention to the details leads to a deterioration of performance; it is the opposite of Zen behavior, where one attends only to the (perceptual) results that one intends to produce and lets the required lower-level perceptions take care of themselves (Herrigel, 1971). However, while it violates the principles of Zen, attention to the detailed perceptions involved in the production of behavioral results can provide interesting hints about the nature of the perceptual control hierarchy.

The Perception of Behavior

The behavior of an actor organized like the hierarchical-control model consists of changes in the values of variables in the actors environment. An observer cannot see what is going on inside the actor; he or she can only see the actor's actions and the effect of these actions on the external environment. The effect of these actions is to cause purposeful behavior of certain variables in the environment: the variables that correspond to perceptions that the actor is actually controlling. The purposefulness of the behavior of these variables is evidenced by the fact that consistent behaviors are produced in the context of randomly changing environmental disturbances. Thus, a typist can consistently type the word "hello," despite changes in the position of the fingers relative to the keyboard, variations in the push-back force of the keys, or even a shift from one keyboard arrangement to another (from QWERTY to Dvorak, for example).

Since the actor controls his or her own perceptions, the observer cannot actually see what the actor is "doing"; the acts "doings" consist of changing the intended states of his or her own perceptions. The observer sees only the variable results of the actors actions-results that might or might not be under control. For example, the observer might notice that a click occurs each time the typist presses a key. The click is a result produced by the typist, and the observer is likely to conclude that the typist is controlling the occurrence of the click. In fact, the click might be nothing more than a side-effect of the typist's efforts to make the *key* feel like it has hit bottom. There are methods that make it possible for the observer to tell whether or not his or her perceptions of the actor's behavior correspond to the perceptions that are being controlled by the actor (Marker, 1989). These methods make it possible for the observer to determine what the actor is actually doing (i.e., controlling).

Hierarchical Control

The hierarchical nature of the processes that generate behavior would not be obvious to the observer of a hierarchical control system. The observer could tell that the system is controlling many variables simultaneously, but he or she would find it difficult to demonstrate that some of these variables are being controlled in order to control others. For example, the observer could tell that a typist is controlling letter

sequences, keypress events, finger movements, and finger configurations. But the observer would have a hard time showing that these variables are hierarchically related. The observer could make up a plausible hierarchical description of these behaviors; for example, finger positions seem to be used to produce finger movements which are used to produce keypresses which are used to produce letter sequences. But finding a hierarchical description of behavior does not prove that the behavior is actually produced by a hierarchical process (Davis, 1976; Kline, 1983).

Hierarchical Invariance

Hierarchical production of behavior implies that the commands required to produce a lower-level behavior are nested within the commands required to produce a higher-level behavior. For example, the commands that produce a particular finger configuration would be nested within the commands that produce a movement from one configuration to another. Sternberg, Knoll, & Turlock (1990) refer to this nesting as an invariance property of hierarchical control. Lower-level commands are like subprograms invoked by programs of higher-level commands. The invariance of hierarchical control refers to the assumption that the course of such a subprogram does not depend on how it was invoked from the program (low-level invariance); similarly, the course of the program does not depend on the nature of the commands carried out by the subprograms (high-level invariance).

Convergent and Divergent Control

The hierarchical-control model satisfies both the low- and high-level invariance properties of hierarchical control. The commands issued by higher-level systems have no effects on the commands issued by lower-level systems, and vice versa. It is important to remember, however, that the commands in the control hierarchy are requests for input, not output. Higher-level systems tell lower-level systems what to perceive, not what to do. This aspect of control-system operation solves a problem that is either ignored or glossed over in most hierarchical models of behavior. how does a high-level command get turned into the lower-level commands producing results that satisfy the high-level command? If commands specify outputs, then the result of the same command is different when there are varying environmental disturbances. The high-level command to press a key, for example, cannot know which lower-level outputs will produce this result on different occasions. This problem is solved by the hierarchical-control model because intended results are represented as a convergent function, which produces a single perceptual signal, rather than as a divergent network, which produces multiple behavioral outputs.

Most hierarchical models of behavior require that a high-level command be decomposed into many lower-level commands to produce an intended result. In the hierarchical-control model, both the high-level command and the intended result of the command are represented by a single, unidimensional signal. The signal that represents the intended result is a function of results produced by many lower-level commands. But the high-level command does not need to be decomposed into all of the appropriate lower-level commands (Powers, 1979). The difference between the high-level command and the perceptual result of that command is sufficient to produce the lower-level commands that keep the perceptual result at the commanded value (Marken, 1990).

Levels of Behavior

The hierarchical invariance properties of the control hierarchy provide a basis for determining whether its behavior is actually generated by hierarchical processes. Hierarchical control can be seen in the relative timing of control actions. In a control hierarchy, lower-level systems must operate faster than higher-level systems. Higher-level systems cannot produce a complex perceptual result before the lower-level systems have produced the component perceptions on which it depends. This nesting of control actions can be seen in the differential speed of operation of control systems at different levels of the control hierarchy. Lower-level systems not only cornea for disturbances faster than higher-level ones; they carry out this correction process during the higher-level correction process. The lower-level control process is temporally nested within the higher-level control process.

Arm Movement

Powers, Clark, & McFarland (1960) describe a simple demonstration of nested control based on relative timing of control system operation. A subject holds one hand extended straight ahead while the experimenter maintains a light downward pressure on it. The subject is to move his or her arm downward as quickly as possible when the experimenter signals with a brief, downward push on the subject's extended hand. The result of this simple experiment is always the same: the subject responds to the downward signal push with a brief upward push followed by downward movement of the arm. An electromyograph shows that the initial upward push is an active response and not the result of muscle elasticity.

The arm movement demonstration reveals one level of control nested within another. The subject's initial upward push (which cannot be suppressed) is the fast response of a lower-level control system that is maintaining the perception of arm position in a particular reference state (extended forward). The behavior of this system is nested within the response time of a higher-level system that moves the arm downward. The higher-level system operates by changing the reference for the arm-position control system. The downward signal push causes the brief upward reaction because the signal is treated as a disturbance to arm position. This is particularly interesting because the signal is pushing the arm in the direction it should move; the lower-level reaction is "counterproductive" with respect to the goal of the higher-level system (which wants to perceive the arm down at the side). The reaction occurs because the lower-level system starts pushing against the disturbance to arm position before the higher-level system can start changing the reference for this position.

Polarity Reversal

More precise tests of nested control were carried out in a series of experiments by Marken & Powers (1989). In one of these experiments, subjects performed a standard pursuit tracking task, using a mouse controller to keep a cursor aligned with a moving target. At intervals during the experiment, the polarity of the connection between mouse and cursor movement was reversed in a way that did not disturb the cursor position. Mouse movements that had moved the cursor to the right now moved it to the left; mouse movements that had moved the cursor to the left now moved it to the right.

A sample of the behavior that occurs in the vicinity of a polarity reversal is shown in Figure 2. The upper traces show the behavior of a control-system model, and the lower traces show the behavior of a human subject. When the reversal occurs, both the model and the subject respond to error (the deviation of the cursor from the target) in the wrong direction, making it larger instead of smaller (any deviation of the error trace from the zero line represents an increase in error). The larger error leads to faster mouse movement, which causes the error to increase still more rapidly. A runaway condition ensues, with error increasing exponentially.

About 1/2 second after the polarity reversal, the subject's behavior departs abruptly from that of the model. The subject adjusts to the polarity reversal, and the error returns to a small value. The model cannot alter its characteristics, and so the error trace quickly goes off the graph. These results provide evidence of two nested levels of control operating at different speeds. The faster, lower-level system controls the distance between cursor and target. This system continues to

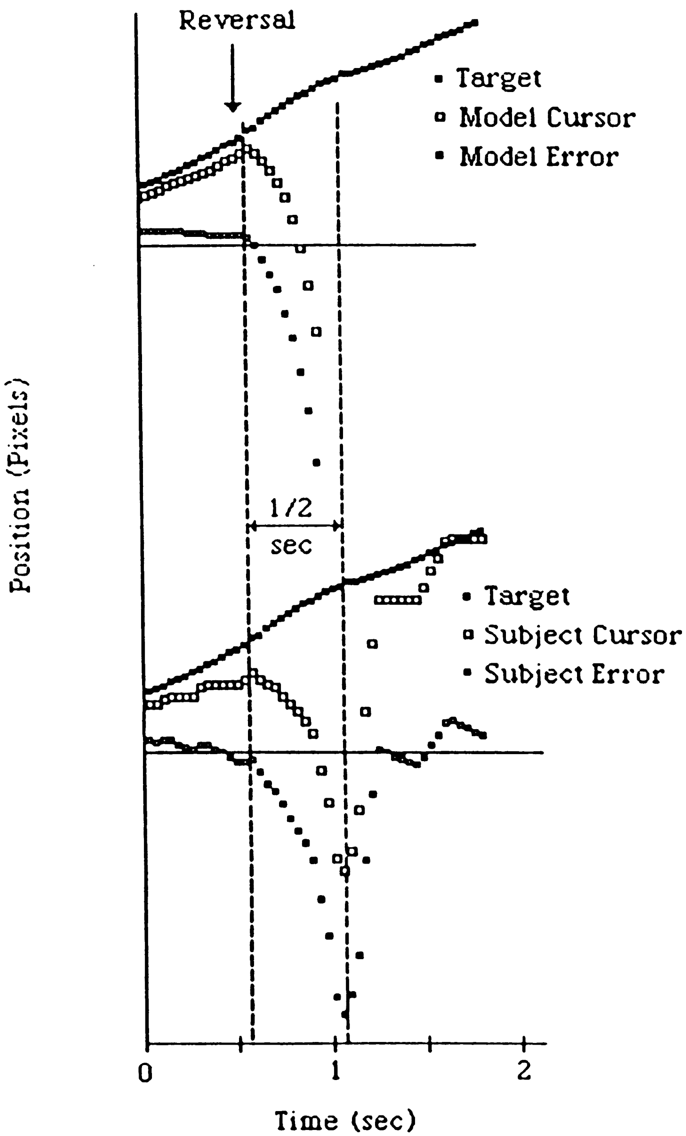

Figure 2. Low-level runaway response to mouse-cursor polarity reversal (after Marken & Powers, 1989, p. 415).

operate as usual, even when, due to the polarity reversal, this causes an increase in perceptual error. Normal operation is restored only after a slower, higher-level system has time to control the relationship between mouse and cursor movement.

Levels of *Perception*

The arm movement and polarity shift experiments reveal the hierarchical organization of control from the point of view of the observer. The hierarchical-control model suggests that it should also be possible to view hierarchical organization from the point of view of the actor. From the actor's point of view, hierarchical control would be seen as a hierarchy of changing perceptions. One way to look at this hierarchy is again in terms of relative timing—in this case, however, in terms of the relative timing of the perceptual results of control actions, instead of the actions themselves.

Computation Time Window

The hierarchical-control model represents the results of control actions as unidimensional perceptual signals. A configuration, such as the letter "h; " is a possible result of control actions, as is a sequence of letters, such as the word "hello." The model represents these results as perceptual input signals, the intensity of a signal being proportional to the degree to which a particular result is produced. This concept is consistent with the physiological work of Hubei & Wiesel (1979), who found that the firing rate of an afferent neuron is proportional to the degree to which a particular environmental event occurs in the "receptive field" of the neuron.

Many of the higher-level classes of perception in the control hierarchy depend on environmental events that vary over time. Examples are transitions, events, and sequences. The neural signals that represent these variables must integrate several lower-level perceptual signals that occur at different times. Hubei and Weisel found evidence of a computation time window for integrating perceptual signals. Certain cells respond maximally to configurations (such as "lines") that move across a particular area of the retina at a particular rate. These are "motion detector" neurons. The neurons respond maximally to movements of configurations that occur within particular time windows. Movements that occur outside of these time windows are not included in the computations of perceptual signals representing motion.

The hierarchical-control model implies that the duration of the computation time window increases at higher levels in the hierarchy. The minimum computation time window for the perception of configurations should be shorter than the minimum computation time window for the perception of transitions, which should be shorter than the minimum computation time window for the perception of sequences. I have developed a version of the psychophysical method of adjustment that makes it possible to see at least four distinct levels of perception by varying the rate at which items occur on a computer display. A computer program presents a sequence of numbers at two different positions on the display. The presentation positions are vertically adjacent and horizontally separated by two centimeters. The numbers are presented alternately in the two positions. The subject can adjust the rate at which the numbers occur in each position by varying the position of a mouse controller.

The results of this study are shown schematically in Figure 3. At the fastest rate of number presentation, subjects report that the numbers appear to occur in two simultaneous streams; the fact that the numbers are presented to the two positions alternately is completely undetectable. However, even at the fastest rate of number presentation, subjects can make out the individual numbers in each stream. At the fastest rate, there are approximately 20 numbers per second in each stream. This means that there is a 50-millisecond period available for detecting each number. This duration is apparently sufficient for number recognition, suggesting that the computation time window for perception of configuration is less than 50 milliseconds. Studies of the "span of apprehension" for sets of letters suggest that the duration of the computation time window for perception of visual configuration might be even less than 50 milliseconds, possibly as short as 15 milliseconds (Sperling, 1960).

As the rate of number presentation slows, the alternation between numbers in the two positions becomes apparent. Subjects report perception of alternation or movement between numbers in the two positions when the numbers in each stream are presented at the rate of about seven per second. At this rate, an alternation from a number in one stream to a number in another occurs in 160 milliseconds. This duration is sufficient for perception of the alternation as a transition or movement from one position to the other, suggesting that the computation time window for transition perception is on the order of 160 milliseconds. This duration is compatible with estimates of the time to experience optimal apparent motion when configurations are alternately pre-

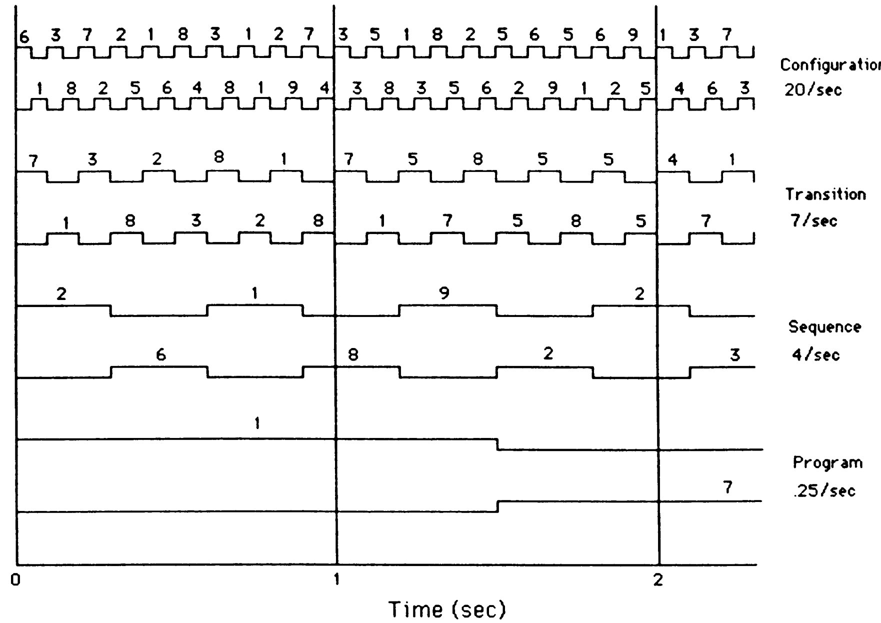

Figure 3. Schematic representation of the results of the number rate adjustment study.

sented in two different positions (Kolers,1972).

The numbers presented in each stream are always changing. However, subjects find it impossible to perceive the order of the numbers as they alternate from one position to another, even though it is possible to clearly perceive the individual numbers and the fact that they are alternating and changing across positions. The rate of number presentation must be slowed considerably, so that each stream of numbers is presented at the rate of about two per second, before it is possible to perceive the order in which the numbers occur. At this rate, numbers in the sequence occur at the rate of four per second. These results suggest that the duration of the computation time window for the perception of sequence is about 05 seconds. This is the time it takes for two elements of the sequence to occur—the minimum number that can constitute a sequence.

The numbers in the rate-adjustment study did not occur in a fixed, repeating sequence. Rather, they were generated by a set of rules—a program. The sequence of numbers was unpredictable unless the subject could perceive the rule underlying the sequence. The rule was as follows: if the number on the right was even, then the number on the left was greater than five; otherwise, the number on the left was less than five. (Numbers in the sequence were also constrained to be between zero and nine). Subjects could not perceive the program underlying the sequence of numbers until the speed of the two streams of numbers was about 025 numbers per second, so that the numbers in the program occurred once every two seconds. The perception of a program in a sequence of numbers requires considerably more time then it takes to perceive the order of numbers in the same sequence.

The perception of a sequence or a program seems to involve more mental effort than the perception of a configuration or a transition. Higher-level perceptions, like programs, seem to represent subjective rather than objective aspects of external reality; they seem more like interpretations than representations. These higher-level perceptions are typically called “cognitions.” Of course, all perceptions represent subjective aspects of whatever is “out there”; from the point of view of the hierarchical-control model, the location of the line separating perceptual from cognitive representations of reality is rather arbitrary. Behavior is the control of perceptions which range from the simple (intensities) to the complex (programs).

Perceptual Speed Limits

The hierarchical-control model says that all perceptions of a particular type are controlled by systems at the same level in the hierarchy. This implies that the speed limit for a particular type of perception should be about the same for all perceptions of that type. The 160 millisecond computation time window for perception of transition, for example, should apply to both visual and auditory transition There is evidence that supports this proposition Miller & Heise (1950) studied the ability to perceive an auditory transition called a “trill.” A trill is the perception of a temporal alternation from one sound sensation or configuration to another. The speed limit for trill perception is nearly the same as the speed limit for visual transition perception found in the number rate adjustment study—about 15 per second. As in the visual case, when the rate of alternation of the elements of the auditory trill exceeds the computation time window, the elements “break” into two simultaneous streams of sound; the perception of transition (trill) disappears, even though the sounds continue to alternate.

There is also evidence that the four-per-second speed limit for sequence perception found in the number-rate adjustment study applies across sensory modalities. Warren, Obusek, Farmer, & Warren (1969) studied subjects’ ability to determine the order of the component sounds in a sound sequence. They found that subjects could not perceive the order of the components until the rate of presentation of the sequence was less than or equal to four per second. This was a

surprising result, because it is well known that people can discriminate sequences of sounds that occur at rates much faster than four per second. In words, for example, the duration of the typical phoneme is 80 milliseconds, so people can discriminate sequences of phoneme sounds that occur at the rate of about 10 phonemes per second. But there is reason to believe that the phonemes in a word are not heard as a sequence; that is, the order of the phonemes cannot be perceived. Warren (1974) showed that subjects can learn to tell the difference between sequences of unrelated sounds that occur at rates of 10 per second. However, the subjects could not report the order of the sounds in each sequence, only that one sound event differed from another. A word seems to be a lower-order perception—an event perception—that is recognized on the basis of its overall sound pattern There is no need to perceive the order in which the phonemes occur, just that the temporal pattern of phonemes (sound configurations) for one word differs from that for other words.

The Relationship Between Behavior and Perception

Configurations, transitions, events, sequences, and programs are potentially controllable perceptions. An actor can produce a desired sequence of sounds, for example, by speaking sound events (phonemes) in some order. An observer will see the production of this sequence as a behavior of the actor. The hierarchical-control model suggests that the actor’s ability to produce this behavior turns on his or her ability to perceive the intended result. Since perception depends on speed, it should be impossible for the actor to produce an intended result faster than the result can be perceived. The observer will see this speed limit as a behavioral limit. An example can be seen in the arm-movement experiment described above. In that experiment, it appears that the time to respond to the signal push is the result of a behavioral speed limit: the inability to generate an output faster than a certain rate. But a closer look indicates that the neuromuscular “output” system is perfectly capable of responding to a signal push almost immediately, as evidenced by the immediate upward response to the downward signal push. The same muscles that produce this immediate reaction must wait to produce the perception of the arm moving downward. The speed limit is not in the muscles. It is in the results that the muscles are asked to produce; a static position of the arm (a configuration perception) or a movement of the arm in response to the signal push (a relationship perception).

Sequence Production and Perception

Some of the most interesting things people do involve the production of a sequence of behaviors. Some recent studies of temporal aspects of sequence production are directly relevant to the hierarchical-control model. In one study, Rosenbaum (1987) asked subjects to speak the first letters of the alphabet as quickly as possible. When speed of letter production exceeded four per second, the number of errors (producing letters out of sequence) increased dramatically, indicating a loss of control of the sequence. The speed limit for sequence production corresponds to the speed limit for sequence perception—four per second.

The letter-sequence study does not prove that the speed limit for letter-sequence production is caused by the speed limit for letter-sequence perception. It could be that the speed limit is imposed by characteristics of the vocal apparatus. However, in another study, Rosenbaum (1987) found the same four-per-second speed limit for production of errorless finger-tap sequences. The speed limit for finger-tap sequence production is likely to be a perceptual rather than a motor limit, because we know that people can produce finger taps at rates much higher than four per second. Pianists, for example, can do trills (alternating finger taps) at rates which are far faster than four per second. Further evidence of the perceptual basis of the finger-tap sequence speed limit would be provided by studies of finger-tap sequence perception. When a subject produces a sequence of finger taps, he or she is producing a sequence of perceptions of pressure at the finger tips. A perceptual experiment where pressure is applied to the tips of different fingers in sequence should show the four-per-second speed limit. Subjects should have difficulty identifying the order of finger-tip pressures when the sequence occurs at a rate faster than four per second.

Confounding Levels

It is not always easy to find clear-cut cases of behavioral speed limits that correspond to equivalent perceptual speed limits. Most behavior involves the control of many levels of perception simultaneously. People control higher-level perceptions (like sequences) while they are controlling lower-level perceptions (like transitions). This can lead to problems when interpreting behavioral speed limits. For example, Rosenbaum (1983) presents some finger tapping results that seem to violate the four-per-second speed limit for sequence perception. When subjects tap with two hands, they can produce a sequence of at least eight finger taps per second. But each tap is not necessarily a separate event in a sequence. Some pairs of taps seem to occur at the rate at which

sequences are experienced as events. A sequence of finger taps is an event in the same sense that the sequence of muscle tensions that produce a finger tap is an event; the order of the components of the sequence cannot be perceived. These finger-tap events are then unitary components of the sequence of finger-tap perceptions.

The fact that cetain pairs of finger taps are produced as events rather than ordered sequences is suggested by the errors made at each point in the finger-tap sequence. Errors occur most frequently at the point in the sequence at which a fast pair is being initiated. Errors rarely occur for the second element of a fast pair. This suggests that the errors occur at the sequence level rather than the event level. The subject’s attempts to produce a key-press sequence too rapidly apparently interfere with sequence rather than event production. Events are already produced at a fast enough rate, and an increase in the speed of sequence production has little effect on the ability to control the component events.

Changing Perception Can Change Behavior: Going Up A Level

The relationship between perception and behavior can be seen when a person learns to perform a task by controlling a new perceptual variable. An example of this can be seen in simple pursuit-tracking tasks. In the typical tracking task, the target moves randomly. When, however, a segment of target movement is repeated regularly, the subject’s tracking performance improves markedly with respect to that segment (Pew, 1966). According to the hierarchical-control model, control is improved because the repeated segment of target movement can be perceived as a predictable event. With the random target, the subject must wait to determine target position at each instant in order to keep the cursor on target. With the repeated target, the subject controls at a higher level, keeping a cursor-movement event matching a target-movement event. The fact that the subject is now controlling a higher-level perception (an event, rather than a configuration) is evidenced by the longer reaction time when responding to a change in target movement. When controlling the target-cursor configuration, the subject responds almost immediately to changes in target position. When controlling target-cursor movement events, it takes nearly 1/2 second to respond to a change in target-movement pattern.

An experiment by Robertson & Clines (1985) also shows improved performance resulting from changed perception. Subjects in the Robertson and Clines study performed a learning task where the solution to a computerized game could be perceived at several different levels. Subjects who were able to solve the game showed three distinct plateaus in their performance. The level of performance, as indicated by reaction-time measurements, improved at each succeeding plateau. Because the same outputs (key presses) were produced at each level of performance, each performance plateau was taken as evidence that the subject was controlling a different perceptual variable.

Behavior/Perception Correlations

Few psychologists would be surprised by the main contention of this paper: that there is an intimate relationship between perception and behavior. However, most models of behavior assume that the nature of this relationship is causal: that behavior is guided by perception. This causal model provides no reason to expect a relationship between the *structure* of perception and behavior. For example, the causal model provides no reason to expect a relationship between the ability to identify a sequence of sounds (perception) and the ability to produce a sequence of actions (behavior). This does not mean that the model rules out such relationships; it just does not demand them.

The control model integrates perception and behavior. Behavior is no longer an output, but instead a perceptual input created by the combined effects of the actor and the environment. Behavior is perception in action. From this point of view, behavioral skills are perceptual skills. Thus, it is not surprising to find some indication of a correlation between behavioral and perceptual ability. For example, Keele and his colleagues (Keele, Pokorny, Corcos, & Ivry, 1985) have found that the ability to produce regular time intervals between actions is correlated with the ability to perceive these intervals. These correlations are fairly low by control-theory standards, but they are expected if the production of regular time intervals involves control of the perception of these intervals.

Conclusion

This report has presented evidence that human behavior involves control of a hierarchy of perceptual variables. There is evidence that the behavior of non-human agents, such as chimpanzees, also involves the control of a similar hierarchy of perceptions (Plooij & van de Rijt-Plooij, 1990). A model of hierarchical control shows how studies of perception and behavior provide evidence about the nature of control from two different perspectives. Perceptual studies provide information about the ability to perceive potentially controllable consequences of actions. Behavioral studies provide information about the ability to produce desired consequences. The factors that influence the ability to perceive the consequences of action should also influence the ability to produce them. In both cases, we learn something about how agents control their own perceptions.

The hierarchical-control model implies that limitations on the ability to produce behavior reflect limitations on the ability to perceive intended results. The speed at which a person can produce an errorless sequence of events, for example, is limited by the speed at which the order of these events can be perceived. But not all skill limitations are perceptual limitations. Controlled (perceived) results are produced, in part, by the outputs of the behaving agent. The ability to produce certain outputs can limit the ability to control certain perceptions. For example, it is impossible to perceive oneself lifting a 300-pound barbell until the muscles have been developed to the point that they are able to generate the output forces necessary to control this perception.

Perception and behavior are typically treated as two completely different types of phenomena. Perception is a sensory phenomenon; behavior is a physical phenomenon. But the concept of control as the behavior of perception suggests that this separation is artificial. Perception and behavior are the same phenomenon seen from two different perspectives. In order to understand how this phenomenon works, it will be necessary to understand how agents perceive (perception) and how they act to affect their perceptions (behavior). Studies of perception and behavior should become an integral part of the study of a single phenomenon: control.

Availability of Software

A HyperCard version of the number-rate-adjustment program can be obtained from the author. Send a formatted 3.5-inch double-density or high-density diskette in a reusable mailer with return postage.

References

Albus, J. (1981). *Brains, behavior, and robotics.* Petersborough, NH: Byte Books.

Arbib, M. (1972). *The metaphorical brain.* New York: Wiley.

Bryan, W. L. & Harter, N. (1899). Studies on the telegraphic language: The acquisition of a hierarchy of habits. *Psychological Review, 6, 345*-375.

Davis, W. J. (1976). Organizational concepts in the central motor network of invertebrates. In R. M. Herman, S. Grillner, P. S. G. Stein, & D. Stuart (Eds.), *Neural control of locomotion (p. 265).* New York: Plenum.

Greeno, J. G., & Simon, H. A. (1974). Processes for sequence production. *Psychological Review,* 81,187-197.

Herrigel, E. (1971). *Zen in the art of archery.* New York: Vintage.

Hubel, D. H. & Wiesel, T. N. (1979). Brain mechanisms of vision. In J. M. Wolfe (Ed.), *The mind's eye: Readings from Scientific American* (pp. 40-52). New York: Freeman.

Keele, S. W., Pokorny, R. A., Corcos, D. M., & Ivry, R I. (1985). Do perception and production share common timing mechanisms: A correlational analysis. *Acta Psychologica,* 60,173-191.

Kline, R. (1983). Comment on Rosenbaum, et al. Hierarchical control of rapid movement sequence. *Journal of Experimental Psychology.. Human Perception and Performance,* 9, 834-36.

Kolers, P. (1972). The illusion of movement. In R. Held & W. Richards (Eds.), *Perception: Mechanisms and models (pp.* 316-323). San Francisco: Freeman.

Lashley, K. S. (1951). The problem of serial order in behavior. In L. A. Jeffress (Ed.), *Cerebral mechanisms in behavior: The Hixon Symposium* (pp.112-136). New York: Wiley.

Marken, R S. (1986). Perceptual organization of behavior: A hierarchical control model of coordinated action. *Journal of Experimental Psychology: Human Perception and Performance,* 12, 267-76.

Marken, R. S. (1988). The nature of behavior: Control as fact and theory. *Behavioral Science,* 33, 196-206.

Marken, R. S. (1989). Behavior in the first degree. In W. A. Hershberger (Ed.), *Volitional action: Conation and control (pp.* 299-314). Amsterdam: North-Holland.

Marken, R. S. (1990). Spreadsheet analysis of a hierarchical control system model of behavior. *Behavior Research Methods, Instruments & Computers,* 22, 349-359.

Marken, R. S., & Powers, W. T. (1989). Levels of intention in behavior. In W. A. Hershberger (Ed.), *Volitional action: Conation and control* (pp. 409-430). Amsterdam: North-Holland.

Martin, J. G. (1972). Rhythmic (hierarchical) versus serial structure in speech and other behavior. *Psychological Review,* 79, 487-509.

Miller, G. A., & Heise, G. A. (1950). The trill threshold. *Journal of the Acoustical Society of America,* 22,637-638.

Palmer, S. E. (1977). Hierarchical structure in perceptual representation. *Cognitive Psychology,* 9, 441-474.

Pew, R. W. (1966). Acquisition of hierarchical control over the temporal organization of a skill. *Journal of Experimental Psychology,* 71, 764-771.

Plooij, F. X., & van de Rift-Plooij, H. H. C. (1990). Developmental transitions as successive reorganizations of a control hierarchy. *American Behavioral Scientist,* 34, 67-80.

Povel, D-J. (1981). Internal representation of simple temporal patterns.

Journal of Experimental Psychology: Human Perception and Performance, 7, 3-18.

Powers, W. T. (1973). *Behavior. The control of perception.* Chicago: Aldine.

Powers, W. T. (1979, August). The nature of robots: Part 4: A closer look at human behavior. *Byte, pp.* 309-323.

Powers, W. T. (1989). *Living control systems.* Gravel Switch, KY: Control Systems Group.

Powers, W. T. (1990). A hierarchy of control. In R. J. Robertson & W. T. Powers (Eds.), *Introduction to modern psychology: The control-theory view* (pp. 59-62). Gravel Switch, KY: Control Systems Group.

Powers, W. T., Clark, R. K., & McFarland, R. L. (1960). A general feedback theory of human behavior: Part II. *Perceptual and Motor Skills, 11,* 309-323.

Robertson, R J., & Glines, L.A. (1985). The phantom plateau returns. *Perceptual and Motor Skills,* 61, 55-64.

Rosenbaum, D. A. (1983). Hierarchical control of rapid movement sequences. *Journal of Experimental Psychology: Human Perception and Performance,* 9, 86-102.

Rosenbaum, D. A. (1987). Hierarchical organization of motor programs. In S. P. Wise (Ed.), *Higher brain functions: Recent explorations of the brain's emergent properties.* New York: Wiley.

Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs, 74(11).*

Sternberg, S., Knoll, R. L., Monsell, S., & Wright, C. E. (1989). Motor programs and hierarchical organization in the control of rapid speech. In M. Jeannerod (Ed.), *Attention and performance XIII: Motor representation and control (pp.* 3-55). New Jersey: Erlbaum.

Stemberg, S., Knoll, R. L., & Turlock, D. L. (1990). Hierarchical control in the execution of action sequences. *Phonetica*, 45,175-197.

Warren, R M. (1974). Auditory temporal discrimination by trained listeners. *Cognitive Psychology, 6,* 237-256.

Warren, R M., Obusek, C. J., Farmer, R. M., & Warren, R P. (1969). Auditory sequence: Confusion of patterns other than speech or music. *Science, 164*, 586-587.

Mimicry, Repetition, and Perceptual Control

W. Thomas Bourbon
(Research Division, Department of Neurosurgery, University of Texas Medical School - Houston, 6431 Fannin, Suite 7.148, Houston, TX 77030)

Abstract

In their attempts to explain, predict, and control human behavior, behavioral scientists typically overlook controlling done by themselves and by the people they study. The literature on perceptual control theory (PCT) describes several reasons for that omission, and in this paper I show another. when they mimic events in the environment, or when they repeat actions that they imagine or remember, people "act like" the kinds of lineal causal systems portrayed in most behaviorist and neuro-cognitive theories of behavior. A PCT model can emulate the behavior both of the person who acts like the lineal causal models and of the lineal models themselves. The results described in this paper show that the lineal causal models used in the behavioral sciences produce behavior that is a special limiting case of the behavior exhibited by the control-system model in PCT.

Mimicry and Repetition are Limiting Cases of Perceptual Control

People are living control systems who control many of their own perceptions. This paper is about two circumstances that have led scientists to think people are not living control systems: (1) when people try to mimic the actions of variables in the environment, and (2) when people try to repeat remembered or imagined patterns of actions. In these cases, the behavior of a control system can be mistaken for that of a lineal causal system whose actions are caused by antecedent events. To show that observers can mistake people for cause-effect systems, I use a demonstration that builds on work described in a previous paper, "Models and Their Worlds" (Bourbon & Powers, 1993), hereafter referred to as "Worlds." In the present demonstration, a person does variations on a simple pursuit-tracking task. In the process, the person unintentionally imitates the performance of two popular cause-effect models of people. Then I show a PCT model that duplicates the person's performance, as well as that of each lineal causal model.

The *Experimental Setting: Pursuit Tracking*

Figure 1A shows the experimental setting from "Worlds." A person uses a control handle to affect a cursor (a short horizontal mark on a computer screen) while two target marks unaffected by the handle move in unison up and down on the screen. Figure 1B shows the environmental variables that affect the cursor and target. For each of 1800 moments sampled during a one-minute run and modeled during a simulation, the following program statement determines the position of the cursor:

c: = h + d,

where c is cursor position, h is handle position, t is the momentary value of the target function generated by the computer, and d is the momentary value of a computer-generated disturbance (zero for some runs).

For the first part of the demonstration, the task was the same as the one described in "Worlds":

> The person's task in all phases of the experiment is to keep the cursor exactly between the target lines. (There is nothing special about that relationship between cursor and target; the person could easily select any other.) This task is known as "tracking" (Bourbon & Powers, 1993, p. 55).

Seen Cursor Position Minus Seen Target Position Equals Zero

Figure 1C shows the results when the person kept the cursor aligned with a moving target. The target moved up and down at a constant velocity, and no disturbance affected the cursor (d = 0). The person moved the handle in a pattern that necessarily, but unintentionally, resembled the pattern for the target.

Perceptual control theorists often use the PCT model to reproduce and predict results like these. Correlations between predicted and actual handle positions often exceed .995, even when the predictions precede the person's data by one year (Bourbon, Copeland, Dyer, Harman, & Mosley, 1990) or five years (Bourbon, 1993a). In those studies, people kept the cursor aligned with the target, but a person could easily select any other relationship to control, as I show next.

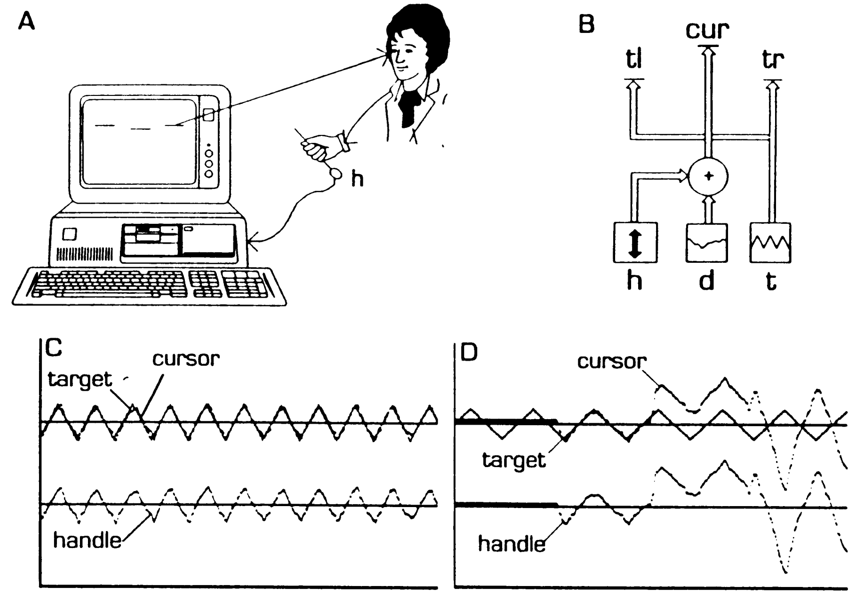

Figure 1. A. The experimental setup, in which a person uses a control handle to keep a cursor in a desired relationship with a moving target on a computer screen. B. Environmental connections among handle position (h), the disturbance function (d), target function (t), target marks (tl and tr), and cursor (cur). C. Results when a person did the tracking task and d was zero. D. Results when a person did the tracking task and, during successive 15-second periods, (a) did not move the handle, (b) kept the cursor aligned with the target, (c) kept the cursor one inch above the target, and (d) moved the cursor to twice the inverse of the target position. (In the plots, "up" represents the handle moving away from the person, and the cursor and target moving upward on the screen. The horizontal axis of each plot represents time, from 0 to 60 seconds.)

Seen Cursor Position Minus Seen Target Position Equals Various Values

Figure 1D shows the results with the person in the same setting as before, but with the target moving at a slower velocity. During successive 15-second intervals, the person (a) did not move the handle, (b) kept the cursor even with the target, (c) kept the cursor an inch above the target, and (d) moved the cursor to positions twice as great as the inverse of the target. The person did not need practice to produce these results. Bourbon (1993b) showed that a simple PCT model can duplicate results like these. When a person and a PCT model adopt and create different intended perceptions, they disprove the common misconception that control systems cannot change their “goals” or intended results.

Predictions by the Three Models from “Worlds”

In “Worlds,” after a run with the conditions shown in Figure 1C, we tested two popular lineal causal models and the model from perceptual control theory. We compared the models’ predictions of what the person would do when the experimental conditions changed. We described the models in detail in “Worlds”; I summarize them in this paper’s Appendix.

Running the Models

In ‘Worlds,” we described the procedures for running (simulating) each model. We used the person’s data from an initial experiment to estimate the parameters for each model, then ran the models under altered conditions. The present demonstration followed the same procedure: I used data from Figure 1C to estimate the parameters of the models, then ran them in simulation. The top row of Figure 2 shows the results of the simulations, which are the same as those in Phase 3 of “Worlds” (p. 65). Each result is a quantitative prediction by a model (described in the Appendix) of what would happen if the person functioned like that particular model.

The PCT model. The PCT model tests the idea that when the person produced the results in Figure 1C, he compared his momentary perceptions against what he intended to perceive. When there was a mismatch between present and intended perceptions, his actions changed to create and maintain a match. If the person acted that way during the first task, then he could probably keep the cursor aligned with the target, even when it followed a new and variable pattern and a random disturbance affected the cursor. His handle positions, which would vary as necessary to oppose the random disturbance, would be unintended side-effects of control and would no longer duplicate the positions of the target or the cursor they control.

In the present simulation of the PCT model (Figure 2A), the reference signal specified the perceptual signal, and any discrepancy between the signals drove the handle to positions that canceled the effects of the disturbance to the cursor. The cursor remained aligned with the target, as was intended, and the position of the handle was an unintended side-effect of control.

The S-R *model.* A stimulus-driven (stimulus-response, S-R) model tests the idea that for the results in Figure 1C, the position of the target

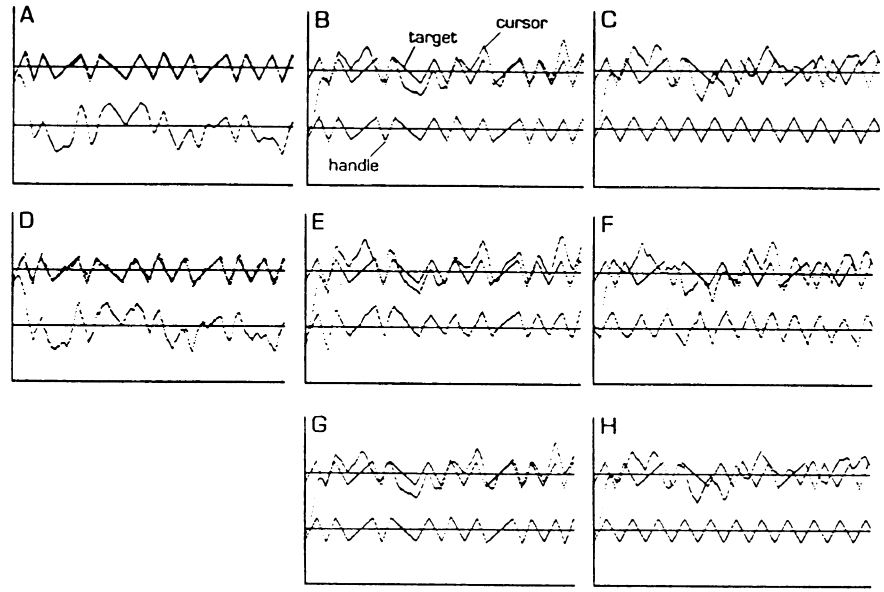

Figure 2. Top row, predictions by (A) a PCT model (cursor position - target position = zero), (B) an S-R model (handle position = target position), and (C) a plan-driven model (handle position = planned handle position). (These models are described in the Appendix.) Middle row, data when the person controlled (D) to keep the seen cursor even with the seen target, (E) to keep felt (unseen) handle position = seen target position, and (F) to keep felt (unseen) handle position = planned handle position. Bottom row, results when the PCT model impersonated the other models, with reference signals for (G) handle position = target position, and (H) handle position = planned handle position. In each run or simulation, the target path was the same, and the same random disturbance affected the cursor. (In the plots, “up” represents the handle moving away from the person, and the cursor and target moving up on the screen. The horizontal axis of each plot represents time, from 0 to 60 seconds.)

reflexively determined the position of the person’s control handle. If the person acted that way during the first run, then his handle could still follow the target when it traced a new and variable pattern and a random disturbance affected the cursor. In that case, the position of the cursor would become an unintended side-effect of control. In the present simulation of the S-R model (Figure 2B), the target determined the position of the handle, and their momentary positions were nearly identical. The cursor “wandered” away from the target and its position was an unintended side-effect.

The plan-driven model. The plan-driven neuro-cognitive model tests the idea that, for the results in Figure 1C, the person's memory of momentary handle positions from earlier practice sessions determined the position of his control handle. If the person acted that way during the first run, then his handle positions should duplicate the ones in Figure 1C, even when the target followed a new and variable pattern, and a random disturbance affected the cursor. Neither the handle nor the cursor would duplicate the pattern of movement traced by the target. The position of the cursor would be an unintended side-effect of control.

In the present simulation of the plan-driven neuro-cognitive model (Figure 2C), the plan for a pattern of target movements ("remembered" from the data in Figure 10 determined the position of the handle. The uncontrolled cursor wandered independently of the target. Its position was a side-effect of control.

The Person Performs Under New Conditions

In the simulations I just described, the three models predicted different results for the person running under the environmental conditions from Phase 3 of "Worlds" (pp. 64-67). Now I report what happened when the person repeated the tracking task three times under those conditions. A random disturbance affected the position of the cursor, and, from one excursion to the next up or down the screen, the probability was 2/3 that the velocity of the target would change to another of three possible values. During the first repetition, the person again kept the cursor aligned with the target; in the other two, he created results in which the position of the cursor became an uncontrolled side effect.

Seen Cursor Position Minus Seen Target Position Equals Zero

First, the person kept the cursor aligned with the target. Figure 2D shows the results. The patterns of positions for the target and cursor were similar ($r = .91$, $n = 1800$ data pairs). The pattern of the person's handle movements necessarily differed from that in Figure 1C. He controlled the relationship of the cursor and target, which was a consequence of actions, but did not control his actions. The relationship between his actions and the movements of the target necessarily varied to eliminate effects of the disturbance on the cursor's position. It is impossible for a person to specify and plan the required actions before a condition as variable and disturbed as this one.

Failed models? In "Worlds," we compared the person's data in this condition against the predictions by the two lineal causal models, which

in the present demonstration is the same as comparing Figure 2D with Figure 2B and 2C. Obviously, the results when the person used the handle to keep the cursor aligned with the target were different from the results of the lineal causal models. Those models controlled the position of the simulated handle, but not the cursor. The lineal models, which can accurately explain the results of the undisturbed condition shown in Figure 1C, failed to predict the results in Figure 2D. The person did not act like a lineal causal system, but *can* he?

Mimicry and repetition. So far, the person has controlled the position of a cursor compared with a target, and the positions of his control handle were unintended and uncontrolled side-effects. What would happen if he did not control the position of the cursor at all, and instead controlled his actions? The position of the cursor would become an uncontrolled side-effect. Next, I show the results when the person ran under the same conditions as shown in Figure 2D but controlled his felt perceptions of hand movements. First he made them match the seen movements of the target, then he made them match a remembered pattern of felt movements. In the first case, his movements mimicked a present perception in another sensory modality. They appeared to fit the S-R model, where "stimuli" (movements of the target) cause "responses" (movements of the handle). In the second case, his movements repeated a remembered pattern. They appeared to fit the plan-driven neuro-cognitive model, where plans or commands from the mind-brain control handle movements, independent of events in the environment.

Mimicry: Felt Handle Position Equals *Seen Target Position*

People sometimes make their actions mimic those of other people. Some children who watch adults playing musical instruments use toy instruments and make exaggerated motions that they believe are the same as the adult's actions. Sometimes an inexperienced person attempts to perform without practice in a marching band or military unit by watching and duplicating the actions of others in the unit. In gatherings, sometimes individuals mimic what they see other people doing. In cases like these, people try to make their felt actions match actions or events they see in the environment.

During the present demonstration, the person made his felt, but unseen, handle movements match the movement of the target. By making his actions duplicate the movements of an environmental stimulus, he played the role of an S-R system; he functioned like a control system, making his presently perceived hand position match the presently seen position of the target. To help him play that role, the target function and the disturbance remained the same for 15 practice runs, and a piece of cardboard shielded his hand from view. He practiced nuking his felt hand movements match the movements of the target on the screen. When he decided that he was "ready," he did the run shown in Figure 2E and 11 similar runs, for a total of 12 runs.

The person's controlled handle movements generally resembled the pattern of target movements. The mean correlation for twelve sets of predicted and actual handle positions was .824 (S.D. = .089, range = .981 to .615, n = 1800 data pairs per set). By accepted standards in behavioral science, that mean correlation is extremely high, but agreement between predicted and actual handle positions is even higher when a person keeps a cursor aligned with a target. In a study with 104 sets of 1800 predicted and actual handle positions, Bourbon et al. (1990) reported a mean correlation of .996 (S.D. = .002). It is easier for people to make one seen environmental variable track another than to make their own felt actions "track" a seen variable.

There were obvious differences between movements of the target and handle. For example, the target always moved at one of three uniform velocities, but the person's handle velocities were not uniform. Also, before reversing direction, the target always moved the same distance above or below the center of the screen, but the person reversed handle movements at varying distances from the center of their range. He did not perfectly duplicate the performance of a pure S-R system. Even after 12 practice sessions, it was not easy for him to judge and control either the velocity of handle movements or the distances he moved the handle before reversing its direction.

During the present trials, the cursor was affected by a random disturbance and by the handle. It "wandered" around the position of the target. Cursor position was an accidental side-effect when the person controlled the position of the handle.

Repetition: Felt Handle Position Equals Remembered Handle Position

Sometimes people make their patterns of actions repeat a remembered or imagined pattern. Many self-improvement and rehabilitation programs urge clients to imagine themselves doing a desired action "perfectly;" then to do the action as imagined. When people attempt to move through a darkened familiar environment, they sometimes try to duplicate movements they remember from when they could see their surroundings. In a group that uses a device like a baton or banner to do synchronized routines, some people who drop the device try to continue making movements remembered from performances when they held it. In cases like these, people try to make the actions they feel match patterns they remember or imagine.

In the present demonstration, the person did the condition shown in

Figure 2, except that he made the pattern of his felt-but-unseen handle movements match the pattern he remembered from the condition shown in Figure 1C. By making his actions duplicate the earlier pattern, he imitated the performance of a neuro-cognitive plan-driven system. To help him act that role, he ran 22 replications of the undisturbed task shown in Figure 1C and kept the cursor aligned with the target. A piece of cardboard screened his hand from view, and he paid close attention to the tactile and kinesthetic sensations that accompanied successful tracking. He intended to repeat the practiced movements from memory when the screen was blanked during the next task When he was ready, the program started. The initial positions of all variables were displayed on the screen, then the screen went blank and he completed the run shown in Figure 2F and 15 additional runs, for a total of 16 runs.

Qualitatively, the pattern of the person's controlled handle movements resembled the one from Figure 1C. Quantitatively, the match between modeled and actual patterns of handle movement was atrocious. The mean correlation for sixteen sets of 1800 predicted and actual handle positions was -0.003 (S.D. = .118, range = .390 to -.223). It was *much* harder for the person to create a precise replica of a highly practiced regular pattern of handle movements than to make either a cursor or his hand movements match a seen target. This result has serious implications for all neuro-cognitive plan-driven models of behavior, but especially for those where people claim that the elimination of sensory "feedback" does not affect planned actions. In the present case, simply concealing the person's hand behind a piece of cardboard eliminated precise repetition of the desired pattern.

There were obvious differences between handle movements during the undisturbed run and this one. In the undisturbed run, where handle position was an accidental side-effect of control, the velocity of the person's handle movements necessarily approximated the uniform velocity of the target; in the plan-driven run, where he controlled the handle's positions, handle velocities were more erratic. Also, during the plan-driven run, he reversed the direction of the handle at varying distances from the center of its range; during the undisturbed run, when the position of the handle was an unintended side-effect, the reversals were more uniform. Even after 22 practice sessions, it was not easy for him to judge and control either the velocity of handle movements or the distances he moved the handle before he reversed its direction

The person labored under other serious burdens that confront every Plan-driven system. Such systems are extraordinarily sensitive to the slightest errors in the timing of actions and to the smallest deviations from the required values of any important variables. A deviation at any time during the running of such a system can quickly lead to actions and consequences that are the *reverse* of what they should be. We discussed this extreme sensitivity to small errors in "Worlds" (p. 59), but we did not show quantitative examples of its consequences. Plan-driven models cannot serve as general models of human behavior.

Comparing the Models and the Person

The person's handle positions (Figure 2E and 2F) were more variable than those of the corresponding lineal causal models (Figure 2B and 2C, respectively), in large part due to his not maintaining uniform velocities for the handle. Also, the person moved the handle through a pattern that was not centered in the range of movement, but the models centered their simulated handles. Finally, the plan-driven model perfectly "remembered" the pattern of target movement from the first run, but the person obviously did not; he reversed the direction of handle movement at the wrong times, compared to the ideal remembered pat tern. When it comes to controlling one's own actions, what happens is not always what the person remembers and intends.

The PCT Model Emulates the Person and the Causal Models

Pure causal systems, like the lineal models I explained earlier, cannot produce unvarying results in a variable environment. In 'Worlds," we described a rationale for making causal models succeed in a variable world:

> To modify cognitive or SR models so that, like living systems, they might thrive amidst change, we must... give each model an internal standard and a process for comparing present perceptions against that standard. But then the models would all be control systems, each controlling its input (Bourbon & Powers, 1993, p. 70).

We cannot modify either a pure S-R model or a pure plan-driven model so that it emulates the PCT model, yet simultaneously preserve its core structure. On the other hand, we can easily modify a PCT model so that it emulates either lineal causal model: All we have to do is change p*, the reference signal for the PCT model.

The PCT Model Emulates the S-R Model

To emulate the S-R model, where the position of the target determines the position of the model's handle, the PCT model makes its perceived handle position match the perceived position of the target. The reference signal, p*, becomes h - t = zero, where h and t are positions of the handle and target. Any perceived discrepancy (error signal) between h and t changes the position of h, according to the following program steps:

p: = h - t
error: = p* - p
h: = h + k • error • dt

With no other change, the PCT model will "impersonate" the S-R model (and the person, when he made his felt handle movements match seen target movements).

Figure 2G shows the results when the "modified" PCT model ran in simulation. It reproduced the results of the pure S-R model (Figure 2B): the disturbed and uncontrolled cursor no longer tracked the target, but handle movements, which were now controlled, accurately tracked target movements. This PCT model also reproduced general features of the person's attempt at impersonating a stimulus-driven system, shown in Figure 2E. However, the agreement between the PCT model and the person would be just as poor as that between the S-R model and the person.

The PCT Model Emulates the Plan-Driven Mode!

The PCT model can emulate the plan-driven model, where the computed or remembered pattern of previous target positions determines the position of the model's handle. In that role, the PCT model specifies that the perceived handle position at any moment matches the computed position. The reference signal, p*, for the PCT model becomes h - H = zero, where h is the present position of the handle, and H is the momentary computed or remembered ideal position. A perceived discrepancy between those positions produces movements of the handle, according to the following steps in the computer program:

p: = h - H
error: = p* - p
h: = h + k • error • dt

With no other change, the PCT model will emulate the plan-driven model (and the person, when he made his felt handle movements match a remembered pattern of handle movements).

Figure 2H shows the results when the "modified" PCT model ran in simulation. It accurately duplicated the results of the pure plan-driven model (Figure 2C). The PCT model also reproduced qualitative features of the person's attempt at impersonating a neuro-cognitive plan-driven system, shown in Figure 2F. However, the agreement between the PCT model and the person would be just as poor as that between the neurocognitive plan-driven model and the person.

Discussion

A person can act like a system where environmental stimuli control its actions, and like one where internal plans and commands control its actions; a PCT model can achieve the same results as the person, but neither a pure stimulus-controlled model nor a pure plan-driven model can duplicate all of the appearances of a person and of the other models. To make either cause-effect model to do that, we would need to radically change its core structure and convert it into a perceptual control system. However, for people or PCT models to act like lineal causal systems, their core structures do not change. All that changes for the person is the intended perception; for the PCT model, only the reference signal changes.

Generality of the PCT Model

In the present demonstration, a person used the experimental arrangement shown in Figure 1A to achieve several different controlled results. In the second stage of the demonstration, three different models of behavior each predicted one of the person's results: the PCT model kept its cursor aligned with a target, the S-R model made its handle movements match target movements, and the plan-driven model made its handle movements match a remembered plan. The success of all three models during that stage does not mean that we need a different model to explain the person's performance in each condition. To the contrary, in the final stage of the demonstration, a PCT model with a simple change in its reference signal duplicated all of the results of the person and the two lineal causal models. Perceptual control theory provides a *general* model of control behavior, while each of the lineal models applies only to a limiting case.

There is no defense for using either lineal causal model as a general model of behavior, but many behavioral scientists do. The settings where scientists believe the environment controls a person's behavior are diverse. They range from behavioral conditioning laboratories, where scientists say environmental stimuli control a person's actions, to social gatherings, where they say people "lose control" of their behavior, with control passing to presumed forces such as a "virus-like emotional contagion" or a "group mind." Instead of proving the legitimacy of a stimulus-response model, those are instances when, for whatever reasons, people intend to perceive their actions matching

perceptions of a selected feature of the environment. Other consequences of a person's actions would, like the position of the cursor in the present demonstration, "go out of control." Events like these often catch the attention of observers, whether they are behavioral scientists or the local constabulary, but those observers are wrong if they assume that the person has "lost control" to "powerful" forces in the environment.

There are also many settings where scientists believe that a plan (command, trait, neural signal, gene, force) from the mind-brain controls a person's behavior. They range from concert halls, where many scientists say that some performers' actions occur too regularly and rapidly for the environment to affect them, to neurophysiological clinics, where they say that people with damaged spinal sensory nerves provide evidence that motor plans determine the course of behavior. Instead of proving the legitimacy of a plan-driven model, these are instances when, for whatever reasons, people intend to perceive their actions matching remembered or imagined patterns of movements. When they do, other consequences of their actions will, like the position of the cursor in the present demonstration, "go out of control."

How Could behavioral Scientists Overlook the Fact of Control?

I have shown that, depending on which perceptions a person controls, an observer can mistake the person for a stimulus-controlled system or a plan-driven system. That is one reason behavioral scientists might have overlooked the phenomenon of control. There are other reasons, and perceptual control theorists have described some of them.

For one thing, when scientific psychology began in the 1800s, psychologists followed a tradition several centuries old. They assumed that the lineal models of cause and effect explaining the actions of inanimate objects also explain human behavior. But as William Powers has written, the "orderly march of cause and effect from stimulus object to sensory receptor, and from muscle tension to the eventual behavioral result, does not exist" (Powers, 1973, p. 4). Powers described a fact that sometimes makes it difficult for informed observers to see the phenomenon of control and virtually guarantees that uninformed ones will not:

> In general an observer will *not*, therefore, be able to see what a control system is controlling. Rather, he will see an environment composed of various levels of perceptual objects reflecting his own perceptual organization and his own vantage point. He will see events taking place, including those he causes, and he will see the behaving organism acting to cause changes in the environment and

the organism's relationship to the environment. The organism's activities will cause many changes the observer can notice, but what is controlled will only occasionally prove to be identical with any of those effects. Instead, it will normally be some function of the effects, and the observer's task is to discover the nature of that function (1973, p. 233, emphases in the original).

Powers has written much more about those ideas (see Powers, 1989, 1992). So have other perceptual control theorists. One of them, Wayne Hershberger (1987a, 1987b, 1988, 1989, 1990), has discussed the idea that when an organism controls its perceptions, observers often notice overt actions that seem either elicited by antecedent environmental stimuli, or emitted from within the organism Psychologists have treated elicited and emitted behaviors as distinct from one another and governed by different "laws"; they sometimes call elicited actions "involuntary" and emitted actions "voluntary." Hershberger emphasizes the fact that organisms voluntarily control many of their perceptions of environmental variables by using involuntary actions to eliminate effects of environmental disturbances acting on those variables. The illusory exclusivity of the two "classes" of behavior makes it difficult for many observers to notice that the organism is a controller.

In a series of ingenious experiments, Richard Marken (1982, 1989, 1992) has illuminated another point made by Powers: when an organism voluntarily controls its perceptions, its actions simultaneously produce many unintended consequences. It is not always obvious which of the many variables an organism affects are "under control." Marken has shown the procedures that an observer must follow to distinguish between intended and unintended consequences of behavior—between controlled and uncontrolled states of the environment.

Marken (1993) also has shown several circumstances where an observer can mistakenly think a perceptual control system is a reflexive stimulus-response system, or a reinforcement-controlled system, or a cognitive system. Mistakes like these are behind many lineal causal models in behavioral science, and they guarantee that scientists will "miss" the fact that organisms control many of their own perceptions. Marken suggests that theorists who advocate any of the three mutually exclusive lineal causal models are similar to the three legendary blind men who encountered an elephant: each observes part of the phenomenon of control, consequently, their various interpretations of *the phe*nomenon are incomplete and incorrect, but understandably so.

Conclusion

In the present demonstrations, a person and a PCT model emulated,

or "acted like," lineal causal models used in nearly all behavioral theories. Similarly, in laboratories and clinics, people emulate nearly any kind of system a scientist thinks they should be. For more than a century, the clinical practices, research methods, and theoretical preferences of behavioral scientists have guaranteed they would not discover this obvious fact: a person is one kind of "thing" that an observer can mistake for any of the many kinds behavioral scientists have imagined. Every person controls perceptions; perceptual control theory explains and predicts the control of perception, even when a person impersonates a lineal causal system.

Appendix

The following behavioral models are from the paper "Models and Their Worlds" (Bourbon & Powers, 1993).

The S-R Model

From the person's data during the run in Figure 1, we calculated the slope (m) and offset (intercept, b) of the regression of the handle on the target. Target position is t and handle position is h. The S-R model for the person consists of

$$h := mt + b$$

and

$$c := h + d.$$

Target position, an independent variable, determines handle position, as a dependent variable. This model represents pure environmental control of behavioral actions.

The Plan-Driven Model

The plan-driven cognitive model "remembers" the average pattern of target movements during the run shown in Figure 1, then "computes" handle movements that perfectly match those target movements. The resulting model consists of

$$h := H$$

and

c: = h + d.

In this model, a computed representation (H) of the pattern of previous target movements (t) causes the handle to move in a pattern identical to that of the computed representation.

The PCT Model

The computational steps for the PCT model are

p: = c - t,

error: = p* - p,

h: = h + k • error • dt,

and c: = h + d,

where p is the perceptual signal, and p* is the reference signal or intended value of p. In "Worlds" (Bourbon & Powers, 1993, p. 61), we explained k, the integration factor that resents the velocity of handle movements when there is error, and dt, the sampling interval (here, 1 / 30 second). The reference signal specifies the perceptual signal; if they do not match, the resulting error causes handle movement.

Acknowledgements

I thank Richard Marken and Greg Williams for their critical reviews of versions of this paper, and William Powers for his encouragement when I designed the demonstrations.

References

Bourbon, W. T. (1993a). On the accuracy and reliability of predictions by control-system theory: Data gathered five years after the predictions. Experiment conducted at the annual meeting of the Control Systems Group, Durango, Colorado, 27 July-1 August 1993.

Bourbon, W. T. (1993b). A modest proposal to all behavioral scientists. Paper presented at the annual meeting of the Control Systems Group, Durango, Colorado, 27 July-1 August 1993.

Bourbon, W. T., Copeland, K. C., Dyer, V. R, Harman, W. K., & Mosley, B. L. (1990). On the accuracy and reliability of predictions by control-system theory. *Perceptual and Motor Skills*, 71,1331-1338.

Bourbon, W. T., & Powers, W. T. (1993). Models and their worlds. *Closed Loop, 3*, 47-72.

Hershberger, W. A. (1987x). Some overt behaviors are neither elicited nor emitted. *American Psychologist, 42,* 605-606.

Hershberger, W. A. (1987b). Of course there can be an empirical science of volitional action. *American Psychologist,* 42, 1032-1033.

Hershberger, W. A. (1988). Psychology as a conative science. *American Psychologist, 43,* 823-824.

Hershberger, W. A. (1989). The synergy of voluntary and involuntary action. In W. A. Hershberger (Ed.), *Volitional action: Conation and control (pp.* 3-20). Amsterdam: North-Holland.

Hershberger, W. A. (1990). Control theory and learning theory. *American Behavioral Scientist, 34,* 55-66.

Marken, R. (1982). Intentional and accidental behavior. *Psychological Reports, 50,* 647-&50.

Marken, R. S. (1989). Behavior in the first degree. In W. A. Hershberger (Ed.), *Volitional action: Conation and control (pp.* 299-314). Amsterdam: North-Holland.

Marken, R. S. (1992). *Mind readings: Experimental studies of purpose.* Gravel Switch, KY: Control Systems Group.

Marken, R. S. (1993). The blind men and the elephant: Three perspectives on the phenomenon of control. *Closed Loop, 3,* 37-46.

Powers, W. T. (1973). *Behavior: The control of perception.* Chicago: Aldine.

Powers, W. T. (1989). *Living control systems.* Gravel Switch, KY: Control Systems Group.

Powers, W. T. (1989). *Living control systems II.* Gravel Switch, KY: Control Systems Group.

The Control Systems Group is a membership organization which supports the understanding of cybernetic control systems in organisms and their environments: *living control systems*. Academicians, clinicians, and other professionals in several disciplines, including biology, psychology, social work, economics, education, engineering, and philosophy, are members of the Group. Annual meetings have been held since 1985. The CSG Business Office is located at 73 Ridge Pl., CR 510, Durango, CO 81301; phone (303)247-7986.

The CSG logo shows the generic structure of cybernetic control systems. A Comparator (C) computes the difference between a reference signal (represented by the arrow coming from above) and the output signal from Sensory (S) computation. The resulting difference signal is the input to the Gain generator (G). Disturbances (represented by the black box) alter the Gain generator output on the way to Sensory computation, where the negative-feedback loop is closed.

Closed Loop

Journal of Living Control Systems

Winter 1994

Volume 4

Number 1

Front cover

Closed Loop

Journal of Living Control Systems

Winter 1994 Volume 4 Number 1

Edited by Greg Williams, 460 Black Lick Rd., Gravel Switch, KY 40328

CONTENTS

Members of the Control Systems Group receive Closed Loop quarterly. For more information, contact Ed Ford, 10209 N. 56th St., Scottsdale, AZ 85253; phone (602)991-4860.

Some *Closed Loop* articles are adapted from the electronic mailing list CSGnet. To subscribe to CSGnet, e-mail Gary Cziko at g-cziko@uiuc.edu or phone him at (217)333-8527.

From the Editor

Closed Loop begins its fourth year with a new subtitle, *Journal of Living Control Systems,* which replaces *Threads from CSGnet.* I hope that this will signal a shift in the contents away from CSGnet conversations, toward research reports, clinical studies, and review articles. My ultimate goal is a peer-reviewed journal which includes contributions to the science of living control systems written by both Control Systems Group members and non-members. To start moving toward that goal, I need submissions from *you.* Why not send a manuscript *before* I come asking you personally? Consider yourself warned!

This issue contains articles by seven "pioneering" (in my estimation) perceptual control theorists, all of whom have been studying and using (and, in some cases, inventing) perceptual control theory for several years. I asked them to write on the theme "PCT: Looking Back, Looking Forward," and, in my opinion, they have done so quite admirably. Now I am sorry that I didn't ask them to write on "Solving the Problems of the World with PCT" Of course, there's always a next time....

How Perceptual Control Theory Began: A Personal History

Mary A. Powers
(73 Ridge Place, CR 510, Durango, CO 81301)

The beginnings of PCT lie in two major developments of the 1920s and 1930s: H. S. Black's concept of negative-feedback control in electronics and Walter B. Cannon's concept of homeostasis in biology. These were brought together in the early 1940s, primarily by Norbert Wiener, a mathematician, Julian Bigelow, an engineer, and Arturo Rosenblueth, a co-worker with Cannon. In 1943, they published the first paper relating engineering control theory to neurophysiology.

Although Wiener called his 1948 book *Cybernetics, or Control and Communication in the Animal and the Machine* and believed in the importance of control theory as a key to explaining some phenomena of living systems, he was far more interested in communication engineering and information theory. This bias was shared by the participants in the 10 Macy Conferences that preceded and followed the 1948 publication of *Cybernetics.* Many of the people who attended these conferences were prominent figures in the life, social, and behavioral sciences, mathematics, physics, and philosophy. Though officially titled "Cybernetics: Circular Causal and Feedback Mechanisms in Biological and Social Systems," the meetings were primarily concerned with issues of information and communication.

Enter Bill Powers, an ex-Navy electronic tech and college physics major in his 20s (hardly the sort of person who got invited to the Macy Conferences) who had then what he has now: an irresistible urge, when confronted with something unfamiliar but interesting, to grab a pencil and a piece of paper and start figuring it out. What was interesting to him in *Cybernetics* was not communication, but rather the idea that the nervous system seemed to be a control system. He thought this was an enormously exciting idea, and he couldn't wait to see where the big scientific guns and gurus would carry it. Because he couldn't wait, he started figuring it out for himself, but he was sure for many years that someone far more competent than he would be coming along with a more extensive and profound analysis. That someone, we now know with 20-20 hindsight, turned out to be himself, 20,30, and now 40 years later.

Why Bill Powers? My purpose here is to suggest a few of the variety of characteristics and circumstances that made him uniquely the person to develop PCT. Since he is a private person, I intend to avoid getting too personal, with the one exception that for the rest of this paper I'm going to call him Bill.

One place to begin is with the satisfaction Bill has always found in figuring out how things work, as mentioned above. This contributed to a professional career at a technical level, with little aspiration to rise beyond the actual hands-on design and construction of control systems and development of computer software into the heady realms of administration and paperwork. And his real career lay elsewhere—since working at the lower levels of an organization means, usually, being able to walk out the door at five o'clock and leave it all behind, the evenings and weekends that others might have spent furthering their professional ambitions were free for PCT.

But sticking to the technical level also meant looking at a lot of emperors and finding them naked. There is a good deal of difference between talking about control systems metaphorically, philosophically, and theoretically, and dealing with them on a practical basis, when you're in there soldering wires and making the damned things work. And Bill made a number of control systems work very nicely indeed.

While this sort of experience contributed to the solidity of the foundations of PCT, PCT at the same time contributed to Bill's successful design of control systems: he would imagine "taking the point of view" of the control system he was designing—if I were this system, what would I be able to perceive, what would I need to perceive, what would "really" be going on? This worked so well that he was convinced he was cheating, fudging over gaps in his technical expertise by using control theory as he was developing it to explain living systems (of course one person's cheating is another person's insight).

Another circumstance fostering Bill's approach to living control systems was his coming of age in the era of analog computers. The digital computer as a metaphor for the workings of the nervous system was immediately more attractive to many than the telephone switchboard it replaced, but in Bill's eyes, it is false at its base. His programs, although digital, are designed to simulate the actual analog functions of the brain, not, as in Artificial Intelligence, to produce brain-like results by whatever means. The contributions his analog models might make to neuroscience have yet to be explored.

While Bill wanted his model to be plausible and workable from the physiological ground up, his main interest was psychology. What he knew of psychology when he began was whatever was taught in undergraduate courses around 1950. Behaviorism held the high ground as far as psychology as a science was concerned. The therapeutic community was largely Freudian, with a dash of humanistic psychology—Carl Rogers and Fritz Perls, and later Abraham Maslow. The treatments available for psychosis were lobotomies and electric shock. Bill was interested in psychology for personal reasons, as almost everyone is, and like many young engineers and other technically inclined people, he discovered what seemed to be a far more fruitful approach in the pages of what for many of us was our favorite magazine, *Astounding Science Fiction*. Many people were drawn to Dianetics because, unlike behaviorism, it didn't try to do away with the mind; in fact, *all* the action was in the mind, accessing and dealing with memories, in a very straightforward and routinized manner. There was an appealing technical flavor to it. Like others who went into Dianetics, Bill got out when L. Ron Hubbard's grandiosity, greed, and paranoia turned off youthful enthusiasm, and when it seemed that this "new science of the mind" was not all that it was cracked up to be.

Soon, the first wave of disillusioned Dianetikers went back to work or school and went on with their lives (I kept running into them at the University of Chicago in the early '50s, and there are four—that I know of—alive and well in the CSG). Bill, who had read *Cybernetics* by that time and thought it to be a much more promising approach than Dianetics had turned out to be, went to work as a medical physicist, and he discovered to his delight that his bosses knew about, used, designed, and could teach him about control systems. This means he did not approach the subject of living systems from the point of view of a control engineer, but rather as a student of control theory, applying what he was learning to both artificial and living systems at the same time. This, I think, is the source of his realization that the reference signal, which in artificial systems is set externally to the system and labeled "input," is, in living systems, internal, and not an externally accessible input at all.

Together with Bob Clark, another physicist, and later Bob McFarland, a psychologist, the first model of hierarchical living control systems was worked out. It was published in 1960 as "A General Feedback Theory of Human Behavior," which presented a six-level hierarchical model. By this time, Bill had left his job and begun graduate work in psychology at Northwestern University, and the association with Clark and McFarland ended. The graduate work ended after one year, done in by total incomprehension on the part of the faculty as to what on earth Bill's master's thesis proposal was about, by wifely financial panic, and by an appealing job offer from the Northwestern astronomy department.

"Feedback theory" was the name of the game for many more years, as a book slowly took shape, was dropped into the wastebasket, was written again, and then again. As this went on, the emphasis shifted from the one immediately obvious component that makes control systems unique, namely feedback, to the overall system of which feedback is a part, and ultimately to that aspect of a living control system that makes it so radically different and so difficult to understand, the control of perception. The only possible way to know what is happening, or what one is doing, or the effects of either on the other, is by perception.

How does a person entirely alone develop a science, without money, a lab, or colleagues? One answer, of course, is that all the equipment was readily at hand. Between children, a dog, a clunky computer, and above all, himself, there was more than enough to observe and think about. The nature of much of that observation was unique, however, and involved a form of introspection in which one does not think about thoughts, but about what one is seeing: What perceptions are necessary to see an object, or movement? From what perceptions does logic emerge, or principles? Thus the six levels of 1960 became nine by 1972, and they now number 11. Bill is the first to admit that the levels he sees are personal, and possibly idiosyncratic, and it is with some dismay that he sees them taken as a final word on the subject, copied down and memorized. But the main point here is that the levels, and much else about PCT, were derived from experience; the theory had to explain not just the performance of subjects, of others, but how the world looks from the only point of view available to anyone, from the inside.

The main thing that Bill has been able to bring to his work, then and now, is a mind with no strings attached except his own initial feeling that control theory could answer some of his questions. It is from that stance that he has read books, taken courses, and otherwise absorbed what was already available in the life sciences. Learning what other people have done has never meant accepting either their premises or their conclusions. As an outsider, he has never had to conform to any particular school of thought or please any particular community of scholars. When confronted with such pressure (as with his master's thesis), he has simply walked away and continued on his own path.

I think it took many years for Bill, and for the other people who have become committed to PCT, to fully realize how radically different control theory is from the rest of the behavioral sciences. There is something about PCT that offends just about every point of view: behaviorists, cognitive scientists, dynamic systems analysts, roboticists, cyberneticists, and even control engineers seem equally unimpressed, or baffled, or annoyed. Well-meaning attempts to integrate control theory into the mainstream have succeeded only in confusing the issue with inaccuracies and gratuitous embellishments. The concept of PCT is expressed as principles which contradict many fundamental assumptions: that behavior is the end point in a chain of events, that

the brain calculates necessary outputs, that the concept of purpose is unnecessary to explain behavior, that reference signals (if they exist at all) can be imposed from outside, that feedback can be given or withheld, that self-regulation is a conscious process only and has nothing to do with homeostasis, and so on, and on.

In 40 years, Bill and his colleagues have developed a rich and comprehensive theory which encompasses and resolves many issues in the behavioral sciences. I will never forget the astonishment, joy, and relief on Bill's face as he looked around at the people gathered at the first CSG meeting in 1985, when he really felt for the first time that control theory was not a lonely and eccentric obsession, but rather a shared enterprise that might, just might, change the behavioral sciences forever. That hope, unfortunately, is still discouragingly far from being fulfilled, but at least it is clear that PCT no longer exclusively depends on the unique life, talents, and circumstances of a single person.

Perceptual Control Theory: Origins, Development, Future

Robert Kenley Clark
(834 Holyoke Dr., Cincinnati, OH 45240)

Preface

The following paper is a condensed summary of my experiences and applications of Perceptual Control Theory from the time I first met Bill Powers, during the period of our collaboration, as well as our initial association with R. L. McFarland. My later separate work with McFarland, Richard J. Robertson, and others is included, as well as a similar condensation of my independent activities through later years.

Having participated in the original development of PCT, I have continued to apply and develop related concepts. During most of these years, I was out of touch with the early group, but I continued to work with these concepts informally on my own. My contact with these early associates appeared to have been completely and permanently broken. Since my employment was unrelated to behavior theory, none of these ideas was written, presented, or published.

Those familiar with PCT can generally infer much of how I would have been applying those concepts and methods. These informal applications continue. I think this is about what all of us expected from the beginning of our association. At that time, none of us thought we had “final answers.” I believe this remains our mutual orientation.

I feel that my general success throughout these twists in my activities is due, at least in part, to my familiarity with those ideas now labeled “PCT.” Those insights Bill and I shared from the earliest days also continue to provide much of my basic orientation.

Before VA Research Hospital

Perceptual Control Theory began when Bill Powers and I met in the early 1950s. Our physical science backgrounds, interests, and orientations resulted in an “instant fit.” We both read science fiction and had been impressed by L. Ron Hubbard’s Dianetics. When a friend learned of my interest, he suggested that I contact Bill. At that time, Bill was actively working on the application of Dianetics.

My own orientation was derived from my family, childhood, and schooling. My father, a professor of Classics, a humanist, and a philosopher, was highly verbal—with a general questioning and skeptical orientation. (A great admirer of Socrates.) My older brother and our mutual friends were also intellectually and academically oriented.

In high school, I belatedly discovered that my peers did not share my interests. This roused my curiosity, so I undertook to learn more about their interests.

At the university, this interest developed in the form of participation in the operations of a social fraternity, various clubs, theater, and other organizations. I helped form a discussion group concerned with international affairs.

My interest became more focused when an English composition course required a "source theme." Ambitiously, I decided to write about my opinion of modern psychology. I reviewed several current books: Gestalt, Behaviorism, Psychoanalysis, Extroversion/Introversion, and others. I found little of interest in most of them—my theme was clearly negative. Later, I found I could attend a course in Abnormal Psychology as a non-credit listener without taking the General Psychology course. This included a visit to a mental hospital. On taking the exam (without having bothered with the reading assignments), I would have rated a C grade. This would have been very hard to do in the physics-mathematics-chemistry curriculum I was following. While my interest in behavior continued to increase, my interest in current psychological theories practically disappeared.

Having completed my triple science major, I entered the graduate physics program at the University of Illinois at Urbana. Before completing my Ph.D., I moved to the Columbia Presbyterian Medical Center in New York, where I joined the Radiological Physics Laboratory. Here I continued preparation for my Ph.D. and performed my thesis research: "Absolute Determination of Beta Ray Activity."

During this period, I had an opportunity to attend a class in "Memory and Concentration." The teacher, Bruno Furst, was the author of *Stop Forgetting*. In this course, I learned a great deal about how memory works, and how one can intentionally put it to work for one's own purposes—a course, in my opinion, that should be introduced very early in the teaching/learning process.

When I met Bill, I found he was then involved in activities related to Dianetics. I still have a couple of papers he wrote at that time, as well as some of the Scientology materials then current. We were repelled by the developing transition of Dianetics into Scientology and other less realistic areas. But our mutual orientation and interests gradually developed into an intensive collaboration. It was never clear to me which of us came up with which proposals—it never was of any importance.

Our first experiment was proposed by Bill. It was generally known that a sudden loud noise results in a reduction of electrical skin resistance. It occurred to Bill that the galvanic skin response (GSR) could be connected to form a positive feedback loop. At the time, Bill had access to suitable equipment and space (at the Cancer Hospital) to put this together and try it out.

Bill assembled the equipment, and the first trials showed the beginnings of oscillation. But I (as a subject) soon found I could more or less "ignore" the noise, no matter how loud it went. In today's terms, my higher-order systems were able to modify the sensitivity of some of the systems involved. So we demonstrated—to each other—that the GSR is not a purely automatic response.

In the process of working with this equipment, I was able to learn partial control of my own skin resistance. I can voluntarily reduce it, but then have to wait awhile for it to return to higher levels. I still have this ability.

VA Hospital—Three of Us

In 1954, I learned of an opening at the VA Research Hospital in Chicago. I was able to arrange for Bill to join me, together with an outstanding tool and instrument maker and a secretary. We reported to Therapeutic Radiology, Diagnostic Radiology, and Radioisotopes. We had excellent shop and electronics facilities.

I continued my "outside interests," which in the long run contributed perspective regarding the behavior of "ordinary people." These interests included: (1) observing hospital management (manager's weekly meetings), working with purchasing, personnel, etc.; (2) Health Physics Society (second president of the Midwest Chapter); (3) National Health Physics Symposium; (4) Association of Physicists in Medicine; (5) other physics-related activities—I helped organize Radiation Control, Inc., and was the president of the company during most of its life.

The Physics Unit at the VA Hospital included primary responsibility for overseeing the installation, operation, and safety of a unique Cobalt 60 teletherapy unit. This later led to our learning computer programming. After work, we drove up to the Evanston campus of Northwestern, where there was an IBM 650 that Bill had arranged for us to use. We would have dinner and then spend a long evening working with the computer. (Home around 12:30 or 100!)

Bill and I continued our close collaboration, gradually developing a theory of behavior on which we generally agreed. There were always some points where we were not fully satisfied, and we frequently exchanged ideas. Our exchanges were so frequent and informal, it was rarely clear to me which of us came up with any particular idea.

However, I am sure that Bill first pointed out the relevance of the negative-feedback control concept and examined its history. Further, he suggested the concept of an interlocking hierarchy of simple versions of such control systems. If such systems are successful in their operation, they would always be close to their normal condition, so linearity would be a valid first approximation.

I'm not sure just when, but soon after joining the VA, we met Bob McFarland, a Ph.D. clinical psychologist in the Neuropsychiatric Service. This was the beginning of a three-way collaboration (Bill, me, and McFarland) that lasted until Bill left (1961). McFarland and I continued at the VA through most of 1963. McFarland had other responsibilities and outside activities, so his participation in discussions of theory was limited.

While joining in our discussions, McFarland's primary contribution was finding ways that our work could be presented. During these years, I, and usually all three of us, attended meetings and presented papers. Informal presentations were not uncommon. However, on several occasions we were specifically invited to present papers.

The first major presentation of what came to be called 'PCT" was a symposium arranged by McFarland at the August 1958 meeting of the American Psychological Association. Each of us presented papers and were joined by J. Arbit and C. Van Buskirk. During the preparation, we reviewed each other's papers and offered suggestions, as Bill and I had done from the beginning of our association. This practice continued throughout our collaboration.

Throughout our affiliation with the VA Hospital, we made presentations at meetings and published papers (see the appended lists of meetings and papers).

VA Hospital–Two of Us

After Powers left to pursue his own objectives (1961), McFarland and I continued to work together. I transferred to the Neuropsychiatric Service, working more closely with McFarland. About this time, Dick Robertson became interested in applying these developing ideas to his own interests.

McFarland arranged for me to participate as co-therapist with a group of VA patients. A couple of events that impressed me were these:

One day, Bob announced that I would have to conduct the meeting alone—he had other commitments. I was concerned that I might "somehow do some harm." Bob pointed out that it is at least as hard to make a mental patient worse as it is to make him or her better, so that I need not be concerned about accidental "errors."

A patient reported repeated problems with another patient (not present) at a water cooler. Suggestions that she try a minor change in the way she spoke to him were immediately and strongly rejected. Indeed, a general characteristic of the group was the rejection of any and all attempts to try something different. This seemed likely to be one result of being locked in some form of internal conflict, while including something that precluded operation of the Reorganizing System. Of course, this is not necessarily true of all mental patients, there are many other possibilities. However, there seemed to be an implied "I am what I am, and that's all I am." Even the least attempt to change was avoided.

A study on the effects of drugs on "laboratory induced anxiety" (conceived by the Chief of the Service) led to my learning more statistics than I had needed for physics. "Statistical inference" was included. I found "non-parametric statistics" quite a different and interesting basis for experimental design. Two papers on computerization resulted from this work.

The first of these was presented at the annual meeting of the Society for Psychophysiological Research in October, 1962. This was held in Denver, which made it convenient for us (McFarland and me) to visit Jay Shurley's Sensory Deprivation Facility in Oklahoma City. The subject floated in a skin-temperature water tank in a dark soundproof room. Excellent design. Both Bob and I tried this out. As I recall it, the conclusion was that after sleeping a bit, the subject's expectations largely determined his experiences. I had the odd experience of the "disappearance" of my left arm from the wrist to above the elbow. It reappeared on the slightest movement, so I concluded that the water temperature matched the skin temperature of this part of my arm so closely that I no longer had any sensory inputs from that area. With muscles relaxed and joints stationary, there were no other signals available.

The second paper was presented at the New York Academy of Sciences Conference on Computers in Medicine and Biology in May, 1963.

The "anxiety" project was the basis of seeking funds for a computerized data gathering and reducing system. Although McFarland had been informed by Central Office that funds would be approved, he learned that the Chief of the Service had rejected them without conferring with either of us. We both resigned in protest.

The Human Systems Institute

In 1963, we formed the Human Systems Institute. This was intended to be a Tax Exempt Organization, but, as I learned much later, we did not know how to get IRS approval—and our attorney was of no help! This was the year that McFarland was the president of the Illinois Psychological Association.

During this period, we had several interesting projects.

For Illinois Bell, we analyzed a management position with a view to providing a training program at some later date. When we reported that this position was obsolete and unnecessary, the training program was dropped. Meanwhile, this produced some income for us.

Another project was computerization of a Career Profile Test for one of Bob's psychologist friends.

In addition, we submitted a grant request in collaboration with IIT Research Institute—an education in grantsmanship!

Also, through his connections, Bob arranged clients for a project in the therapy of adult stutterers. This led to meetings with several speech therapists, as well as a presentation at a national meeting of speech therapists. Interestingly, in the informal meeting afterward, one of their members asked if I had been a stutterer. "Not to my knowledge." He then stated that he was a "cured" stutterer, and that my purely theoretical analysis fitted his experience exactly! Theory confirmed by experience!

In 1965, funds were running out, so HSI had to terminate. I returned to teaching as Associate Professor of Physics at the Chicago Circle Campus of the University of Illinois.

The Mosier Safe Company

In 1968, I improved my economic condition by moving to Cincinnati and joining the Mosier Safe Company, becoming Manager of Applied Research. My primary responsibility was the development of an automated identification system. In this connection, I investigated signatures, voiceprints, and fingerprints, as well as several other concepts. I participated in evaluation of other systems that were offered to Mosier.

In addition, I supervised Computer Applications, Test & Evaluation, Materials Lab, and Special Projects. I brought the heads of these several groups together for discussions of the various projects and related matters. I was surprised to find that some of these people expected me, as "Boss," to "know all the answers." However, they generally seemed pleased that I recognized their capabilities and respected them as individuals. (PCT paid off in terms of general attitude and cooperation.)

Assembly/editing of "Technological Forecasts," written by the engineers of R&D, was also my responsibility.

When I joined Mosier, the management was in the process of implementing the 'Profession of Management Program," produced by some industrial psychologists (names no longer available to me). This

was a very ambitious program, taught first to top management, who then taught it to their subordinates, and so on down the line. This material was based on the concept that management is, somehow, an identifiable skill that can be learned independently of other knowledge and skills. Thanks in part to my PCT-related background, I was able to review this rather extensive material in a couple of months of spare time. It was pretty much a mix of the obvious and the unnecessary. For example, they placed great emphasis on "communication." Sensible enough; but they neglected to indicate when, where, and to whom to communicate about what. It was also pretty clear that an effective manager must know quite a bit in addition to general principles of management. The behavior of the old-timers was as might have been expected: they learned the special language and could recite it when necessary. But they rarely made any application to their previous methods.

I was able to attend additional internal courses in management and finance. It was interesting to see how the accounting was handled in the transfer of income from Mosler to its owner, American Standard.

I was quite surprised at a pricing decision that was made after a very coherent presentation by Mosler's quite competent marketing people. After seeming to understand the survey data and the logic that clearly demonstrated an optimum pricing strategy, the key vice-presidents went for a minimum initial price! And the analysis really wasn't that difficult.

Perhaps the closest I came to direct application of Perceptual Control Theory was in solving a paint matching problem. Here I learned more than I really wanted to know about paint manufacture and application. The problem was that furniture for bank lobbies was manufactured in plants in two different cities. The color depended on several interacting variables. To control each of them in proper balance would have been unmanageable. One of the important variables was adequate stirring of the paint, which was controlled by the individual painters. My solution was to provide sets of reference chips to each painter and to their quality control people. This, of course, was the right concept, but incomplete. Another level of control was needed. Therefore, I required painting a test chip in each production batch. This freshly painted chip was to be compared to the local quality control reference chip, then forwarded to me for final approval. So far as I ever learned, this seemed to have solved the problem.

Is this a two- or three-level hierarchical control system? Where is the test of the controlled variable? As long as the completed units match, there is no disturbance. The controlled variable is revealed only when the completed units fail to match.

Mosler's Central R&D was terminated in 1976.

Being a bit old to find a job in technical management, which would have been my preference, I accepted an opportunity to join "DFS," Diversified Financial Services, Inc., a small company. Max Redlich, the owner, had a background in life and health insurance, as well as the retail furniture business. As a Financial Consultant, he had several clients who were more or less on the edge of bankruptcy. My part was helping with the handwork, accounting, etc. There were several interesting experiences during this period, without a direct relation to PCT, but relating to people—where PCT is always helpful. Later I became a co-owner, and a third person was added. We were working with a computer programmer to develop a program for retail business management. But a key client, refusing to accept our advice, was forced into bankruptcy. Here I learned a good bit about how bankruptcy works: how one can deal with the IRS, how to work with banks in refinancing loans, etc. A PCT background is very helpful in understanding the interactions among people in these sometimes-tense situations.

When we found we had gone too far for a specific client, it became necessary to dissolve the company. Max returned to his initial field, life and health insurance, as well as pension planning and administration. 1 stayed with him to help with the planning and administration.

During this period, I "officially" retired—that is, I started to receive my Social Security benefits. So that I could work with Max as an independent consultant, my wife Mary Ann and I formed an S-Corporation in 1980. Nothing else really changed.

When the insurance agency we were affiliated with suddenly had its General Agent replaced, Max and I, as individuals, both moved to The Lincoln National Life Insurance Company. I was to provide a communication link between the home office and agents having pension clients. To me, it looked like an unnecessary linkage, but then I was new to the agency system. And there were funds available.

I learned a lot of things about the insurance business and had a lot of interesting experiences with insurance people.

The first clue I got to the agency system was when I asked for my job description. The General Agent tried to provide one, but it was clear that he was entirely unacquainted with the concept. It became clear that I was really expected to sell pension plans—but that was never spelled out.

1 found that the agency system is not a hierarchical structure. Rather, the General Agent (who is paid on the basis of the production of his affiliated agents) provides facilities and services to the individual agents. The agency provides forms, advertising, sales materials, etc.—produced by the home office. In some cases, the General Agent charges the affiliated agents for the use of office space and/or other items. But he gives them no direction beyond basic "training" in the company's requirements and the use of the many forms. The agency might provide additional seminars in sales presentations and methods.

As a form of sales training, the agency brought in a consultant who presented a series of seminars and workshops produced by the Wilson Learning Corporation of Eden Prairie, Minnesota. This series was based on a two-dimensional classification of "Social Styles." I found it generally consistent with PCT, but the basic concept of feedback control was only included by implication. This approach could be considered for revision and, possibly, integration with PCT's higher orders.

During this period, I completed both the Chartered Life Underwriter and the Chartered Financial Consultant courses. These included about a dozen quite respectable college junior- and senior-level one-semester courses covering such topics as accounting, taxation, economics, investments, and other appropriate topics.

While associated with Redlich, I had become fully licensed for sales of life and health insurance, as well as pension plans. I later qualified for sales of plans having investment aspects. I still retain this licensing. In addition, I became a "Registered Representative" with Lincoln National for selling mutual funds. I later discontinued this because of the continuing paperwork required.

I also learned about the selling process—from the agent's viewpoint. Essentially, an insurance agent is a true entrepreneur. And his/her most important asset is at least 200 satisfied clients. While an income is obviously essential, the most effective agents are those who enjoy providing for their clients' desires. Their clients think of them as friends rather than salespersons. The successful agent accepts the prospect's solution and sells him/her the implementation. (If the customer wants a Cadillac, you don't try to talk him/her into a Volkswagen—no matter that it would be much better for him/her.)

However, my personal orientation tended to be too much one of trying to solve the prospect's problem and then selling him or her my solution. At this point, I ceased active sales efforts.

AARP

While I was still involved with insurance, I learned that the AARP (American Association of Retired Persons) includes Chapters. I joined one nearby. Their newsletter was being distributed at meetings instead of being mailed. I joined with others to sell advertising for non-profit-organization postage. With a mailing list of over 200, this helped increase the membership and meeting attendance. I became the Legislative Chairman, reporting activities of both the State

Legislative Committee and the National Legislative Committee. Later, I was appointed to the Ohio State Legislative Committee. In this capacity, I formed a Legislative Council consisting of the presidents and legislative chairmen of the 19 Chapters located in the five counties in the southwest corner of Ohio. We arranged to videotape the monthly meetings for presentation on the public cable system. In this position, as a Registered Lobbyist, I learned a great deal about the details of the legislative process. I also worked with several of the Chapters in conducting Candidate Forums during election season.

In the course of these activities, I had learned about the internal operations of the AARP. I saw the operation of their structure of volunteer leaders as they were guided by a staff of permanent employees. This was very instructive.

In July 1987, after over three years, my appointment to the Ohio State Legislative Committee terminated. For several reasons, I ceased to work further with the AARP organization, but I do retain my membership.

Contact!

Sometime in the fall of '87, the phone rang—and I heard a voice from the past! Bill Powers had found me in Cincinnati!

We exchanged a few letters in which we discussed some of the ideas I had been developing over the years since Bill had left. I was working from Parts I and H of the 1960 publication, and from the ideas McFarland and I had developed. At that time, I had not even heard of Bill's 1973 book, *Behavior: The Control of Perception (BCP).* I was interested in the Orders above Fourth, which were not yet well worked out. From my association with people of highly varied backgrounds, it occurred to me that "Fifth Order," as I had conceived it, could be subdivided into "Modes" corresponding to the various Orders. This was based on classifying the topics of ordinary conversation. The irregular correspondence with Bill did not develop further at that time.

Greg Williams visited me while he was in Cincinnati at Christmas time in 1987. From him, I learned of *BCP* and other events. Greg and I corresponded irregularly for a couple of years. I sent him some of my notes and preliminary drafts, including a discussion of "Fifth Order." He reciprocated with copies of *Continuing the Conversation, Closed Loop,* and various papers. He also sent me a copy of *Living Control Systems: Selected Papers of William T. Powers.* This covered the years 1957 to 1988. Bill sent me a second, autographed copy about the same time. *Behavior: The Control of Perception* arrived from the publisher a bit later.

Through Greg, I learned about CSGnet. I got a modem and connected to the net in September 1991. In December 1992, I posted my

first discussion of the "Decision Making Entity" (DME), including an extension of the concept of "Modes of Fifth Order." My contact with the net continues.

Museum

Meanwhile, in July, 1989, I learned that a small museum was being formed. The Archimedes Rotorcraft and V /STOL Museum involved rotary wing aircraft, in which I had long been interested. On joining the group, I found that they assumed that non-profit organizations are automatically tax exempt. From my insurance studies, I knew that this was insufficient. The regional head of IRS Tax Exempt Organizations was more than willing to give us the guidance we needed. He was particularly helpful with the documentation needed to get the official recognition letter. It took about three months instead of the usual 18 to 24 months!

I was a Board Member, wrote the By-laws, and became the Secretary, Treasurer, and Editor/Publisher of our monthly newsletter. While the membership was not large, it was international in scope. The newsletter included historical notes written by a past Director of Flight Dynamics at the Wright-Patterson Air Force Base. He and I both provided technical papers related to the engineering and operation of rotary wing aircraft. A small gift shop was included.

For two successive years I participated, as Financial Chairman, in planning the Annual Convention of the Popular Rotorcraft Association. I managed the publication of the Convention Program, as well as selling local ads for it.

After about 3 years, I resigned all connection with the Museum. It had become clear that more time was needed than I cared to spend on this activity.

Civic Activities

During the latter period of my insurance activities, and continuing through the AARP and museum periods, I became interested in local community organizations. The first of these was the Forest Park Business Association. I found that this group of business people did little long-term planning of the Association's affairs! After being elected to their board, I worked with another Board Member in rewriting the Bylaws. Later, I dropped the Board Membership, but I still continue my membership in the organization.

When I decided to study the City of Forest Park, I became Legislative Chairman of the Business Association. This facilitated my contacts with the City of Forest Park. For the past year, I have attended all scheduled meetings of the Administrative Staff and the City Council, most of the meetings of the several Commissions, and some of the meetings of the Council's Standing Committees.

PCT concepts and methods have been very helpful in working with these people—with no need to try to teach PCT to them. On several occasions, I have helped resolve developing conflicts. "Moving up a level" is always useful—although not always easy! Another helpful approach is emphasis on "Reality"—that is, current perceptions, otherwise known as: "What's happening *now?"*

A recent campaign for election to City Council was most interesting—I knew most of the 13 candidates and found that they campaigned almost exclusively on the basis of their personalities! Virtually no policy or other proposals were discussed! That is, there was very little consideration of lower-level problems and their relation to higher-level policies.

Throughout this period, I found the interactions among personalities at least as interesting and pertinent as considerations of organization and operation. Personalities and their interactions certainly pertain to the higher levels of human systems.

Recently, I became a member of the Civil Service Commission. I joined this Commission because it works closely with the Human Resources Department and the City Manager in personnel-related matters. It is, therefore, a place where PCT is directly pertinent both in personnel decision-making and in conflict resolution.

Future

When one turns to the future, one finds a mixture of projections and hopes—both truly imaginary.

In seeking to reach a larger and more understanding audience, I expect there will be extensions, elaborations, and modifications of the basic theory. The levels of the hierarchy will be studied in more detail and become more meaningful and available for application. While a specialized vocabulary has its place in technical discussions, expressing the concepts in more ordinary terminology will be necessary for more general understanding. As this is accomplished, both the number of participants and the applications of PCT will increase steadily.

In a view of the future, we find applications that are waiting for study and development. These might include the following:

Economics. Here we can expect to find recognition of the multiplicity of peoples' motives. Projections that go beyond static or linear methods will be examined. Situations where decision-makers have incomplete and/or incorrect information will be analyzed. Alternative specifications for the various levels of the hierarchy—especially those above purely mechanical systems—will be proposed.

Sociology. Group behavior where individuals have differing objectives will be analyzed. How cooperation occurs and how conflicts are resolved will be studied. Various forms of organization will be examined from the standpoint of the effects on the individual's freedom.

Education. Improved understanding of the "learning process" will be developed. Many new methods and procedures are being studied—mostly from the standpoint of teaching, rather than learning. PCT will be particularly helpful in these efforts.

Mental Illness and Psychotherapy. To the extent that such problems are the consequence of inadequate (incorrect?) learning, PCT has already been helpful. As PCT becomes better known in the general community, these important matters will become better understood. PCT methods will also assist in distinguishing between organic and functional problems.

In the course of these developments, I expect to find modifications and/or additions to the basic theory. These might include the following:

Emotions. The relation between imagination and emotion will be clarified. The physiological effects of different emotions—friendship, affection, loneliness, aversion, and many others—will be explored. This important area has barely been touched.

Memory. The relation between memory and the hierarchy will be extended to include those memories not directly used for operation of the control systems. How formation of memories can be improved and controlled will be examined, as will how and why availability of memories can be limited, even while their existence is beyond question.

Conflict Theory. This important topic will be extended to include conflicts between organizations as well as internal conflicts between anticipated (imagined) events. This will help clarify the operation of the DME.

Reorganization. Recognition of the special requirements for this critical process will be clarified. The role of intrinsic error in initiating Reorganization will be analyzed. Relations among memories, the planning process, and the operation of the DME will be clarified.

To accomplish all these developments and applications, as well as others, will take much time and effort. It will require multiple participants, mutual understanding, and cooperation. I anticipate and hope for the development of the kind of interactions I enjoyed in working with Bill Powers, Bob McFarland, and others some 30 to 40 years ago—but on a much larger scale! On with the show!

Meetings Attended as a Speaker:

Symposium—American Psychological Association Meeting, Washington, D.C., August 30, 1958. Papers by R. K. Clark, W. T. Powers, J. Arbit, V. Buskirk, and R. L. McFarland.

Symposium on a General Feedback Model of Behavior, All-University Seminar, Urbana, Illinois, November 1958. Papers by R. K. Clark and others.

Symposium in Neural Mechanisms, Information Theory and Behavior, VA Hospital, Battle Creek, Michigan, March 10 & 11, 1960. Papers by R. K. Clark, W. T. Powers, R. L. McFarland, and others.

Northwestern University Society for Neuroelectrokinetics, Evanston, Illinois, December 1960. Paper by R. K. Clark and others.

Westsuburban Psychologist's Association, Moosehart, Illinois, February 1961. Paper by R. K. Clark and others.

Bio-Medical Engineering Colloquium, Northwestern University, Evanston, Illinois, March 30, 1961. Paper by R. K. Clark.

Symposium—American Psychological Association Meeting, New York City, September, 1961. Papers by R. K. Clark (two), R. L. McFarland (two), and R. J. Robertson.

American Academy of Psychotherapists—Seventh Annual Conference, Chicago, Illinois, October, 1962. Invited paper by R. K. Clark.

Cleveland Physics Society, Cleveland, Ohio, November 1962. Invited paper by R. K. Clark.

Society for Psychophysiological Research—Second Annual Meeting, Denver, Colorado, October 13 & 14, 1962. Papers by R. K. Clark and representatives of other research groups.

New York Academy of Sciences—Section of Biological and Medical Sciences, Conference on Computers in Medicine and Biology, New York City, May 1963. Paper by R. K. Clark, R. L. McFarland, and M. Bassan.

Papers Presented and/or Published:

Clark, R. K., McFarland R. L., & Powers, W. T. (1957). A general feedback theory of human behavior. *University of Chicago Counseling Center Discussion Papers, 3(18).*

Clark, R. K., McFarland, R. L., & Powers, W. T. (1957). A general feedback theory of human behavior: A prospectus, *American Psychologist, 12,* 462.

Clark, R. K. (1958, August). *Conceptual framework of a general feedback theory.* Paper presented at the meeting of the American Psychological Association, Washington, DC.

Clark, R. K. (1958, November). *Verbal structures in a general feedback model of behavior.* Paper presented at the All-University Seminar, Urbana, IL.

McFarland, R. L., Powers, W. T., & Clark, R. K. (1959). A preliminary report on a clinical rating scale to measure participation in group psychotherapy derived from a hierarchical feedback model. *[Baltimore VA Hospital! Newsletter for Cooperative Research in Psychology, 1(4).*

Clark, R. K. (1960, March). *A general feedback theory of human behavior.*

Part IBasic concepts. Paper presented at the Symposium in Neural Mechanisms, Information Theory and Behavior, VA Hospital, Battle Creek, MI.

Powers W. T., Clark, R. K., & McFarland, R. L. (1960). A general feedback theory of human behavior: Part I. *Perceptual and Motor Skills, 11,* 71-88.

Powers W. T., Clark, R. K., & McFarland, R. L. (1960). A general feedback theory of human behavior: Part II. *Perceptual and Motor Skills,* 11, 309323.

Clark, It K. (1960, December). *A general theory of human behavior from the viewpoint of physical science.* Paper presented at the meeting of the Northwestern University Society for Neuroelectrokinetics, Evanston, IL.

Clark, R. K. (1%1, March). *Human behavior as an organization of feedback systems.* Paper presented at the Bio-Medical Engineering Colloquium, Northwestern University, Evanston, IL.

Clark, R. K. (1961, September). *A brief overview of general feedback theory.* Paper presented at the meeting of the American Psychological Association, *New* York.

Clark, R. K. (1961, September). *The group therapy process scale and the personal assessment program.* Paper presented at the meeting of the American Psychological Association, New York.

Clark, R. K. (1962). A general theory of human behavior from the viewpoint of physical science. *Newsletter for Research in Psychology, 4(2).*

Clark, R. K., & McFarland, R. L. (1962, October). *How can the scientist help the psychotherapist?* Paper presented at the Seventh Annual Conference of the American Academy of Psychotherapists, Chicago.

Clark, R. K., Chessick, R. D., & McFarland, R. L. (1962). High speed data processing—Compromises and considerations. *Psychophysiology Newsletter, 8(4).*

Clark, R. K. (1962, October). *Feedback system analysis of behavior.* Paper presented at the meeting of the Radiation and Medical Physics Society of Illinois, Chicago.

Clark, R. K. (1962, November). *A "systems oriented" theory of behavior.* Paper presented at the meeting of the Cleveland Physics Society, Cleveland, OH.

Clark, R. K., & McFarland, R. L. (1963). Systems concept of stimulus.

Perceptual and Motor Skills, 17, 99102.

Clark, R. K., McFarland, R. L., & Bassan, M. (1964). Integrated data collecting and processing systems in psychophysiology. *Annals of the*

New York Academy of Sciences, 115, 905914.

Chessick, R. D., McFarland, R. L., Clark, R. K., Hammer, M. & Bassan, M. (1966). The effect of morphine, chlorpromazine, pentobarbital and placebo on "anxiety." *Journal of Nervous and Mental Disease, 141(5).*

The Early Days of Perceptual Control Theory: One Person's View

Richard J. Robertson
(5712 South Harper, Chicago, IL 60637)

Introduction

One of the net losses of modem civilization, it seems to me, is the loss of adventure in contemporary life. To be sure, in our generation we have had astronauts going to the moon, people living on the sea floor, explorers of new life styles, and pioneers in science, but these are uncommon and rare people. I believe that we have lost adventure in the lives of ordinary people, like the pioneers who pushed toward the ever-receding frontier. Their lives were filled with challenges simply in the course of pursuing a livelihood. That frontier is gone—the tangible, geographic one. What we have left are the pinnacles gained by the few explorers in the fields of endeavor that are not yet purely cut and dried.

So it struck me with quite some surprise, recently, that throughout the years while I have been bemoaning the lack of adventure in modem life, I have been on an exciting voyage of discovery without realizing it. While fiddling with the television controller the other night, I happened to stumble on the movie "Columbus." As I watched Columbus journey from place to place begging for a hearing and being put down by most of the smart and powerful of his time, I began to get a feeling of familiarity. I know a captain and a crew who have had a similar experience. Let me tell you about it.

I seem always to have enjoyed taking the historical approach to understanding things in which I am interested. I want to see what follows from what. Thus, it feels natural to try to pin down why it was that, in 1957, when three guys came to the University of Chicago Counseling Center to give an hour's lecture on a new approach in psychology, I was ripe for that one hour to start gears rolling that would give direction to the rest of my life. It happened in our Thursday afternoon open seminar, during my internship. I had been interested in psychology ever since I read through an ancient text by Pillsbury that an uncle had left lying around our house from his college days, in my first or second year of high school. But in graduate school I had been a somewhat indifferent student. In fact, when I took my M.A. in Human Development in 1952, the department suggested that I should probably regard it as a terminal degree. They said I didn't have much ability to conceptualize.

I felt that they might be right. I was aware that I could not keep straight all the distinctions between positive reinforcers of desirable behavior, and negative reinforcers of undesirable behavior, and positive reinforcement of stopping undesirable behavior, and cessation of aversive reinforcement of desirable behavior, and several further refinements. (I still can't keep them straight, but, thank goodness, now I don't have to.) I knew that I got bored and tended to drift away when trying to memorize these concepts and their definitions.

I had found Freud's ideas more stimulating, for the most part, but I noticed eventually that his "explanations" of behavior take the form: The reason A causes B is that when you observe B on many occasions, you usually find A preceding it. In contrast to this, in my physiology courses I was learning explanations for things like blood pressure in terms of stroke volume, heart rate, and resistances in the system. That struck me as more like what I thought "explanation" should mean. I didn't find explanations like that in any of my psychology courses. I had not been doing terribly well in many of them, either; I seemed to keep asking the wrong questions-like, 'What would be happening in the brain when a person is having a given experience?" So I took the advice of the departmental counselor, and my M.A. and I went out into the business world.

However, after a few years of employment as an industrial trainer and, later, job analyst, I realized that I would soon die of boredom at such work. So I talked my way back into the university, partly thanks to the good graces of Carl Rogers. I had taken a number of his courses and had been intrigued by his point that one learned much more about the behavior of an individual by trying to view the world through that person's eyes, rather than by surmising about what was going on inside by observations from outside.

In his theory of personality, Rogers (1959) declared that we all live in a world of our own perceptions. This idea had a profound place in his conception of what therapy is about-that by reflecting a person's message, you will help him or her see more clearly what reality he or she is perceiving (and coping with). This concept might have helped make me receptive, later on, to the idea that behavior is the control of perception. I'm not sure about that; it seems connected to me. However that might be, what I am sure of is my state of excitement after hearing those three guys talk on a "General Feedback Theory of Behavior" on that Thursday afternoon. A year or two later, when I had my Ph.D.

in hand, I took my first job four-fifths time, so as to be able to volunteer one day a week at VA Research Hospital in Chicago, where Bob McFarland and Bob Clark were working with the theory. Their new conceptions were that alluring.

Early Research on a Feedback Theory of Behavior

Once I had begun to learn the theory from Clark and McFarland, they gave me a job running subjects on a gray box with four red lights and four pushbuttons on the front that Bill Powers had left behind when he moved on to study and work at Northwestern University. This box presented a game that the subject had to learn totally through experience and could only win by stumbling onto anticipating the moves of the machine. After running a subject, I would sit and measure the distance between blips on the readout tape and plot the resulting pattern of reaction times on a graph.

The results often showed a neat, descending curve with three reaction-time (RT) plateaus, reflecting the three levels of skill that had to be mastered in a successful performance: (1) the pattern of key-light connections (control of the finger-position configurations); (2) the order of finger pushes (control of the sequence of events); (3) the plateau of negative reaction times where the subject was anticipating the machine (control of the time relationship).

I say that the results "often" showed this neat picture. But not always. The RT patterns of some subjects were simply chaotic—my term for a graph showing no plateaus, in which the RT's appear to be scattered at random. The same phenomenon occurred years later, when I was teaching at Northeastern Illinois University and ran the experiment again, this time on computerized apparatus that Bill Powers instrumented for me. Once again, there was the same mix of neat three-plateau curves and random patterns. But this time, I had obtained verbalizations from the subjects as they went through the task. Certain of them seemed quite significant. I found that only those with the patterned graph could articulate the concept that the way to win the game was to beat (anticipate) the machine. In both the earlier and later programs, the subjects with random patterns fell into two groups. Either they never did win, or if they stumbled accidentally on a win, they could not say how it worked.

The original research was published in-house in Northwestern University Psychiatry Research Papers. I took an illustrative curve from that study for my article on control theory in the Wiley *Encyclopedia of Psychology* (Robertson, 1984). The replication was published in *Perceptual and Motor Skills* (Robertson & Glines, 1985) as a rebuttal to behaviorist Keller's (1958) attempt to discredit Bryan and Harter's

(1899) description of plateaus in learning telegraphic communication. These experiments, conducted over a period of 20 years with subjects of presumably different demographic characteristics and different instrumentation show what I consider a very robust result. The curve taken from the original data for the *Encyclopedia* article and that taken from the later data for the *Perceptual and Motor Skills* article both show similar plateaus and increasing variability before plateau shifts. Despite different overall length of the graphs, these features look similar, as do also the mean RT's per plateau. That suggests to me that we are observing something fundamental about the way people, in general, learn tasks of this sort. First you learn a set of configurations, then the sequence in which these configurations apply, then you play with the sequence in time until it dawns on you to precede the machine. Reaction times decrease at each step: when you know which key punch turns off which part of the display, you can be ready to strike as soon as it comes on; many subjects try to time it to the instant and thus stumble on anticipation by being too fast.

Early Applications: Clinical and Rehabilitation Work

Bob McFarland suggested an explanation of the other phenomenon depicted in the 'Towers Game" experiments: the characteristic increase in RT variance just prior to the drop to the next plateau. He proposed that the subject was experimenting with his/her performance when mastery of the first level did not result in a win. Later on, in Powers' (1973) discussion of reorganization, we have a theoretical explanation. He states: 'The effects of the outputs of the reorganizing system must be such as to change the properties of behavioral systems... as a result, of course, visible behavior would change its character, as would experienced behavior." (page 185)

Powers made the implication that the organism does not know exactly what change must occur. Random excitation caused by the reorganizing system results in various alterations of action. Then the action that begins to bring the desired objective under control becomes the basis of a new control system. The increased variability of RT's prior to the drop to the next plateau certainly appear to be instances of such reorganization outputs. The subjects, in their spontaneous verbalizations while doing the experiment, say things like, "Hmm, I know which key turns out that light, but do they come on in a fixed order or any old order? Hmm, I think they always come on in the same order. Oh, I've got it, they come on in this order" (rapidly extinguishing the machine display in fixed sequence). Their RT's often slow up while they pause to think, then become very fast when they try out their hypotheses.

At the time I began to think about reorganization during the original experiment, I was involved in my professional life as a clinical psychologist in a rehabilitation hospital. I would often have patients with aphasia or other cognitive deficits whom I would try to help rebuild lost capacities. What better occupation in which to experiment with reorganization? I was also learning from Clark and McFarland that they had had some good results in working with veterans by applying the scheme of the control system hierarchy (the original version, Powers, Clark, & McFarland, 1960a, b). Noting a lack of some particular skill in a patient, they would examine the immediately lower orders in his hierarchy and construct drills on weak aspects of the presumed foundation.

Following their lead, I would approach a patient who, for example, could not draw a straight line, by getting him to make any mark on a piece of paper. Then I would get him to draw the mark from one dot on the paper to another immediately adjacent dot, and then move the dots progressively farther apart. Next we would employ the technique to copying letters of his name, eventually making a signature. The idea of reorganization underlying this is to encourage perceiving the task as a sequence of moves across the whole signature, rather than as isolated drawing movements.

Shortly thereafter, I went into Veterans Administration research, where various small successes continued to show the promise of developing and applying "reorganization theory" (if I may be permitted to glorify it with a name). I will cite two of them which illustrate how just the bare bones of a good concept can lead to useful applications. One day, one of my research assistants came running fearfully into the office saying that a veteran, to whom she had tried to administer our 20-odd-page questionnaire, had chased her out of his room, menacingly. A moment later, he walked in cursing about invasion of his privacy and with a few other complaints about the questionnaire. I noticed the large metal plate in the front left portion of his skull and surmised that as he looked at the thick questionnaire, the idea of replying to it as a whole might have felt overwhelming. I said, "Oh, sure, no problem, but you wouldn't mind telling me your birth date here, would you?" (As I indicated the beginning question.) He complied with my request, and I then asked if he minded telling me the next piece of information, and so on, until we had completed the entire questionnaire.

I cite this as an example of how a good concept/theory provides lines of action that would not otherwise occur to a person. In this case, my impression that this man lacked a lot of computing ability in the brain area that is usually attributed with sequence-controlling properties, plus my experience with the usefulness of "order-reductions" in McFarland's and Clark's training efforts, led to the technique of pointing the patient at each question as a separate event and urging his attention away from the task as a whole.

The other story from this era is about the satisfaction that comes from using theory to make sense out of an otherwise puzzling observation. One of the VA patients happened to be a former state table tennis champion. Under his leadership, we soon had a large number of staff and patients involved in a round of ping pang tournaments. During the tournaments, a number of us noted a curious change that would quite regularly occur in the play of a contestant when he began to recognize that he was clearly outmatched. He would first concentrate very hard, then begin to alternate between wild shots and cautious play. It occurred to me that these periods of variability, if we could graph them, would look like the patterns of RT variation preceding a new plateau in the Powers Game experiment. The participants themselves acknowledged this aspect of play as part of their experiments to obtain eventual increases in skill. In this view, what would have seemed a lapse into sloppiness on the part of a losing player took on an opposite significance.

Powers' Book and Further Developments

I left VA research to take a position in the Psychology Department of Northeastern Illinois University. I soon began to offer a seminar to a few select students in which I used Powers, Clark, & McFarland (1960a, b) as the main text, supplemented by notes on my experiences in learning the theory and a few other reprints from Clark and McFarland. After giving the class a few times, I had four students who really grasped the theory and began to use it in constructive ways in their own lives. One of them made what I consider a profound use of the concept of reorganization in clinical work—one that I continue to find invaluable in my own practice of psychotherapy.

She was doing a fieldwork project in a psychosomatic ward of a large general hospital. She was allowed to talk with patients as a kind of supportive "mentor," being a graduate of the ward herself. One day, one of the patients was threatening to withdraw from the program, complaining of severe anxiety, saying, "Nothing seems right, all my ways of thinking about things are up for grabs." My student had a powerful "Aha!" at that point. She said, 'Wait a minute, that is just what you should be experiencing. If you were still reacting in the way you used to, you would be doing what got you here in the first place. The fact that nothing seems right is because you have changed. You are no longer perceiving things as you did, and so your experience does not feel familiar." With this insight, the patient chose to tolerate her unsettled state a while longer and eventually proved a success in the program.

It was at about this time that *Behavior: The Control of Perception* (Powers, 1973) came out. We devoured it eagerly, and it helped us draw several further applications from this experience with the patient on the psychosomatic ward. It occurred to us, after reading Powers' chapter on learning, that this patient's complaint of everything being up for grabs was in a certain sense similar to the variability in performance in the learning experiments and the ping pong tournaments. Having become disillusioned with her old ways of thinking about her experiences, she wavered through various new percepts. In Powers' theory of reorganization, this would be the result of random signals in the existing hierarchy. They would, of course, interfere with the functioning of some of the existing systems. As she began to settle upon new, more effective ways to view her situation, she reported that her anxiety dissolved. We began to form a wholly new idea about anxiety as a result of these observations.

Clinical reports in the field of psychotherapy have frequently noted that patients complain of anxiety as they move deeper into self-examination. Clinicians have typically treated anxiety as a condition to be gotten rid of, rather than as an indication of an underlying process of reorganization. However, it seemed to us that—if Powers's concept of the control systems of a person as a huge interconnected hierarchy were right—then of course when the reorganizing system begins to alter the parameters of some systems in the body, many other systems that interact with them would be plunged into varying degrees of error. It struck us that anxiety could be the name given to that condition. Later on, I ran across other clinicians whose experience had also suggested that anxiety in treatment often appears to be a precursor of major change. I have subsequently had many instances where simply offering this view of anxiety has helped a person to refocus on their desired changes, instead of on the symptoms of anxiety, and to achieve a good outcome.

Another application of PCT in my clinical work has been to encourage patients very firmly to keep stating their goals or objectives as specifically as possible, a method I share with David Goldstein and Ed Ford, although each of us seems to do it in a slightly different way. My favorite procedure is to ask the person repeatedly to state what he or she wants to perceive or experience in literal terms. For example, "So, you want to hear your boss say, 'You're the hardest worker here,' is that right? How close has he come to saying just that? What do you do that gets him to say anything like that?" It continues to surprise me how often a person is looking for a particular feeling but has hardly any idea of the kind of events that must occur for him or her to get that feeling.

Once patients get the idea that the good feelings being sought are closely tied to specific events, they usually take off with the concept, making further applications on their own. I wrote an applied control-theory psychology text for my students in an introductory mental health course, to pass along these observations. It gave a simplified sketch of the theory, showing how it had led to these and other applications. I also offered it to a number of "pop" psychology publishers. I heard from several of them that the first reader or two liked it, but it was always rejected at higher levels. A fair number of my students grasped enough of the main idea to make their own applications, as the patients had done, so it seemed to have served some purpose anyway.

Later Research

Once I had *Behavior: The Control of Perception* as a text, my classes grew slightly larger, and I found some students who were interested in learning more about doing research from the control-theory point of view. We settled upon research on the self as a fertile ground. Carver and Scheier's book (1981) had come out by then. They reviewed many studies in social psychology dealing with various aspects of the self, which they interpreted, more or less, as having aspects of a higher-order control system. Since I had been interested in the self since way back in my days with Carl Rogers, I found their work of considerable interest. However, I wasn't satisfied with the relatively murky views of this concept that one finds in the literature on it. I proposed it as a subject for deeper investigation in one of my first advanced courses in perceptual control theory.

I have an unusually clear (for me, at least) recollection of the progress of that series of discussions. I think it illustrates well the development of a theoretical question through intermediate steps to a research program, so I would like to spell it out in detail. We began by speculating that at least some of the previous concepts of the self in the psychological literature seemed suggestive of control systems, even on the part of writers who had never heard of control theory. A good example is Epstein (1973), 'The Self-Concept Revisited, or a Theory of a Theory." He proposed that the term "self" is used to describe a conception, or theory, that a person develops to explain him- or herself to him- or herself for the purpose (among others) of knowing how to make difficult decisions.

In our discussions, we began to play with the idea that a self could then be thought of as a control system of the highest order. What would it control? We examined Powers's scheme of the learned hierarchy for types of variables controlled at the various levels. Variables like intensity, configuration, relationship, and sequence are, in a certain sense, very concrete. That is, you can construct objective measures for them, as Bill did in constructing tracking experiments where the relation-

ship of "equal" or "in line" can be viewed directly between cursor and target on the screen. Now, what would be the counterpart of that at the level of a system controlling that a person would continue to be the same consistent person? It struck us at some point that we were seeing that in action all of the time. We noted that when we talk about ourselves, a portion of our talk consists of telling each other what kind of person we are. "I am a quiet person," "I am a talkative person," "I am a shy person," etc. These are the kinds of attributes, collections of which some writers called "self image."

At this point, we derived, from the theory presented in *Behavior: The Control of Perception*, an implication that had not been clear in previous studies of the self. In Powers' discussion of how you can determine whether you are observing a control system in action, he described the "test for the controlled variable." If a phenomenon is under feedback control, you will see it corrected back to its prior state if you disturb it. During one of our class discussions, one of the students made the self-image remark, "I am a shy person." I simply said, "No, you're not," as an attempt to apply Powers' test. Her jaw dropped, her eyes widened, and she said, with indignation in her voice, "I certainly am!"

Looking back, I marvel at how much more work it took to incorporate this simple experience into a workable experiment. I proposed at this point to David Goldstein that we work on it. But as traditionally trained psychologists, we seemed to have to go through a series of successive approximations to move from a traditional research format to a rigorous presentation of this original, simple, informal test of the controlled variable. I will describe the history and present the research here, since it is unlikely to be published anywhere else.

We began with a design typical of hundreds of studies on various aspects of the self. We got subjects to describe themselves and their ideal selves on an adjective checklist and then had them estimate where their own scores would fall on a wheel-like circumplex of eight personality factors (sociable, accepting, submissive, assertive, etc.) from Conte & Plutchick (1981). A week later, we gave them a doctored "personality profile report" in which some of the factors they had rated as self-descriptive were affirmed but others were reversed. Our rationale was that the false descriptions would result in a sense of error which they would then take some action to correct. We provided the opportunity to do that by describing the doctored feedback as a new experimental instrument and invited them to correct any attributions they thought the testing had got wrong. We then scored any statements they made to correct "wrong" descriptions as favoring the hypothesis and failure to contradict as against the hypothesis. Anyone familiar with the typical research in this area will recognize that this design follows a very usual format, as for example the study by Frey and Stahlberg (1986) that Runkel (1990) took apart in great detail in his text on psychological research methods.

Table 1 shows the results of three samples of subjects with whom this first design was employed. These results do not support the hypothesis. In the first sample, there were more instances of acceptance of false attributions than corrections. The second sample results were favorable to the hypothesis. In the third sample, there were more instances of correcting undoctored attributions. At that point, we faced a question common to many research projects of this sort. Was the hypothesis (that people would correct falsified self-descriptive attributions) disproved by the results, or was the experiment inadequate to the task? Like many researchers who have invested time and money, we preferred to believe that the method was inadequate.

Table 1. Reactions to receiving false attributions.

Subject sample number	N	Corrections of disturbed categories	Acceptances of false attributions	Corrections of undisturbed attributions	Ambiguous
1	10	9	10	6	5
2	12	13	9	6	8
3	12	8	2	18	8

Note: Each subject responded to three questions, hence frequencies show number of chances to correct; that is, three times the number of subjects.

One member of the research team, a schoolteacher familiar with students like our subjects, speculated that many of them did not have a very robust self concept. We went back to the drawing board, determined to control for a confounding factor such as ego strength. The data were already at hand in unanalysed information that we had gathered during the project. We tallied up the discrepancies for sample 1 between "actual" and "ideal" ratings that the subjects had given themselves on the circumplex measure, then defined a measure that we called "self-knowledge" as the inverse of the total. We split the sample at the median on this measure and cross-tabulated it with the correction data. This time, indeed, we found that the subjects with the higher "self-knowledge" performed according to the hypothesis, as compared with those low on "self-knowledge," as indicated in Table 2.

Table 2. Sample 1 subjects' reactions to receiving false attributions, by high and low self-knowledge groups.

Discrepancy	N	Corrections of disturbed categories	Acceptances of false attributions	Corrections of undisturbed attributions	Ambiguous
high	5	1	5	7	2
low	5	8	4	0	3

Note: High discrepancy is equivalent to low "self-knowledge," and vice versa.

A chi square on these results was significant at the .05 level, and we could presumably have had it published somewhere, had we stopped at this point. But we made the mistake of trying to replicate these results. The data for samples 2 and 3 came out in the opposite direction. The operational hypothesis was thus invalidated. I might note that, in all of these samples, there were contrary instances, and whatever differences were noted were only between group means, a fault that Kunkel (1990) has pointed out in almost all psychological research aimed at investigating properties of human beings *qua* human.

After some intermediate steps which are not worth describing here, we began to see that the concept of testing for a controlled variable calls for an entirely different research design—and for results that should be universal. The first inference we drew from careful thought about PCT was that the instance of disturbance of the self image, and its correction, if any, should be immediate in real time. We had realized that there is no particular justification for assuming that any individual is controlling exactly the same aspect of his or her self image a week later, as compared to the initial selection of adjectives.

We developed a format closely similar to the initial, informal situation from which the inquiry had started. We had subjects work in pairs in which one partner would do a Q-sort self-description with items from the original adjective check list. The other partner had been secretly instructed to read off the first item and say, "Why, no, you're not like that," immediately upon completion of the Q-sort. The complicit partner then wrote down the other's reply. We then had judges score replies like "I am so" as for, and all others as against the hypothesis. Table 3 presents these results.

Table 3. Subjects' replies to contradiction of self image in four samples, using the second design.

Subject sample number	N	Corrections	Non-corrections	Ambiguous
1	8	8	0	0
2	8	7	0	1
3	10	8	1	1
4	9	9	0	0

Summing the results of 35 subjects in 4 samples, we have 32 instances of correcting, one failure to correct, and two unscorable replies. This finding appears considerably more robust than one usually finds in typical research on this topic. However, we were uneasy about the deceptive aspect of the way in which the self image had been disturbed. Therefore, we designed a format in which subjects again made self-descriptions, but this time we asked them to imagine what they would say to someone who looked at their description and said that it was not accurate. As a control, we asked them to do the same with an arbitrary list of neutral adjectives. We had their answers rated by student judges according to whether they objected or not to the aspersion, as well as to the neutral terms. Table 4 shows those results.

Table 4. Written responses to hypothetical disturbance of self-chosen self-descriptive and neutral attributions.

Reactions to relevant statements				Reactions to neutral adjectives			
Subject	Correct	Modify	Accept	Subject	Correct	Modify	Accept
1	1	4	0	1	0	4	1
2	3	2	0	2	0	2	3
3	5	0	0	3	0	3	2
4	0	5	0	4	2	2	1
5	5	0	0	5	2	3	0
6	4	1	0	6	0	0	5
7	4	0	1	7	0	2	3
8	5	0	0	8	0	4	1
Total	27	12	1	Total	4	20	16

This methodology is rather simple and perhaps primitive, but its strength lies, I believe, in that it applied some rigor to something anyone can witness in everyday life. I have repeated the informal experiment now on hundreds of occasions with almost universally consistent results. Anyone else who wishes can do the same. It doesn't require any particular lab set-up or complex instrumentation. All that is required is to wait until a potential subject makes a self-descriptive remark and then contradict or interfere with it in some way and observe the result. I am satisfied that the objective has been achieved. There is almost invariably a strong correction to a disturbance of self-description when a person declares himself or herself to have such and such a characteristic. From that, I conclude that it is feasible to regard the "self" as a type of control system, and the "self image" as a type of controlled variable monitored by such a system.

Toward the Future

I have found it extremely exciting—and I still do—to be "along for the ride" in this paradigm revolution concerning the nature of behavior. I get a thrill when I experience the sense of simplification by seeing an odd collection of psychological "phenomena" as special cases of the same underlying process. For example, some of my most satisfying experiences in working on the textbook (Robertson & Powers, 1990) were insights such as when it occurred to me that "self-fulfilling prophecies," "experimenter bias," and learned helplessness" could all be seen as special cases of control of expectations. (The reference settings in each instance were established in the particular events used to define these "phenomena.") In the literature where they are introduced, they are offered as unique human processes, unrelated to each other.

The psychological literature is full of such cases. It is equivalent to the condition that would have obtained in physics before Isaac Newton. The motions of planetary bodies, apples falling from trees, and cannon balls would have all had to be explained with separate and unrelated "laws of nature." The lack of a unifying theory allows, nay, introduces many false complications into psychology and, at the same time, diverts energy from investigating the true complexity of living processes to the pursuit of many trivial distinctions and measurements.

I must acknowledge, of course, that drawing upon theory to attempt to simplify the underlying structure of phenomena is only the first step in gaining knowledge. Proposed simplifications are speculations that need confirming. That is, they need confirming in those instances where the surface phenomena continue to be interesting after one takes a look at them from a PCT point of view. I suspect that in many instances they will not. At least, I have stopped being amazed to have it pointed out that people regularly act, quite automatically, to bring about experiences that match their expectations—of whatever sort.

I believe, further, that there is a tremendously varied and exciting realm of possibilities for different directions in which to test out, and draw applications from, PCT. The study of some of the lower-order systems is well along in the various tracking experiments done by Powers, Bourbon, Marken, and others. The existence of higher-order systems, postulated by Powers, is to my mind established in the self research results. But there is much to be done in investigations of the intervening levels.

Some of that is already implicit in Bourbon's results with two-person interactions. I don't know if anyone is yet sketching out (or cataloging—is that a more apt term?) the principles and programs that different people implement in doing the tasks. Likewise, it will be very interesting as young investigators construct and test models of how different people choose a strategy for dealing with a task, the mastery of which is unknown to start with. The Powers Game is one type of activity where subjects' choices of strategies for mastery will be amenable to the test for the controlled variable. There obviously are many more.

There also needs to be research on the reorganization phenomenon in all kinds of learning situations. Many fertile questions about how reorganization proceeds have appeared on CSGnet in the past year. I would hope to see some of them instrumented and pursued in the near future.

There are many observational facts in psychology that might well be recast into PCT terms in an approach to find the underlying mechanism. For example, compare Plooij's (1987; 1989a, b) work on developmental sequences with that of Piaget. Piaget presents some excellent step-by-step descriptions of how behavior gradually becomes more complex in many skill dimensions. But his "theory" proposes "explanatory" concepts like "equilibration," which resolve into "it happens because that's how it happens," when analyzed. In contrast, Frans Plooij, also describing some invariant sequences in behavioral development, has related them to Powers' hierarchy of controlled perceptions and has shown how the more complex are combinations of the prior steps in development.

There are many observational facts in developmental, clinical, and social psychology, but very few genuine attempts to propose underlying mechanisms, outside of PCT and the field of psychobiology. In that field, control theory is beginning to be applied, but is hobbled by the kind of misunderstanding of it that Powers has spent so much time pointing out. And certainly, there are many potentially rewarding applications of PCT to follow those being begun by a growing number of us.

For the person who gets personal satisfaction from seeing an unknown shore come into view for the first time, as well as from the company of fellow voyagers, PCT has emotional and intellectual satisfaction to offer. I find it rewarding to interact personally at our small face-to-face PCT meetings with people whose specialties are distant from my own, as well as with those in the same field. I'm glad at such times that PCT hasn't attracted a great horde of people whom I wouldn't be able to know as individuals. I don't enjoy seeing anonymous faces, talking about things about which I have no chance to stop and ask what they really mean. For that reason I haven't attended an American Psychological Association convention in many years. I feel some sadness in knowing that this is already beginning to change in the CSG. When I see Bill Powers laboriously leading a querulous interloper in CSGnet through the careful first steps of understanding how behavior is the control of perception, I often get an impulse to shout, "Save yourself for those who come of their own accord!" And as to the fact that so many well-established people can't be converted, we have already sufficiently understood how PCT already predicted it. Let's not waste any more time about that. On with finding new facts and making new discoveries!

References

Bryan, W. L., & Harter, N. (1899). Studies on the telegraphic language: The acquisition of a hierarchy of habits. *Psychological Review,* 6, 345-375.

Carver, C. S., and Scheier, M. F. (1981). *Attention and Self Regulation:* A *Control Theory Approach to Human Behavior.* New York, Heidelberg, Berlin: Springer Verlag.

Conte, H. R., & Plutchick, R. (1981). A circumplex model for interpersonal personality traits. *Journal of Personality and Social Psychology, 40,* 701-711.

Epstein, S. (1973). The self-concept revisited, or a theory of a theory. *American Psychologist, 28,* 404-416.

Frey, D., & Stahlberg, D. (1986). Selection of information after receiving more or less reliable self-threatening information. *Personality and Social Psychology Bulletin, 12,* 434-441.

Keller, F. S. (1958). The phantom plateau. *Journal of the Experimental Analysis of Behavior, 1,* 1-13.

Plooij, Frans X. (1987). Infant-ape behavioral development, the control of perception, types of learning and symbolism. In J. Montangero, A. Tryphon, & S. Dionnet (Eds.), *Symbolism and Knowledge* (pp. 35-64). Geneva: Archives Jean Piaget Foundation.

Plooij, F. X., & van de Rijt-Plooij, H. H. C. (1989a). Vulnerable periods during infancy: Hierarchically reorganized systems control, stress and disease. *Ethology and Sociobiology, 10,* 279-296.

Plooij, F. X., & van de Rijt-Plooij, H. H. C. (1989b). Evolution of human parenting: Canalization, new types of learning, and mother-infant conflict. In J. 13. Hopkins, M.-G. Pêcheux, & H. Papousek (Eds.), Infancy and Education: Psychological Considerations (Special issue]. *European Journal of Psychology of Education, 4,* 177-192.

Powers, W. T. (1973). *Behavior: The Control of Perception.* Chicago: Aldine.

Powers, W. T., Clark, R. A., & McFarland, R. L. (1960a). A general feedback theory of human behavior: Part I. *Perceptual and Motor Skills, 11,* 71-88.

Powers, W. T., Clark, R. A., & McFarland, R. L. (1960b). A general feedback theory of human behavior: Part II. *Perceptual and Motor Skills, 11,* 309-323.

Robertson, R. J. (1984). Control theory. In R. Corsini (Ed.), *Encyclopedia of Psychology,* Volume 1 (pp. 288-291). New York: Wiley.

Robertson, R. J., & Clines, L. A. (1985). The phantom plateau returns. *Perceptual and Motor Skills, 61,* 55-64.

Robertson, R.J., Goldstein, D. M., Mermel, M., & Musgrave, M. (1987, October). *Testing the self as a control system.* Paper presented at the meeting of the Control Systems Group, Kenosha, WI.

Robertson, R. J., & Powers, W. T. (Eds.). (1990). *Introduction to Modern Psychology: The Control Theory View.* Gravel Switch, KY: CSG.

Rogers, C. R. (1959). A theory of therapy, personality and interpersonal relations, as developed in the client-centered framework. In S. Koch (Ed.), *Psychology: A Study of a Science,* Volume 3, N.Y.: McGraw-Hill.

Runkel, P. J. (1990). *Casting Nets and Testing Specimens: Two Grand Methods of Psychology.* New York: Praeger.

Confessions of a Non-Pioneer

Tom Bourbon
(Research Division, Department of Neurosurgery, University of Texas Medical School - Houston, 6431 Fannin, Suite 7.148, Houston, TX 77030)

The people who can write as legitimate PCT “pioneers” are Bill and Mary Powers and Bob Clark. The fact that Greg Williams includes me in a list of pioneers says much about the past, present, and probable future of PCT; the message is not encouraging. Much of what I have written here reveals the lack of any contribution by me during the early years of PCT. If what I say is of any value, it is probably by way of documenting how long some of the supposedly “new alternatives” and “new objections” to PCT have really been around.

1960-1973

In 1960, when Powers, Clark, and McFarland published their first papers on control system theory (CST) in *Perceptual and Motor Skills,* I was an undergraduate student, changing my major and my institution for the third time. I had completed four courses in psychology at a small private college, all under one Skinnerian radical behaviorist. Since my high school days, I had been interested in studying how people and their environments affect one another. I saw little chance of working on that subject under a devoted rat runner. Seeking broader exposure to psychology and physiology, I transferred to the University of Texas at Austin in 1960. I eventually earned a B.A., then completed a Ph.D. in 1966.

During those years, while the original papers on CST languished in obscurity, I learned a psychology that was a mixture of Hull-Spence behaviorism, the ecological psychology of J. J. and Eleanor Gibson, information theory and its applications in psychology by people like Garner and Hake, classical psychophysics, and the early work on signal detection theory. The “information processing perspective” was in its birth throes.

In physiology, I came across work on control and regulation, but I read much of that material on my own. The physiologists and physiological psychologists with whom I studied presented some material on physiological "feedback," but then the term referred to pathways that descended from higher centers, "feeding back" down the nervous system toward the sensory receptors—feedback was activity feeding back down the lineal causal pathway for behavior, serving some unknown function. I thought there was more to the idea than that but could not work it out. (Incidentally, when they diagrammed a control process, nearly all physiologists used a version of Norbert Wiener's diagrams, in which the system controls its output.)

I learned my psychophysics under Chuck Watson, a young faculty member just graduated from Indiana University, a point of origin of signal detection theory, and under Lloyd Jeffress, one of the authorities on binaural hearing. In Watson's graduate seminar on instrumentation and electronics, I located material on servomechanisms and negative-feedback control systems. There was no hint that such a device might serve as a model for a living system, but I thought it might. I did nothing more with that vague idea at the time. Jeffress was one of the grand figures in my educational experience. I completed my dissertation under him on an obscure topic in auditory signal discrimination. During all that time, while I pursued trivia, Bill Powers, and Mary, worked alone on *Behavior: The Control of Perception.*

I finished my degree during a "seller's market," when new Ph.D.s had many opportunities awaiting them. Like all of Jeffress' students before me, I was expected to go to a decent university where I would become an audition-psychophysics person, teaching, researching, and publishing on those topics for the remainder of my career. I disappointed everyone by going to a university that offered only the B.A. and M.A. in psychology. It was a place where there would be no pressure to publish or perish, which meant I would be free (as free as 12- to 15-hour teaching loads each semester would allow) to pursue my own line of study. After my education in psychology and physiology, I was convinced I had been conned—there was nothing in it that could explain the behavior of real organisms in the real world (apply the PCT interpretation of "real").

By 1973, my students were enduring the results of my personal attempts to fit feedback into the ecological psychology of the Gibsons, and to tie it all together in a mix with information theory and signal detection theory. (All of those poor students. I should have been tarred and feathered!) I was convinced that the environmental and ecological levels and layers described by the Gibsons must represent classes or levels of perception in a person, not objective features of the environment—basic psychophysics and physiology left no other possibility. But I was getting nowhere with working it all out.

I read Bill's 1973 article in *Science*. I can still remember the electric feeling as I devoured it, saying over and over to myself, "Of course! Of course!" He made it all so transparent and simple. When I received the flyer from Aldine announcing "the book," I ordered it immediately. The notes scribbled all over my old copy document my excitement with the book, and my conviction that it was a work of genius, and my determination to do something, no matter how minor it might be, or when it might be, to help "put out the word" about CST. As things turned out, it was quite a while before I did anything other than make CST part of my teaching in an obscure university.

From the first day when I talked about CST in my classes, some students were interested in the theory and excited by its implications for behavioral science. (Of course, many others—and my faculty colleagues—were not at all excited.) David Goldstein joined our faculty to teach courses in developmental psychology and cognitive psychology. He encountered students who kept talking about CST and about how it was related to topics in his classes. In self-defense, he came and asked me what it was all about. During the time from 1973 to 1992, David was the only one of my former colleagues who ever tried to learn what the commotion was really about; most of the others were content to ignore me and take cheap shots at my students. Eventually, David was in trouble running up large phone bills by calling Bill Powers, and he helped bring Bill to our campus for a visit. By then, I had convinced myself that, despite my determination to help, I was equipped to do very little.

In an attempt to learn how to do computer modeling, in the late 1970s I attended an NSF-sponsored short-course on modeling with NDTRAN, a systems dynamics modeling program patterned after the more elaborate and expensive DYNAMO. A few thesis students dared to tackle some modeling problems in CST using NDTRAN. In 1980 and 1981, we went on the road to talk about those projects, first at a meeting of the Society for General Systems Research (SGSR), then at a meeting of the IEEE Society for Man, Cybernetics, and Society. In both places, everyone seemed to think CST was an old-hat version of cybernetics— not their cup of tea. At SGSR, they were interested in entropy, chaos, nonlinear systems, and lots of elaborate verbalisms; at IEEE, they cared about optimization and models of optimal control. CSGnet is not the first place where we have encountered resistance from devotees of those ideas.

Frustrated, in 1982 I went alone to Columbus, Ohio, to a meeting of the American Society for Cybernetics (ASC). I hoped to locate at least a few people who might be interested in CST. When I wandered into the general meeting hall the night before the meeting was to start, I was surprised to see Bill Powers, messing around in a tangle of cables, setting up his homemade computer. At ASC, nobody else was interested in CST. I was dismayed to see that they could so easily reject Bill's ideas without even a hint of a fair hearing or an attempt to understand. Bill was ready to give up on the cybernetics people, whom we both thought ought to be the group most likely to understand CST. I was afraid he might give up altogether—the years of rejection and, for the most part, "going it alone" had taken a toll. In retrospect, I know he would never have quit.

1983-1994

I hit on an idea that I thought might help. It had nothing to do with modeling, and it is probably my one significant contribution to the present state of CST (now known as perceptual control theory, PCT). I knew that Bill and Mary kept a map on their kitchen wall next to the refrigerator, with a pin marking the location of each person who had called or written to ask about or discuss CST. I asked Bill to send me the short list of names that accompanied the pins. I contacted all of them and invited them to attend the 1983 meeting of ASC, in California. Then I contacted Bill Reckmeyer, president of ASC, and told him we needed three sessions on the program. To my amazement, Reckmeyer gave us the sessions. In 1983, a band of seven CST people used smoke and mirrors to create the impression that we were everywhere at the ASC meeting. One evening, at dinner in a Chinese restaurant, we gave Bill a "certificate" from the "off-the-wall group" of control system theorists, commemorating the 10th anniversary of the *Science* article and *Behavior: The Control of Perception.* The next year, Bill Benzon organized the CST contingent at the meeting of ASC in Philadelphia. There, ASC people were "into" poetry, "second-order cybernetics," and autopoeisis—they gagged on CST.

In Philadelphia, we could not all fit in a single elevator. We knew we were really making progress! I flew back to Texas with a good feeling about what was happening. The next year, the CSG held its first meeting, at Kenosha, Wisconsin. I'm sure other "pioneers" have written about the CSG meetings.

We were not finished with ASC. That group had convinced the organizers of the annual Gordon Research Conferences that the ASC brand of cybernetics was scientific and deserved to be the subject of two Gordon Conferences. The first was in Wolfboro, New Hampshire, in 1986. Bill, Mary, and I were the only CST people there; second-order cybernetics, autopoesis, aesthetics, and deconstructionism ruled. At the second conference, in California, there was a formal CST session.

After the California gathering, the Gordon people dropped the ASC from the Gordon Conferences (not, I hasten to add, because of the CST session —the problem was too much emphasis on aesthetics and too little on science).

One evening during the first Gordon Conference, Bill, Mary, and I escaped onto the boat dock and were talking about a paper he had started. We agreed to collaborate on it. It eventually became "Models and Their Worlds," which was rejected several times by legitimate" journals, then published in *Closed Loop* in 1993, seven years after we began our collaboration.

Along the way, Bill gave me copies of some of his programs, on which I hacked around and half-way learned how to do programming and "real" modeling. It seemed obvious that my poor programming might benefit from time spent with Bill, and I planned to spend a week or so visiting him in Northbrook. Greg Williams and Bill Williams learned of my plan and arranged to be in Northbrook at the same time. That was the start of a three-year series of Northbrook "mini-conferences" on CST. During each of them, Dick Robertson and Wayne Hershberger dropped in for awhile. I was always desperate for travel money to attend the mini-conferences. One summer, I submitted a proposal for faculty development money to support a trip to "the laboratory of William T. Powers." In the proposal, I said I would hold down expenses by "sleeping at the laboratory." I didn't tell them the laboratory was in the room behind the kitchen in Northbrook. To keep things "honest," I gave Bill a plaque that proclaims 'The Laboratory of William T. Powers," a place I believe is one of history's great centers of intellectual accomplishment, as I said in the Foreword to Bill's *Living Control Systems II.*

My Students

While I was teaching from 1967-1992, I directed 55 master's theses. After 1973, 14 students dared to complete theses that involved PCT. Some used it as a "perspective" for reinterpreting other work in psychology. Others used it as a source of tasks or behavioral measures for research projects. A few of the hardiest used formal PCT modeling in their theses. No matter the degree to which they used PCT, they all encountered far more flack and nonsense than the typical graduate student in our department—my former colleagues never did appreciate PCT. I had, and always will have, great respect for all of those graduate students, and for the many undergraduates who also faced what was often outright scorn from their peers, who all knew (perhaps with a little help from *my* peers?) that PCT was folly and those who followed it were fools.

Back when I began teaching psychology, I was "safe." Many of my earlier graduate students completed doctoral programs and have become clinicians, research scientists, faculty members and administrators. But in the PCT era, I watched one student after another go off to doctoral programs, wishing they could continue to study PCT—saying they would manage to keep their interest in PCT—then caving in under the pressure to do what they *must* do to survive. During the final few years of my teaching career, I could not accept the error I experienced over seeing students become excited about PCT, then asking me the inevitable question, "Where can I go for a Ph.D.?" At the end, I was watching undergraduates and graduate students simply drop out of psychology altogether, rather than study in a traditional department. Many of those who wanted careers in clinical practice opted for certification in areas other than psychology. I no longer thought it was fair for me to expose students to material that could only end their hopes for professional careers in behavioral science. I left teaching for what looks as though it will be a full-time effort to obtain funding for research on topics other that PCT.

Beyond 1994

Now, 34 years after the first papers on CST-21 years after *Behavior: The Control of Perception* and Bill's article in *Science* a handful of people do PCT modeling. In one way or another, most of us either abandoned, or never pursued, a traditional career path, out of our conviction that PCT is a revolutionary theory. There seems to be no other way to go about this business. Unless others soon pick up the task of PCT modeling, there is very little future for PCT as a science, no matter how many dedicated people use PCT in their applied work. The modelers never had doctoral students who went on to perform modeling. We never had doctoral students, period!

Without more contributors to the modeling, we run the risk of PCT becoming one more among a multitude of "perspectives" or "frameworks," another gloss added to the theories a person held before encountering PCT. There is abundant evidence of that phenomenon on the CSG computer network. There, I see one person after another stake a claim that this, that, or the other theory "says the same things as PCT"; or that PCT is fine, as far as it goes, but that such and such theory is necessary for going further; or that blah theory is more fundamental than PCT and can generate PCT. Time and again, 1 am struck by the fact that those invocations of other theories sound familiar; they are often the same things I heard in the 1960s—the same theories I recognized as part of the immense con job that passed for my education and training in scientific psychology.

In addition to more modelers, we need more people to gather solid empirical evidence to demonstrate and document the phenomenon of control at every level, including social, neurological, biochemical, and applied. We must demonstrate the phenomenon of control, *then* invoke the model, not the other way around. As things stand, even some supporters of PCT show very little interest in empirical work on control—a few even dismiss that work as trivial and say it "adds nothing to our understanding of PCT." With friends like that....

Finally, anyone who suggests that another theory, or an improvement or addition to PCT, is better than any present working version of the PCT model must produce the evidence—a working model that does the job better. Given evidence like that, there is no question about which model works best. That's the only way this game can be played.

My Life as a Control Theorist

Richard S. Marken
(Life Learning Associates, 10459 Holman Ave., Los Angeles, CA 90024)

My life as a control theorist began in the spring of 1974 when I was roaming through the library at the University of California at Santa Barbara. I had just completed the requirements for my Ph.D. and was preparing to start what would turn out to be a very pleasant career as a professor of psychology at a small college in Minneapolis called Augsburg College. My major emphasis in graduate school had been the study of human perception; my thesis research addressed a rather arcane question in auditory psychophysics: what is the time and frequency resolution of the auditory system for detecting tonal stimuli? I got the teaching position at Augsburg based on my familiarity with computers rather than auditory psychophysics (the latter not being in great demand at the time). Before leaving beautiful Santa Barbara, I would occasionally roam through the library to see what was new in the psychology section—particularly in the "perceptual psychology" section. It was during one of these tours that I ran across a book by William T. Powers called *Behavior: The Control of Perception (BCP).*

I was attracted to the book by its title. Being a student of perception, I found it downright puzzling. I had thought about perception quite a bit in graduate school, but I would never have thought of it as something to be controlled (whatever that meant). I checked out the book and found that it had something to do with control theory, feedback, and cybernetics. I was immediately impressed by the clarity of the author's presentation; he seemed to know what he was talking about. I was intrigued, and a little frightened; intrigued because what Powers was saying seemed to be relevant to "real" human behavior in a way that all my graduate studies in psychology had not been; frightened because, in a way I could not yet articulate, Powers seemed to be calling into question some things that I took for granted. I probably spent about two hours with *BCP,* but it had made a bigger impression than I then knew.

My encounter with *BCP* was just a vague memory when I ran across it again in 1977—this time in the library at Augsburg College. But now fate would intervene to allow *BCP* to change my life. Various chance factors made it possible for me to learn what we now refer to as perceptual control theory (PCT), the theory of behavior described in *BCP.* I was, at this time, preparing a rebuttal to a talk given by a colleague (and good friend) about the meaning and implications of Skinnerian behaviorism. B. F. Skinner was receiving a great deal of attention after publication of *Beyond Freedom and Dignity.* Discussions of "determinism" vs. "free will" seemed to be all the rage, and the possibilities of "behavioral engineering" were being actively debated. Like most "cognitive psychologists," I thought that there was something wrong with behaviorism—I just couldn't articulate the problem to my satisfaction. I was planning to base the rebuttal to my friend's pro-Skinner talk on Ulric Neisser's discussion of Skinnerianism in his recently published *Cognition and Reality.* Neisser was one of the "biggies" in the then relatively new field of cognitive psychology, and his arguments against behaviorism made sense. But I could tell that these arguments were more like opinions than scientifically based conclusions.

BCP reappeared at this time. I saw its relevance to my planned rebuttal right away, since the discussion of PCT was framed as a critique of behaviorism. *BCP* seemed to provide a scientific (rather than an emotional) alternative to behaviorism that was missing from the cognitive view. After weeks of going back and forth—whether I should base my rebuttal on the conventional "cognitive" view or on the PCT view of behavior (which I still only vaguely understood)—I opted for PCT. This turned out to be a good choice because it started me on the path to learning PCT; it was a bad choice because I had only the vaguest idea what I was talking about—which made for a pretty ineffective rebuttal.

But I quickly lost interest in debating Skinnerians because another fateful development helped me move forward in my understanding of PCT. In early 1978, some friends of mine at Stanford introduced me to the "personal computer." It was love at first sight. I immediately bought an RCA Cosmac computer kit (with 8K of memory, mass storage on a cassette tape, and "hex pad" data input). By 1978, I had convinced Augsburg to buy a couple of what were then considered to be very fancy Apple II computers. At the same time, I (again by chance) ran across an article by Powers in *Psychological Review* ("Quantitative Analysis of Purposive Systems"). A number of tracking experiments were described in that article, and I was able to replicate them fairly easily on the Apple II (using the game paddle as cursor controllers and Apple Basic as the programming language). Shortly after that, I found Powers' series in *BYTE* magazine, which reassured me that I had been doing the experiments correctly and taught me how to do some of the modeling needed to evaluate the results of the experiments.

My experiments with the personal computer helped me understand PCT in a way that would have been impossible (for me) with words

alone. These experiments made it possible for me to understand what "the control of perception" meant; it meant the end of psychology as we knew it. At this time, I was also writing a textbook on experimental psychology and statistics. I was leading a double life: my PCT experiments on the personal computer showed that the basic assumption on which all psychological research is based—the assumption that perceptual input is the cause of behavioral output—is wrong; and I was writing a textbook explaining how to do psychological research based on this assumption. By 1979, I knew that what I was saying in the textbook was wrong. But I finished the book (it was published by Brooks/Cole in 1981 as *Methods in Experimental Psychology)* in order to get tenure (I did) and to show the "right" way to explain the wrong way to study behavior (I believe I succeeded, though the book was not a bestseller). Once the book was finished, my attachment to conventional psychology was finished as well.

What finally "put me over the edge" and convinced me that PCT is a revolutionary new approach to understanding behavior was the apparently trivial (but completely astounding) realization that, in a control loop, the input to the loop is *not* the cause of the output. This can be demonstrated most easily in a compensatory tracking task where you are to keep a cursor aligned with a target. When control is good you are able to keep the cursor almost exactly on target, despite the fact that there are disturbances that would tend to move the cursor away from the target. You keep the cursor on target by moving a control handle appropriately; to the left to keep the cursor from moving off to the right, and to the right to keep it from moving off to the left. Most people looking at a subject performing this task would say that the deviation of the cursor from the target "tells" the subject which way to move the handle in order to keep the cursor on target; deviation of cursor from target is the "stimulus information" that is used by the subject to make the appropriate responses. But Powers showed that there is almost no relationship (correlation) between deviations of cursor from target and movements of the handle that controls the cursor. Yet there is a nearly perfect relationship between the unseen disturbances to the cursor and handle movements. These results seem "magical"—completely contrary to the "input-output" or "cause-effect" model of behavior—yet they are exactly what is predicted by PCT.

I started trying to do experiments to see if I could find a flaw in the PCT view of the tracking situation. After all, manual tracking studies were well known to me and had been done for years; how could anyone have missed this incredibly surprising fact—that inputs don't cause outputs, that what subjects see doesn't determine what they do. It was during this period that I hit on the idea of having the subject do two runs (at different times) with exactly the same disturbance present both times. This was easy to do with the computer. The idea was this: even though the correlation between cursor and handle movements is low, it might be that something about the cursor is still the cause of handle movements; it's just something that is not picked up by the correlation. For example, handle movements might be caused by some function of the cursor movements, or by the cursor movements from some time in the past, or by some odd weighting of several cursor positions, etc. Any of these aspects of cursor movements might be the cause of the handle movements, and if so, they would not show up in a simple correlation between cursor and handle movements. However, they would show up in a correlation between cursor movements on two different trials where essentially the same handle movements had occurred; the cursor movements on these trials would correlate because something about them must be the same if the handle movements were the same. I knew that I could get the subject to make nearly the same handle movements by presenting the same disturbance on two different occasions; the cause-effect model would predict that the cursor movements should also be nearly the same on these two occasions. In fact, they were not the same at all.

I designed several other tests of the input-output model of tracking, and the results were always exactly those predicted by PCT: no effect of perception on behavior; behavior is the control of perception. The basic assumption of experimental psychology—indeed, the basic assumption of all social science is wrong. This was heady stuff. But the excitement was tempered considerably by my growing realization that work on PCT was going to be very lonely indeed. As I began to present the results of my research to other psychologists (in publications, at meetings and seminars), it became increasingly clear that, while psychologists love to talk about scientific revolutions and to call every new theory in psychology "revolutionary," they don't want a real revolution—and you don't get much more revolutionary than PCT. My presentations on PCT were met with polite interest and, sometimes, nodding agreement, but it was clear that no one really wanted to stop what they were doing and start psychology all over again, from scratch.

It was also becoming clear that there were not many psychologists besides myself who were doing research based on PCT. In fact, the only PCT research publications of any quality that I knew of were by Powers himself. So I wrote to Powers in 1979 and went to visit him in 1980 (he was living in Northbrook, Illinois, at the time, relatively close to Minneapolis). Bill turned out to be as brilliant in person as on paper—and a truly wonderful human being too; kind, helpful, humble—surprising qualities in a person who is just about always right about everything. Through Bill, I learned that there were some other scientific psychologists actively interested in PCT. Eventually, we developed a bit of a network of PCT afficionados. With the invaluable assistance of Bill's wife Mary (a very accomplished PCT afficionado herself), this disorganized group of scholars, who shared little more than an interest in PCT, finally got together in one place—a retreat near Kenosha, Wisconsin—for the first meeting of the Control Systems Group.

In 1985, I left teaching and returned with my family to California. I left teaching only because I could not, in good conscience, continue to hypocritically teach a curriculum that had to be taught if students were to learn "psychology." I could have stayed at Augsburg as long as I liked—and I was encouraged to stay—teaching one or two "special" courses a year on control theory. But I felt that this was not fair to the students or to control theory. When I did teach such courses, students wondered why I was teaching a course that challenged everything they were being taught in the other psychology classes; it seemed as if I were engaged in a personal feud with my colleagues. I also found that teaching PCT in the context of the conventional psychology curriculum gave the impression that PCT is a new explanation for the "facts" being learned in the other classes; it took me several years to realize that this is actually not the case—that PCT is a totally new approach to understanding behavior, a new start for psychology. Existing psychological "facts" are not facts at all, from the PCT perspective: they are usually based on statistical data, so they are not true "all of the time" (often not even a good proportion of the time), and they are not true of any individual person, but only of a non-existent "average person." I realized that in order to do PCT properly, one has to stop doing conventional psychology and start doing PCT—period. In 1985, I stopped doing conventional psychology.

Since leaving teaching, I have made my living as a "human factors" engineer by day while continuing my PCT research at night and on weekends—time permitting. I have managed to publish several papers on PCT since leaving teaching, but my interest in publishing in the conventional psychology journals has almost completely evaporated. Not only is it nearly impossible to get past the review process with a PCT paper, there is virtually no response to these papers when they are published. I am no longer surprised or saddened by this response to PCT; it is quite understandable in PCT terms; psychologists can be expected to control for doing psychology in ways that achieve their higher-order goals—which seem to include publications, recognition by peers, tenured faculty positions, and best-selling textbooks. PCT is obviously not a way of doing psychology that will help a psychologist achieve these goals; in fact, PCT is a disturbance to the kind of psychology that does allow psychologists to achieve these goals. Efforts to "convert" psychologists to PCT are no more likely to be successful than efforts to convert believers to atheists (or vice versa). Nevertheless, it is still fun to discuss and argue about PCT concepts, and the most exciting forum for doing this now is on the Internet.

In 1990, a computer network dedicated to discussion of PCT was formed; there are now approximately 120 people in at least five countries participating in this network, known as CSGnet (the Control Systems Group network). This network might not increase the number of "converts" to PCT, but it *will* provide a forum for sharing ideas and results that come out of PCT research and modeling.

I see two important paths for the future development of PCT. One is the scientific path: much more research needs to be done on the basic PCT model. If I were the head of the Living Systems Research Institute, with many graduate students to help with the research, I would have no shortage of projects to suggest. I think it's important to study the control of higher-order variables—sequences, categories, programs, and even principles. I have begun some simple studies of the ability to control sequences and programs. These studies should be perfected and extended, and a start should be made at modeling the perceptual functions involved in the control of these complex perceptual variables. I would also like to perfect methods for monitoring the value of the reference for a controlled variable. It is important to be able to distinguish variation in a controlled variable that results from poor control from variation that is intended. I also think it is important to study intra- and inter-personal conflict in some detail. Conflict is the basic human problem, from a PCT perspective; we have to understand its essentials in order to know how to deal with it.

The second path is therapeutic. PCT implies a specific approach to therapy based on the idea of getting consciousness "above" the level of the internal conflict—to the level of the systems that are setting the incompatible goals. It should be possible to teach and apply this approach to therapy clearly and consistently. We need to develop therapists who can reliably apply the "method of levels" and who can teach it to others. This means that PCT therapists will have to understand the science of PCT at least as well as the scientists understand the therapy. In fact, PCT should break down the barriers between scientific and clinical approaches to psychology. Any person who is able to do PCT therapy should also be able to do at least some basic PCT science, and any person who is able to do PO' science should be able to do some basic PCT therapy. The difference between PCT science and therapy should only be a difference in emphasis, not a difference in scientific integrity or human compassion.

Of course, there are many more directions in which PCT can expand, but I see them all turning around these two poles—the scientific and the therapeutic. Much more needs to be done with modeling complex,

multidimensional control processes; for example, a model hand might be a nice sequel to the Little Man's pointing arm. This kind of modeling will probably be of most interest to those traveling down the scientific path—but those on the therapeutic path would do well to try to understand why such models work. There are also great possibilities for PCT in the realm of social relations; PCT principles should allow people to develop social organizations that allow people to maintain individual control—to the collective benefit of all individuals inside and outside of the organization. This is a "therapeutic" application of PCT that would surely benefit from the scientific modeling of group behavior using PCT.

Perceptual Control Theory: Looking Back, Looking Forward

David M. Goldstein
(801 Edgemoor Rd., Cherry Hill, NJ 08034)

Looking Back

I finally became a pioneer in something! I guess this is what happens if you live long enough. Let me put on my coonskin hat and remember.

When I was in graduate school at the University of Connecticut in 1973 or 1974, Michael Turvey alluded to Bill Powers in one of his graduate perception courses, but he couldn't really explain how a control system works which made any sense to us (or him). Later on, in a paper with Carol Fowler, he revealed his lack of understanding in a clear way. I sent Bill a copy of the Turvey and Fowler paper. Bill wrote a very long letter trying to explain where the authors went wrong. They never answered. From the exchange between them, I came to appreciate the difference between substance and style.

Tom Bourbon got me into this! His students kept asking me how Jean Piaget related to Bill Powers, whom I didn't remember hearing about in graduate school (but actually did). I read *Behavior: The Control of Perception* out of self-defense and became addicted. At the time when this happened, I was an assistant professor at the Stephen F. Austin State University in Nacogdoches, Texas, and I taught graduate and undergraduate courses. I also had a part-time private practice of psychology.

At that time, I didn't understand Pa well enough to apply it to doing psychotherapy. But I did see how it related to biofeedback therapy. The paper which I wrote in Wayne Hershberger's book summarizes some of these ideas.

I invited Bill Powers to talk at SFA, and he accepted! The reaction of the other faculty members was really interesting and pretty typical of the reactions we have come to expect. Bill's description of Nacogdoches as "the backwater of the world" still sticks in my mind. (How should we describe Durango?) Somewhere in the official records of the SFA newspaper is an article describing Bill's visit. Unfortunately, the video tape made of this event was lost.

Corresponding with Bill by mail on different topics has kept up my interest in PCT. It takes a long time to understand this PCT stuff! We used to say "two years." Perhaps with all the new computer demos, books, tapes, etc., it does not take as long. I was amazed by how patient Bill was when explaining PCT ideas. I was impressed by the fact that he seemed to have thought about most of the questions/issues which occurred to me in my field and had come to some conclusions about them. His willingness to think about fields of knowledge far from his own reminds me of the idea of a Renaissance man. These qualities have become obvious to all those participating in CSGnet.

The first meeting of the control theory group which I recall attending was in Philadelphia in 1981. This was actually part of an American Society for Cybernetics meeting. I got to meet Rick Marken, and we had fun trying to get a computer program to do a pursuit-tracking task on my Commodore-64. Rick's special talents with modeling and his colorful personality were apparent even then. I presented a pursuit-tracking study at that meeting; the research subjects were special education students. The data were analyzed with a transfer function approach devised by Bill. The parameter estimation/modeling approach which we use today did not exist then.

I think I attended all of the annual CSG meetings in Wisconsin except for the last one. The following experiences stand out in my memory:

(1) During one of the meetings, the group name "Control Systems Group" came into use and stuck.

(2) Dick Robertson and I collaborated on applying the PCT approach to looking at the self-image. We explored the idea that the self-image is a systems-level controlled perception. We were not able to get the two papers we wrote on the subject published. One was a research-oriented approach, while the other one was a theoretical integration of the self-concept area. (Dick, should we try *Closed Loop?)*

(3) Ed Ford introduced us to the idea of Quality Time for improving relationships and took us on some very pleasant long walks. I learned about Reality Therapy from him, as well as from Diane Gossen and Perry Good. For a while, it looked as though Bill Powers and Bill Glasser would make a dynamic duo. But....

(4) Dick Robertson, Clark McPhail, Chuck Taylor, and I had some fun times playing tennis during the afternoon breaks.

(5) I started using the Q-Methodology approach as a statistical way to identify a person's self-image. In fact, we described Bill Powers using this approach and obtained the different ways in which subgroups of people at one meeting viewed him. Bill thought it was an interesting "projective technique." From Bill's reactions to this, given his lack of fondness for statistics, I have come to the conclusion that the intensive study of the individual case is the best chance we have of finding "facts." The question Bill always asks: For what percentage of people

will this statement be true? For what percentage of time will this be true of a person?

For this essay, I analyzed the results of the Q-Methodology study. Bill Power's self-image and five different subgroups can be compared as follows:

Bill's self-image	*Group 1 (n = 9)*	*Group 2 (n = 6)*	*Group 3 (n = 7)*	*Group 4 (n = 5)*	*Group 5 (n = 5)*
obstinate (+3)	yes	no	no	no	no
good-natured (+2) yes	yes	yes	yes	yes	
assertive (+2)	yes	yes	yes	yes	yes
quarrelsome (-2)	yes	yes	yes	yes	yes
resigned (-2)	yes	yes	yes	yes	yes
submissive (-3)	yes	no	no	yes	no

From the PCT perspective, the group which knows Bill best is the one which matches his self-image most closely. This translates to 1, then 4, then 2, 3, and 5.

Here, the descriptor "obstinate" was defined by the sentence: "It is difficult to get me to do something I don't want to do." The descriptor "submissive" was defined by the sentence: "I do what other people want me to do."

The results lead one to think about the relationship between a person's theory and the person's self-image. There is obvious agreement among the groups in how they perceive Bill, but it seems clear that group 1 knows Bill the best, that is, comes the closest to his self-image.

(6) I started to apply the Q-Methodology approach to therapy cases. One such case was written up and published in *Operant Subjectivity,* the journal of the Q-Methodology people. A more general look at Q-Methodology from a PCT viewpoint was also published in this journal in a second paper.

The publication of the PCT textbook edited by Dick Robertson and Bill Powers provoked me to want to go back to teaching introductory psychology. I taught one course at Glassboro State College as an adjunct. I was allowed to use the Pa book as long as it was supplemented with a "standard" one. The reactions of the students were interesting. They found PCT understandable but challenging. All exams were take-home essays, and I have kept the answers for future reference.

The creation of CSGnet was the brainchild of Gary Cziko. I first heard about it at the CSG meeting which took place in Pennsylvania. The net has done wonders by allowing people scattered all over the country and world with an interest in PCT to talk and learn from each other. It captures the feeling of the CSG meetings in terms of the intellectual stimulation and willingness to listen to people from other fields of interest.

Unfortunately, a combination of practical factors have made it impossible for me to attend the meetings taking place in Durango. I have greatly missed them. Thank God for CSGnet and the telephone!

At one point I had become discouraged that the PCT approach was so hard to apply to therapy cases. I presented a clinical case on CSGnet, after which Bill offered to teach me about the method of levels by applying it to me. This took place over a number of months. It resulted in a more sophisticated methodology for studying self-image than I previously had used and persuaded me that PCT had some unique contributions to make to therapy.

The "method of levels" plays an important role in PCT psychotherapy. It is a way of raising a person's awareness so that he/she can become aware of background experiences (perceptions) to the one which started out the conversation. As a result of attending to background experiences, a person's awareness rises to higher levels of perception.

The method of levels is the process that Bill Powers went through which resulted in the different levels in PCT. The method of levels is a "bottom-to-top" procedure. The therapist starts where the patient's awareness is ("bottom") and, by looking for background experiences, helps to move the patient's awareness higher ("top"). There is no assumption that the specific levels mentioned in Bill Power's books are the ones which will be found for a particular case.

Those who want to see what PCT psychotherapy looks like when the specific levels are used as the basis of therapy should read Ed Ford's *Freedom From Stress.* Ed follows a "top-to-bottom" strategy. He has people identify the important system-level experiences in their life. Then, for each system-level concept, he has them identify the important principle-level perceptions which are the means of achieving it. Then, for each principle-level perception, he has them identify the important program-level perceptions to achieve a given principle-level perception.

For the past several months, Bill Powers has been doing the method of levels with me. We have communicated using e-mail. Here are some of the things I learned as a result of doing the method of levels with Bill. I was the "patient" and Bill was the "therapist." We went through two rounds of the method. One start-off topic was my reaction to the method of levels as I understood it. The second start-off topic was my reaction to the experience of playing tennis. Here's what I learned:

(1) Each statement which the patient makes has potential background experiences which the therapist can ask the patient to address.

For some reason, I used to think that the background experiences would only show up after rather large segments of conversation. Each statement goes into and through the therapist, who is looking for background experiences along with the patient.

When Bill and I were doing the method of levels, the notational convention emerged to put the background material in brackets. (Like this.] Bill started doing this. Then, when I was writing my e-mail post to him, I would put background experiences which I noticed in brackets as I became more sensitized to what a background experience was like.

(2) The background experience feels more like an observation than an inference. I got the best results if I could observe the feeling or thought which was in the background. It did not feel as if I drew a conclusion or made an inference. It felt as if I made an observation. This helped to give me confidence that the background experience was something which was just as real as the topic which started the conversation. Prior to doing this method-of-levels exercise, I had my doubts about the reality of the background experience.

(3) The therapist is much more active in identifying the background experience than I understood to be the case from Bill's general description of the process before we did the exercise. I don't think that, in most cases, the patient will be doing this on his/her own, at least in the beginning. As the process goes on, the patient does become better at identifying background stuff. The therapist identifies a background experience for the patient and asks if the patient wants to address it, or prefers a different topic, or prefers to continue on the same topic.

The length of the therapist's answer makes a big difference. If it is too short, the patient feels alone in the enterprise. If it is too long, the patient is focusing too much on the therapist.

(4) The identification of a background experience feels a lot different than receiving an interpretation. Bill and I wound up calling this more traditional approach to therapy "psychologizing." The result of psychologizing was that I felt annoyed to have to address stuff which seemed to come out of the blue from Bill. When we were following the method of levels, I felt as though I was addressing my stuff. Psychologizing reliably resulted in blocking the flow of the conversation and progress. I am sure that giving interpretations has useful roles in therapy, but I am more aware of the negative side-effects it can have than I was before the exercise.

(5) The method of levels is not as abstract or difficult to do as I had thought. In fact, I observed that really good ordinary conversations sometimes follow the method of levels. One person says something. The other person tunes into the background stuff and addresses it. I no longer believe it is restricted to highly intelligent, verbal adults who are intellectually oriented. In fact, since the therapist meets the patient

wherever the patient's awareness is located, it probably is applicable to any age group for which other verbal therapy approaches would be attempted. The therapist has to be willing and able to adjust to the state of the patient.

(6) There are no fireworks emotionally or intellectually when the level of awareness is raised. The changes feel much more subtle. The patient is not transformed into someone who the patient never was. It is true that the patient might become aware of stuff of which he/she was formerly unaware. As I went up levels, the feeling aspects of the experience seemed to diminish in intensity. At the lower levels, the feelings were stronger and more salient parts of the experience.

(7) I did become aware of an internal conflict. Becoming aware of the conflict did not result in the immediate resolution of the conflict. I did give myself a daily assignment which I carry out to help me resolve the conflict. I did not previously identify this internal conflict. I can see how it has resulted in some significant inconsistencies in the way I am/behave. If this were a real therapy session, the therapist would probably have to spend time helping the patient figure out ways to resolve the conflict, once it was identified.

(8) The end result of the exercise was to start to examine my self-image. When we got to this point, the method of levels was more difficult to apply. It was here that I decided to continue the exploration on my own and that the joint exercise has stopped. I am now applying the self-image exercise procedure, based on PCT ideas, which I presented at the last CSG meeting I attended.

Recently, I have applied a more sophisticated version of the Q-Methodology studies, taking into account some of the criticisms expressed by Bill and other CSG people. Instead of using items consisting of single words drawn from a standardized set, I use sentences unique to the person being studied. And I resurrected the how/why technique, which I had presented at one of the CSG annual meetings, to have the subjects take each sentence and generate meaningfully related sentences. All of the sentences created became the universe from which a smaller set of sentences, about 20 to 50, would be chosen for the sort.

A second innovation is the way I have been selecting "conditions of instruction"—the sorting instructions given to the subject whose self-image I am studying. The conditions of instruction are chosen so as to sample as widely as possible from the different emotionally packed episodes which have been discussed in therapy. For example, I might instruct a patient, "Describe the way you are at the time of your divorce." Or, "Describe the way you are when you are riding your bicycle."

The interesting thing about the in-depth self-image studies I have conducted so far is that multiple self-images have emerged. It might

be "normal" to be multiple. However, at a level "above" the multiple self-images in most people (even people with multiple personalities), there is a single "observer" self with many of the characteristics of the reorganization system. I hope that one of these therapy case studies will see the light of day in a journal.

For the past three years, I have been the Clinical Director in an adolescent residential treatment center in New Jersey. For the first time, I have been able to apply PCT ideas on a wider basis than in private practice or the classroom. From this experience, I am beginning to learn how to "soft-sell" PCT to clinicians who have different viewpoints and to others. I have introduced 'Post-Critical Incident Counseling," which is PCT-based, brief (15 to 30 minutes), and fills the gap between our behavior-modification-based point/status system and the traditional therapies.

Looking Forward

Now I'll exchange my coonskin hat for my herbal tea leaves (no caffeine, please), with which I shall forecast the future with unerring accuracy.

Closed Loop will become a "real" journal. This seems to be happening already. The participants on CSGnet do not seem to be at a loss for words. I see more and more PCT research being done. People of the PCT persuasion will become more and more involved in following their own hunches. They will become less self-conscious and defensive and feel less of a need to persuade others of the merits of PCT. The research will speak for itself and attract others. People from all walks of study will want to publish in the *Journal of Living Control Systems.* PCT will become the equivalent of the universal language in Hesse's Glass Bead Game *(Magister Ludi).*

The PCT approach will be applied at the biochemical level. Advances in genetics research will combine with PCT ideas. A perceptual signal is "a copy of" a reference signal, just as DNA can create copies of itself. Do we have control systems operating in the genome? Bill Powers is already working with one person in this area. It is very exciting!

In the tea leaves, I see a set of neuropyschological tests based on PCT ideas. The levels of the control system hierarchy are calling out for someone to make them into a set of tests and, at the same time, test some of Bill's ideas about levels and relationships among levels.

A PCT research institute will be established.

PCT tasks will be utilized in research studies even by non-PCTers. They will be impressed by the ability to predict performance in tracking tasks. This will become a tool which they will use, and they will relate the performance tasks to all kinds of things which PCTers wouldn't.

PCT methodology will become more accepted and refined. The intensive study of the individual case will become the way to go. Unlike behavior-modification people, PCT people will study controlled perceptions. The methodology of researchers and clinicians will be merged into one new scientific approach when studying living control systems.

I'll write a book on the PCT approach in clinical psychology. Preliminary title: *Everything I Know about Psychology 1 Learned in Kindergarten or after Graduate School.* (There is something in me which wants to see merit in what non-PCTers have done and are doing. This results in my being less pure than some other PCTers. Oh, well!)

New people will take over part of the functions which Bill Powers has been doing all by himself for all of these years. Fortunately for us, Bill has longevity in his family. However, with all of the "young Turks" coming on board, he will gladly let go of some of his functions. His wiseman function will, however, be retained.

The PCT approach will become widely known: I see PCT as being *the* approach of the future. We have a common language in terms of which people from a diversity of fields can talk to each other.

Bill Powers will live to 100+, will give a keynote invited address at an American Psychological Association annual conference, and will be recognized as the one of the greats in psychology. Finally, Bill will receive the recognition he deserves from the old guard. The history of psychology will become divided into pre- and post-PCT—BC and AC, for short.

And we all will live happily ever after. The people of the world, starting with parents, will stop trying to use brute force to control other people. We will all become very sophisticated at peaceful ways of conflict resolution.

Perceptual Control Theory at 40

William T. Powers
(73 Ridge Place, CR 510, Durango, CO 81301)

As this issue of *Closed Loop* is the first one carrying the subtitle *Journal of Living Control Systems,* readers encountering our approach to this subject for the first time might need an overview of perceptual control theory (PCT) to get started. So this paper will be Yet Another Introduction to PCT. I will slant it, however, toward those coming into to this subject from the physical sciences; the relationship of PCT to physical approaches has been discussed at some length lately on CSGnet.

Rather than just reviewing the history or the principles of PCT, I'll try to develop an argument that leads from conventional views of behavior to the new view that PCT gives us, emphasizing in the end the odd role that organisms, seen through the eyes of PCT, play in a world otherwise dominated by physical laws. The point will be to show that control theory provides us with the germ of a radically new understanding, a break with all traditional theories of behavior—and many new ones as well. The future progress of PCT depends on understanding just how different a view of behavior we get by understanding the logic of control, the logic of a controlling organism's relationship to its environment.

The Etiology of Perceptual Control Theory

All living systems are sensitive to their environments; all act on their environments. This is ancient knowledge. The puzzle presented to the behavioral scientist is only how that sensitivity becomes converted into action. What are the rules, if any?

The most obvious and straightforward scientific approach to this question was realized long ago. In the physical sciences, if you want to know the properties of an assemblage of matter, you apply experimental forces and other influences to the object and observe what it does as a consequence.

In the worlds of physics and chemistry, this is a relatively easy task. Objects tend to be simple and have few properties; they are normally homogeneous or made of simple repeating units. It is not hard to make sure that experimental effects on them are the only effects of any importance. All similar objects made of the same materials behave in essentially the same way, and they will continue to do so no matter how many times an experiment is repeated—in fact, measurements of properties can be almost indefinitely refined by repeating them. A physical or chemical experiment can be clearly described and can be replicated by anyone who wishes to check the results. The reasoning about the meaning of an experiment can be communicated in clear and formal language, and even the reasoning process itself can be made public by being expressed in mathematical terms that anyone can learn. The history of a material object is entirely expressed in its present condition; the path by which it got into that condition is irrelevant, and only the current environment is of any importance in determining what will happen in the future to that piece of matter.

These confidence-building thoughts about the physical-science approach were, of course, tried out on organisms. The results were anything but confidence-building. A behavioral scientist reading the preceding paragraph might well experience mounting despair and envy of the physicist. While it is true that organisms are made of matter and must therefore obey all of the laws of physics and chemistry, it is not true that they are homogeneous or made of simple repeating units. They are, in fact, immensely more complex internally than the objects studied by physicists and chemists. They are too sensitive to their environments for any scientist to be sure of having control of everything important that happens to them. Not only are they sensitive, but they adjust themselves internally to external circumstances. It is not possible to perform the same experiment over and over on an organism to refine measurements of its properties—just imagine giving the same physics test over and over to refine a determination of a student's state of knowledge of physics, or giving a weight-lifting test to an athlete, day after day, to refine measurements of the athlete's strength.

The initial attempts to apply the methods of physical science to organisms were moderately successful at answering questions about perception. When the same methods were extended to the study of behavior, the results were not so encouraging—in comparison with expectations, they could only be called failures. Organisms were so subject to unpredictable influences, it seemed, that extraordinary precautions had to be taken to eliminate unwanted and unpredicted behaviors. This seemed at first to be a technical problem, to be overcome by greater attention to controlling the environment during experiments. As the years went by, however, it became apparent that no amount of attention to detail was enough. Not even the simplest phenomenon, such as a blink of an eye in response to a puff of air, could be made to occur with complete reliability. More complex behaviors simply went

all over the map. The dream of creating "Newton's laws of behavior" was apparently unattainable.

This did not cause a loss of faith in the methods of physics. Most behavioral scientists continued to assume that behavior was created by environmental influences. This assumption led to an attempt to find suggestions of regularity in behavior through statistical means, and then to a conclusion that this was the only possible means of exploring behavior, because behavior is inherently variable. The basic concept was retained: what organisms do is caused by what is done to them by the surrounding environment. But the requirements for formal language, public means of reasoning, ability of anyone to reproduce results, and refinement of measurements by continued experimentation were mostly impracticable. Despite the failure of the physical-science approach, the assumption was that the failure of organisms to behave as predictably as planets was due to technical difficulties, not errors in basic principles. The alternative conclusion, that something was wrong with applying the physical principle of cause and effect to the behavior of organisms, was simply not considered.

This alternative eventually came into play by a roundabout path.

Control theory was invented by engineers of the 1930s trying to build devices that would behave like human beings carrying out a specific kind of task: a control task. Even though the engineers did not realize it (many still do not realize it), the concept of control introduces a new principle, one that denies the basic idea that organisms do what the environment makes them do. While cause and effect still work in control theory as anywhere else, the organization of a control system creates apparent cause-effect dependencies that are different from the actual ones. Part of understanding control processes in organisms is the understanding that conventional cause-effect interpretations can be more misleading than informative.

Organisms are sensitive to their environments, and they act on their environments. The old assumption was that the sensing was the primary process, with the acting following from it. But that is an arbitrary assumption. It is just as plausible to assume that the acting is the primary process, and that the sensing, at least in certain critical regards, follows from the acting. It is even more plausible to say that sensing and acting are processes that go on simultaneously, in continuing streams that can't be clearly separated into cause and effect. This is basically what the inventors of control theory discovered: a type of system in which behavior affects the inputs on which behavior appears to depend. This is the type of system they had to use to imitate the human behavior called controlling.

This discovery led eventually to cybernetics, which endorsed this concept of closed causation without exploring more than its general philosophical implications. Years had to pass before more detailed implications came to light. Still more years had to pass before the basic concepts of control theory could be boiled down to a systematic model of control behavior—now called PCT—that could replace the old systematic cause-effect model based on the approach of physics.

The basic difference between the physical approach and that of control theory is that the physical approach deals with properties of energy and matter, while control theory deals with the properties of *particular organizations* of energy and matter. George Herbert Mead pointed out early in this century that physics doesn't deal with forms, with the entities into which we divide the world of experience. The physicist explores what is the *same* between a horse-cart and an ox-cart. The systems approach is concerned with what is unique to each vehicle, with differences in behavior brought about not by the differing physical or chemical composition of different objects, but by the differing organization of forms made of the same materials differently arranged.

It stands to reason, therefore, that physical laws will have a different significance when seen in the context of an organized system. We can admit that they make the behavior of the system possible, without also admitting that they *explain* the behavior of the system. Physics and chemistry can explain how it is that a neural signal liberates energy that causes a muscle fiber to contract, and how it is that this contraction leads to accelerations, velocities, and positions of limb segments connected to a joint spanned by the muscle. But they can't explain how it is that this signal arises under just these circumstances to reach that particular muscle. Physics and chemistry can't even be applied until the organization is specified. It is at the level of organizational understanding that control theory confronts older conceptions of the organization of behavior in living systems.

The Phenomenon of Control

I am going to avoid semantic arguments about what "control" really is. I will use the term in a particular sense; if others interpret it in a different sense, they will have difficulty following this exposition. I use it in this sense: A system is said to control a variable if it acts on that variable, in the presence of other unpredictable influences of comparable size on the same variable, so as to maintain the variable in an arbitrary state. The "arbitrary state" might mean a state of constancy, or any arbitrary pattern of change. The critical aspect of this definition is that physical influences that normally account entirely for the state of the variable are no longer effective, while the action of the control system causes the variable to behave independently of those other physical influences. When that is true, the variable is called a controlled variable.

The first important fact about control to notice is that the controlled variable is being acted upon by many forces, only one of which is attributable to the control system. The driver of a car can apply a lateral force to the front end of a car by turning the steering wheel. But there are many other influences that create forces acting laterally on the car at the same time: crosswinds, bumps in the road, tilts in the roadbed, unevenly inflated tires, and asymmetries in the aerodynamics of the car's shape, to mention a few.

If we observe, as we commonly do, that the path of the car does not follow strictly from the sum of all of the external forces acting on the car's mass, we can only conclude that it is the driver's contribution that makes the difference. If we see the car moving in a straight line, we can only conclude that the sum of all forces, including the one that the driver can alter, is zero.

So, if any of the external influences is seen to vary, but the path of the car does not vary as Newton's laws and engineering principles would predict, we have to deduce that the driver must be producing a varying force that just cancels the sum of the external forces. Indeed, we can observe the driver continually making adjustments of the steering wheel angle, while the car continues in a straight, or very nearly straight, line.

Likewise, if we observe the car moving along a smooth curve, but we see that the sum of all extraneous forces would tend to make it move along some other path, we can deduce that the varying forces created by the driver add just enough more force in just the right way to produce the curved path. If we see the car moving along a straight expressway, then turning off to take an exit ramp, then making other turns until it ends up parked in a parking lot, we can be quite sure that normal external forces would not have made the car follow just that path (an easily tested assumption). We can be sure that the varying forces created by the driver's motor actions on the steering wheel must have been exactly those necessary to add to the natural forces to create this overall result.

To anyone accustomed to normal physical or engineering analyses of the motions of objects, there must be a jarring note in this account. What is generally done is to observe all of the independent contributing forces and the initial conditions, and then to deduce through physical laws what the resulting motion must be. The driver's steering forces and the external forces due to winds, road tilts, and so forth simply occur as they occur, and the car's path is the outcome.

But here we are speaking as if one of the determining forces, the varying force being generated by the driver, is being adjusted so as to create a *preselected* outcome. Instead of the outcome varying randomly as the unrelated applied forces make it vary, the outcome conforms to some predetermined pattern. One of the causal forces which adds to the other forces continually changes in just the way needed to maintain that pattern. We would appear to be saying, and we are in fact saying, that the outcome we observe is being produced *on purpose.*

The vast majority of behavioral scientists has always rejected this interpretation. When the concepts of PCT were first being developed, this resistance was massive and almost universal (it is considerably less today). To say that outcomes are produced intentionally has seemed to most scientists to call for a reversal of cause and effect, or for giving the future an effect on the present. Many have argued that if all of the causal influences are known, the outcome must be whatever it is, and to call it "intentional" adds no explanatory power. Clearly, the outcome is an effect of converging causes, not a cause of the converging causes. Even if behavior does seem to entail intended outcomes, a scientist must stick to normal cause and effect, and find some other explanation.

There have been centuries of attempts to find some other explanation. But prior to the advent of control theory, all other explanations, we now know, were spurious. Even now there are many who strongly resist admitting that outcomes are indeed intended, and that organisms are the loci of these intentions. This resistance is misplaced, because now we can explain exactly how it is that an outcome can be controlled.

How Control Works

Once again: All living systems are sensitive to their environments; all act on their environments. So far we have talked only about actions and other physical influences on the environment. To see how control works, we must now talk about how organisms sense their environments.

Sensing is a process by which an external variable comes to be represented as a neural (or chemical) signal inside an organism. This looks like normal physical causation, but it is not like most causal processes. There is *amplification* involved. Metabolic processes in an organism maintain the sensing nerve-endings in hair-trigger states of readiness to fire. Only a tiny added stimulus is needed to cause a neural impulse to be generated, and metabolic processes instantly restore the sensor to the brink of firing again. So a small continuing stimulus causes the sensory nerve ending to fire again and again, at a frequency that corresponds to the amount of stimulation. The signals that leave the nerve ending involve the expenditure of many times the energy that causes the sensory ending to fire, nearly all of the energy being supplied from stores within the organism itself.

These neural signals can be further amplified, and eventually they can be routed to effectors such as muscles that provide a final amplification up to levels that can have significant effects on physical

processes in the environment. The result is that organisms can create physical forces of large magnitudes which are produced without any significant reverse effect on the physical variables being sensed. This creates a novel relationship between the organism's output forces and other physical processes.

I remember inventing my first perpetual motion machine, at the age of perhaps 12. I had read that a certain kind of motor could be used either as a motor or as a generator. So I thought of putting fan blades on two of these motors and using one to blow air onto the other, the idea being that the generator would supply the current needed to run the motor while the motor supplied the wind that would run the generator. It took a few more years of education to realize that one has to think of physical processes quantitatively, not just qualitatively. It makes a difference *how much air* can be blown, and *how much* current can be generated, and *how fast* the driving fan can be spun by the available current. High school physics was enough to show me the embarrassing truth: that in physical systems, there are balances that are maintained: balances of forces, of momenta, and of energies. The world studied by physicists is rigorously constrained by these balances, these conservation laws. You can't get any more out of a physical system than goes into it. This is how I and most other people learned to think about physical processes.

This is also true of organisms, of course. No more energy comes out than goes in. But the energy that goes in is of a different form from the energy that comes out it is the chemical energy in food and air, obtained independently of the physical processes involved in behavior, and stored for future conversion into actions. So when an organism, a person, comes across some natural physical process in its environment, it is in a position to throw a monkey-wrench into the machinery by spending some of its store of energy.

Let's switch examples now. Suppose a person sees a fat child and a thin child sitting on opposite ends of a teeter-totter. The end with the fat child on it is, of course, on the ground, and the thin child is high in the air. The upward force of the ground on the fat child's side, plus the upward force from the thin child pressing down on the other end, just equal the fat child's weight. The physical system is in equilibrium.

Now the person places a hand on the thin child's end of the teeter-totter and pushes down, spending a bit of metabolic energy from the last few days' meals and several thousand breaths of air. The thin child descends and the fat child rises. If the amount of downward push follows a certain law, the teeter-totter will end up horizontal and stationary again.

What is the required law? If the force applied is large when the fat child is low and small when the fat child is high, with a continuous transition between the two states, there will be one state in the middle where the force is just right to bring the teeter-totter to the horizontal with all forces in equilibrium. But what could make the force applied by this helpful person follow that law?

Suppose we tried to mechanize this effect. When the fat child's end goes down, a cable pulls a weight at the center of the teeter-totter toward the thin child's end, and vice versa. The history of perpetual motion machines is full of such clever devices. All such devices, however intricate and devious their designs, fail because you can't get more out of a physical system than went into it.

But the person helping balance the teeter-totter is exerting a force of just the right amount without any linkage from the teeter-totter that produces that force. The only link from the teeter-totter to the person is through the person's visual sense, which registers the angle of the teeter-totter as feeble neural signals inside the person's brain. This requires only intercepting some of the light reflected from the physical apparatus and the children, a process that supplies only an infinitesimal amount of energy to the person and exerts no measurable force at all, either way.

The neural signals that now represent the angle of the teeter-totter are further amplified, and they finally enter muscles where the greatest (by far) amplification of all occurs, producing a force that acts downward on the teeter-totter. This force is greatest when the fat child is accelerating upward, smallest when accelerating downward. Stored energy is used by the person in applying the force to the moving teeter-totter. That's vital; none of this could work without the independent source of energy that comes from the eggs and roast beef and peanut butter sandwiches that the person has been eating.

What happens in the end is that the neural signals representing the angle of the teeter-totter come to some particular state representing the horizontal position, and the force applied to the teeter-totter is just the difference in weight of the two children. The physical system is now being maintained in a state far from equilibrium, but if you include the helpful person in the physical system, everything is in equilibrium again: forces, momenta, and energy inputs and outputs.

The factor that determines where this equilibrium will occur is now in the person, not the external physical system. There is some particular condition of the sensory signals that corresponds to the observed equilibrium. If the sensory signals indicate a deviation from this condition, the force will either increase or decrease in the direction that tends to restore the equilibrium. The rule is simple: if the angle slopes downward toward the fat child, increase the force; if upward, decrease it. This rule, which is applied inside the brain of the person, is what determines the equilibrium point.

There's one more factor to consider. The person balancing the teeter-totter might decide to maintain the board at some angle other than horizontal. This amounts to redefining the condition of equilibrium. In a control system model, this is done by providing an adjustable reference signal against which the signal representing angle can be compared. This occurs inside the person's brain. The final amplification of signals that drives the muscles is applied to the *difference* between the reference and sensory signals, so the opposition to even small deviations from equilibrium can be very strong.

With the addition of the variable reference signal, the person can now cause the teeter-totter to behave in any arbitrary way at all, as long as the available muscle forces are large enough and the person doesn't exhaust the stores of metabolic energy. As the reference signal varies, the teeter-totter's angle varies in exact correspondence. It can be made to vary regularly or irregularly, quickly or slowly, with or without a child sitting on either end—or not at all, even though the children climb on and off the board. The angle of the teeter-totter is now completely determined by a reference signal inside the person's brain, and the normal physics of the teeter-totter is totally overridden. The person is inserting extra force, extra momentum, and extra energy—whatever is required to make the desired behavior appear.

This same analysis could have been applied to the driver of the car. The lateral position of the car is represented in the driver's brain as some sort of neural signal. Another neural signal, a reference signal, specifies the lateral position that is to be maintained, and amplification of the difference between the two signals produces muscle forces that act on the car to make its lateral position, as sensed, match the specified position. Varying the reference signal will then cause the lateral position of the car to change in a parallel way, independently of other forces acting on the car. The normal physics of car motion is overridden; external forces lose their determining effects.

Organisms in Control

In the world of physics, there are physical objects linked to each other by properties of the environment and physical laws that cause the behavior of one object to depend on the behavior of other objects. Even in the most complex of physical systems, there is a kind of natural bookkeeping that accounts for all of the interactions. The sum of all forces acting on and inside the system, counting both actions and reactions, is zero. The sum of all changes in energy content, including energy inputs from outside and energy outputs to the outside, is zero. All momenta add up to zero, or at least a constant.

If we want to make one variable in a physical system depend on another one, the normal approach is to establish a physical link. This link connects forces from one object to another object, which involves transfers of energy and momentum and sometimes flows of matter. The new link participates in the balances of the system; it can generate no new energy, and it can create no unbalanced forces. The affected object is in physical equilibrium with the affecting object. If A is pushing on B through the new linkage, then B is pushing back on A with exactly the same force.

An organism is, of course, a physical system subject to all of the same laws and balances. But the organism can create linkages among objects in its environment which, at first glance, seem to violate physical principles.

First, the organism can move about in its environment and dispose itself to create forces on many different objects in many different ways. This means it is in a position to affect objects that are not normally affected by such actions.

Second, the organism can orient its sensors to create internal signals representing many aspects of physical objects around it. The visual sense is particularly potent in this regard: simply by looking in different directions, the organism can create internal signals that stand for the states of objects in many different ways: their position, velocity, size, color, relation to other objects, shape, and so forth. It can do this without affecting those objects in any measurable way.

As a result, an organism can position its muscles and limbs, and its sensory apparatus, in ways that create arbitrary linkages between the objects it can sense and the objects to which it can apply forces. Furthermore, because of the high amplification that takes place inside the organism, this linkage can be made one-way—that is, one object can be made to affect another object without being affected by the reverse path through the same link. There is a violation of the normal energy balance in the physical system, because any normal physical link would require energies, forces, and so on to remain in balance.

The unbalances are made up by the organism from its internal energy stores, and from the way it braces itself against the world as it exerts forces. If we consider the physical environment and the organism as a single system, there is, of course, no violation of any physical principles. The point, however, is that the physical environment linked to an organism can no longer be treated as if no organism were present.

Consider the car and driver again. With no driver in the car, but with the car rolling along the road, physical influences on the car can be calculated according to normal physical principles. From the speed and direction of the wind and the aerodynamic properties of the car, the wind force acting on the car can be calculated. Similarly, forces arising from tilts and bumps and soft tires can be calculated. All of these forces

can be added up, and their effects on the car can be computed. From these forces and the properties of the car and road, the motion of the car can be computed with, in principle, as much exactness as we please.

But now put a driver in the car. Suddenly, the path of the car ceases to follow from the sum of all external forces and the properties of the car and the road. Instead, we find that a new physical linkage has been created. Now when the wind blows and the road tilts, the result is a movement of the steering wheel which *prevents* the car from obeying the physical laws that previously applied.

Even more important, we find that the physical linkage that has been created is *not* between the steering wheel and the wind or the tilt of the road, but between the steering wheel and the lateral position of the car. What the driver is sensing is the *outcome* of all of the applied forces (which now include the effects of turning the steering wheel). The driver watches the visual appearance of the hood of the car against the road ahead and acts to maintain that visual appearance in a specified state (either constant or changing in a specified way). The only thing that gives the car's lateral position an effect on the path of the car is the fact that the driver is sensing that lateral position, internally specifying an intended state for that perception, and producing steering forces based on the difference between what is actually sensed and what is intended to be sensed.

From outside the driver, this critical perceptual linkage is invisible, undetectable in terms of any changes in the physical world. Nothing in the world changes measurably because of being sensed. Nothing in the physical outside world indicates the driver's internal reference signal that specifies the intended state of the perception. As far as any external measurements are concerned, the force that turns the steering wheel has no observable external physical cause. It is an arbitrary force generated for no physically observable reason.

The strangest thing about this force is that after it is added to all of the other independent forces that are applied to the car at the same time, the result is an outcome that is repeatable with great accuracy for long periods of time, even if the external forces change and even if there are changes in the properties of the car and the road. When all of the external forces change, the outcome does not change; instead, the remaining force applied to the car changes in just the way that keeps the outcome the same. The cause changes in order that the effect be preserved.

An organism can attend to any perceivable aspect of the environment. If the forces that the organism can generate are comparable to the external forces that exist at the same time, that aspect of the environment can be made to conform to the organism's intention for it, and to cease behaving as the natural forces on it would otherwise dictate.

The actions of the controlling organism supersede the physical laws that normally govern that part of the environment, in the respect that the organism is controlling.

Conclusions

Organisms are physical systems, and they exist in a physical world. But the laws of physics do not explain their behavior or its effects on the physical world. Organisms force the world around them into highly improbable forms, states of motion, and organization, and they act in a way that keeps normal physical forces from having their normal effects. It is organization, not physics, that explains how they do this.

To understand human behavior in these new terms is to seek a kind of explanation completely different from what behavioral scientists, modeling their approach after physics, have sought. This is what PCT is about, and where its promise for the future lies.

The Control Systems Group is a membership organization which supports the understanding of cybernetic control systems in organisms and their environments: *living control systems.* Academicians, clinicians, and other professionals in several disciplines, including biology, psychology, social work, economics, education, engineering, and philosophy, are members of the Group. Annual meetings have been held since 1985. The CSG Business Office is located at 73 Ridge Pl., CR 510, Durango, CO 81301; phone (303) 247-7986.

The CSG logo shows the generic structure of cybernetic control systems. A Comparator (C) computes the difference between a reference signal (represented by the arrow coming from above) and the output signal from Sensory (S) computation. The resulting difference signal is the input to the Gain generator (G). Disturbances (represented by the black box) alter the Gain generator output on the way to Sensory computation, where the negative-feedback km/ is closed.

This reproduction of *Closed Loop* was created by Dag Forssell in 2009. Addresses and phone numbers have not been updated. Most are obsolete.

Posted at www.pctresources.com

Spell checked, but not proofread. Please report any errors to webmaster. Email address at website.

Back cover

Closed Loop

Journal of Living Control Systems

Spring 1994

Volume 4

Number 2

Front cover

Closed Loop

Journal of Living Control Systems

Spring 1994 Volume 4 Number 2

Edited by Greg Williams, 460 Black Lick Rd., Gravel Switch, KY 40328

CONTENTS

Members of the Control Systems Group receive Closed Loop quarterly. For additional information, contact Mary Powers, 73 Ridge Place, CR 510, Durango, CO 81301, phone (303)247-7986.

Some Closed Loop articles are adapted from the electronic mailing list CSGnet. To subscribe to CSGnet, e-mail Gary Cziko at g-cziko@uiuc.edu or phone him at (217)3338527.

From the Editor

Nobody sent any *new* papers for publication in this issue of *Closed Loop,* despite my impassioned plea in the last issue, so I have reprinted an *old* paper—so old that it is now in the public domain, and its author is long dead. "What Is Man?" was envisioned by its author as a serious treatise on psychology. It originally appeared in 1906 as an anonymously published and privately printed volume with a press run of only 250 copies. It was most certainly not a best-seller. In fact, it was virtually ignored until two days after the death in 1910 of the famous humorist Samuel L. Clemens, when the New York *Tribune* ran an article announcing that it had been penned by none other than Mark Twain himself. Since then, various commentators and critics have often suggested either a central or an ephemeral place for "What Is Man?" in the Clemens canon, with no consensus opinion. But there is no question about the way Clemens viewed it: he was determined—despite the protestations of his wife—to present its arguments to several of the most influential thinkers of his era, and he termed "What Is Life?" his "Bible" and "gospel."

When I serendipitously came upon "What Is Man?" a few years ago, I was immediately fascinated by the apparent parallels between certain ideas expressed in that work and certain tenets of perceptual control theory (PCT). Some of these parallels should be evident if "What Is Man" is considered in the light of the other article in this issue, adapted from CSGnet postings. In that article, Rick Marken and Bill Powers address, from the standpoint of PCT, issues of human autonomy and responsibility similar to those central to the dialogue in "What Is Man?"

In nontrivial ways, the PCT position appears to me to be close to that of Twain's Old Man: humans are machines acting according to hereditary (for Twain, "temperament") and environmental influences (perhaps overstressed by Twain at times), and not according to an unconstrained free will responsible only to the whims of a transcendent Self or Soul; there are no purely unselfish or altruistic actions, only actions which lead to inner satisfaction (of perceptual reference states, PCTers would quickly add); human behavior is not fundamentally different from the behavior of other animals, but humans—sometimes because of inadequate psychological models!—have often used their more highly refined abilities in ignoble ways impossible for "lower" animals. There are substantial differences as well, mostly due to the closed-loop basis for PCT: organisms are machines of a far different sort than Twain suspected. Despite its defects, I value "What Is Man?" as an intriguing and entertaining exploration of human psychology which should be especially interesting to PCTers because in some places it almost reads like a parody of PCT. I even suspect that Samuel Clemens would have wanted to join the Control Systems Group, had he lived long enough! Perhaps his well-documented long-term error signals with regard to the "damned human race" would have been assuaged if he had only known that PCT was going to appear?

I welcome comments from anyone with error signals of their own resulting from my inclusion of "What Is Man?" in this journal. Any and all contributions will be considered for possible publication in a future *Closed Loop.* (Yes, you guessed it. I will indeed go to great lengths to incite you to write *something* for *Closed Loop,* even if that something is vitriolic.)

For the copy-text of "What Is Man?" I used *What Is Man? and Other Essays,* first published in 1917 by Harper & Brothers (New York). I am heavily indebted to Dag and Christine Forssell for providing me with a computer-readable version. An amended version of the text, including manuscript sections not previously published, was edited by Paul Baender and published by the University of California Press in 1973 in *"What Is Man?" and Other Philosophical Writings.* Below is a list of selected references for those interested in learning more about the history and critical reception of "What Is Man?"

References

Allen, J. (1954). *The Adventures of Mark Twain.* Boston: Little, Brown. See pages 294-298, where it is argued that "What Is Man?" resulted from efforts by Clemens to determine the ultimate roots of human malice ("something the higher animals did not have"), despised by Clemens as both cruel and unnecessary.

Emerson, E. (1985). *The Authentic Mark Twain: A Literary Biography of Samuel Clemens.* Philadelphia: University of Pennsylvania Press. See pages 218-219 for a discussion of sources (including Darwin, Huxley, Lecky's History of European Morals, and William James) for some of the ideas in "What Is Man?"

Kaplan, J. (1966). *Mr. Clemens and Mark Twain: A Biography*. New York: Simon and Schuster. See page 340 for comments on an 1897 letter from Clemens to a British psychologist, in which Clemens enthusiastically outlined the main ideas in "What Is Man?" Kaplan also says that Mrs. Clemens refused to hear parts of "What Is Man?"

Lauber, J. (1990). *The Inventions of Mark Twain.* New York: Hill and Wang. See pages 288-289, documenting correspondence between Clemens and H. G. Wells regarding "What Is Man?"—Clemens was "deeply disturbed" when Wells noted that the claim that humans are machines was not an original idea.

Wagner-Martin, L. (1993). *What Is Man?* In J. R. LeMaster & J. D. Wilson (eds.), *The Mark Twain Encyclopedia* (pp. 783-785). New York: Garland. Wagner-Martin suggests that Clemens saw "the human condition worsening rather than improving despite material advantages and sophisticated technology" and wrote "What Is Man?" upon realizing that "all the gentle moralizing of his earlier much acclaimed works" came to almost nought.

What Is Man?

Mark Twain (1835 - 1910)

I

a. MAN THE MACHINE. *b.* PERSONAL MERIT

[The Old Man and the Young Man had been conversing. The Old Man had asserted that the human being is merely a machine, and nothing more. The Young Man objected, and asked him to go into particulars and furnish his reasons for his position.]

Old Man. What are the materials of which a steam-engine is made?

Young Man. Iron, steel, brass, white-metal, and so on.

O. M. Where are these found?

Y. M. In the rocks.

O. M. In a pure state?

Y. M. No—in ores.

O. M. Are the metals suddenly deposited in the ores?

Y. M. No—it is the patient work of countless ages.

O. M. You could make the engine out of the rocks themselves?

Y. M. Yes, a brittle one and not valuable.

O. M. You would not require much, of such an engine as that?

Y. M. No—substantially nothing.

O. M. To make a fine and capable engine, how would you proceed?

Y. M. Drive tunnels and shafts into the hills; blast out the iron ore; crush it, smelt it, reduce it to pig-iron; put some of it through the Bessemer process and make steel of it. Mine and treat and combine the several metals of which brass is made.

O. M. Then?

Y. M. Out of the perfected result, build the fine engine.

O. M. You would require much of this one?

Y. M. Oh, indeed yes.

O. M. It could drive lathes, drills, planers, punches, polishers, in a word all the cunning machines of a great factory?

Y. M. It could.

O. M. What could the stone engine do?

Y. M. Drive a sewing-machine, possibly—nothing more, perhaps.

O. M. Men would admire the other engine and rapturously praise it?

Y. M. Yes.

O. M. But not the stone one?

Y. M. No.

O. M. The merits of the metal machine would be far above those of the stone one?

Y. M. Of course.

O. M. Personal merits?

Y. M. *Personal* merits? How do you mean?

O. M. It would be personally entitled to the credit of its own performance?

Y. M. The engine? Certainly not.

O. M. Why not?

Y. M. Because its performance is not personal. It is a result of the law of its construction. It is not a *merit* that it does the things which it is set to do—it can't *help* doing them.

O. M. And it is not a personal demerit in the stone machine that it does so little?

Y. M. Certainly not. It does no more and no less than the law of its make permits and compels it to do. There is nothing *personal* about it; it cannot choose. In this process of "working up to the matter" is it your idea to work up to the proposition that man and a machine are about the same thing, and that there is no personal merit in the performance of either?

O. M. Yes—but do not be offended; I am meaning no offense. What makes the grand difference between the stone engine and the steel one? Shall we call it training, education? Shall we call the stone engine a savage and the steel one a civilized man? The original rock contained the stuff of which the steel one was built—but along with it a lot of sulphur and stone and other obstructing inborn heredities, brought down from the old geologic ages—prejudices, let us call them. Prejudices which nothing within the rock itself had either *power* to remove or any *desire* to remove. Will you take note of that phrase?

Y. M. Yes. I have written it down: "Prejudices which nothing within the rock itself had either power to remove or any desire to remove." Go on.

O. M. Prejudices which must be removed by *outside influences* or not at all. Put that down.

Y. M. Very well; "Must be removed by outside influences or not at all." Go on.

O. M. The iron's prejudice against ridding itself of the cumbering rock. To make it more exact, the iron's absolute *indifference* as to whether the rock be removed or not. Then comes the *outside influence* and grinds the rock to powder and sets the ore free. The *iron* in the ore is still captive. An *outside influence* smelts it free of the clogging ore. The iron is emancipated iron, now, but indifferent to further progress. An

outside influence beguiles it into the Bessemer furnace and refines it into steel of the first quality. It is educated, now—its training is complete. And it has reached its limit. By no possible process can it be educated into *gold*. Will you set that down?

Y. M. Yes. "Everything has its limit—iron ore cannot be educated into gold."

O. M. There are gold men, and tin men, and copper men, and leaden men, and steel men, and so on—and each has the limitations of his nature, his heredities, his training, and his environment. You can build engines out of each of these metals, and they will all perform, but you must not require the weak ones to do equal work with the strong ones. In each case, to get the best results, you must free the metal from its obstructing prejudicial ores by education—smelting, refining, and so forth.

Y. M. You have arrived at man, now?

O. M. Yes. Man the machine—man the impersonal engine. Whatsoever a man is, is due to his *make*, and to the *influences* brought to bear upon it by his heredities, his habitat, his associations. He is moved, directed, COMMANDED, by *exterior* influences—*solely.* He *originates* nothing, not even a thought.

Y. M. Oh, come! Where did I get my opinion that this which you are talking is all foolishness?

O. M. It is a quite natural opinion—indeed an inevitable opinion—but you did not create the materials out of which it is formed. They are odds and ends of thoughts, impressions, feelings, gathered unconsciously from a thousand books, a thousand conversations, and from streams of thought and feeling which have flowed down into your heart and brain out of the hearts and brains of centuries of ancestors. *Personally* you did not create even the smallest microscopic fragment of the materials out of which your opinion is made; and personally you cannot claim even the slender merit of *putting the borrowed materials together.* That was done automatically—by your mental machinery, in strict accordance with the law of that machinery's construction. And you not only did not make that machinery yourself, but you have *not even any command over it.*

Y. M. This is too much. You think I could have formed no opinion but that one?

O. M. Spontaneously? No. And *you did not form that one;* your machinery did it for you—automatically and instantly, without reflection or the need of it.

Y. M. Suppose I had reflected? How then?

O. M. Suppose you try?

Y. M. *(After a quarter of an hour.)* I have reflected.

O. M. You mean you have tried to change your opinion—as an experiment?

Y. M. Yes.

O. M. With success?

Y. M. No. It remains the same; it is impossible to change it.

O. M. I am sorry, but you see, yourself, that your mind is merely a machine, nothing more. You have no command over it, it has no command over itself—it is worked *solely from the outside.* That is the law of its make; it is the law of all machines.

Y. M. Can't I *ever* change one of these automatic opinions?

O. M. No. You can't yourself, but *exterior influences* can do it.

Y. M. And exterior ones *only.*

O. M. Yes—exterior ones only.

Y. M. That position is untenable—I may say ludicrously untenable.

O. M. What makes you think so?

Y. M. I don't merely think it, I know it. Suppose I resolve to enter upon a course of thought, and study, and reading, with the deliberate purpose of changing that opinion; and suppose I succeed. *That* is not the work of an exterior impulse, the whole of it is mine and personal: for I originated the project.

O. M. Not a shred of it. *It grew out of this talk with me.* But for that it would never have occurred to you. No man ever originates anything. All his thoughts, all his impulses, come *from the outside.*

Y. M. It's an exasperating subject. The *first* man had original thoughts, anyway; there was nobody to draw from.

O. M. It is a mistake. Adam's thoughts came to him from the outside. *You* have a fear of death. You did not invent that—you got it from outside, from talking and teaching. Adam had no fear of death—none in the world.

Y. M. Yes, he had.

O. M. When he was created?

Y. M. No.

O. M. When, then?

Y. M. When he was threatened with it.

O. M. Then it came from the *outside.* Adam is quite big enough; let us not try to make a god of him. *None but gods have ever had a thought which did not come from the outside.* Adam probably had a good head, but it was of no sort of use to him until it was filled up *from the outside.* He was not able to invent the triflingest little thing with it. He had not a shadow of a notion of the difference between good and evil—he had to get the idea *from the outside.* Neither he nor Eve was able to originate the idea that it was immodest to go naked: the knowledge came in with the apple *from the outside.* A man's brain is so constructed that *it can originate nothing whatever.* It can only use material obtained *outside.* It is merely a machine; and it works automatically, not by will-power. *It has no command over itself, its owner has no command over it.*

Y. M. Well, never mind Adam: but certainly Shakespeare's creations—

O. M. No, you mean Shakespeare's *imitations.* Shakespeare created nothing. He correctly observed, and he marvelously painted. He exactly portrayed people whom *God* had created; but he created none himself. Let us spare him the slander of charging him with trying. Shakespeare could not create. *He was a machine, and machines do not create.*

Y. M. Where *was* his excellence, then?

O. M. In this. He was not a sewing-machine, like you and me; he was a Gobelin loom. The threads and the colors came into him *from the outside;* outside influences, suggestions, *experiences* (reading, seeing plays, playing plays, borrowing ideas, and so on), framed the patterns in his mind and started up its complex and admirable machinery, and *it automatically* turned out that pictured and gorgeous fabric which still compels the astonishment of the world. If Shakespeare had been born and bred on a barren and unvisited rock in the ocean his mighty intellect would have had no *outside material* to work with, and could have invented none; and *no outside influences,* teachings, moldings, persuasions, inspirations, of a valuable sort, and could have invented none; and so Shakespeare would have produced nothing. In Turkey he would have produced something—something up to the highest limit of Turkish influences, associations, and training. In France he would have produced something better—something up to the highest limit of the French influences and training. In England he rose to the highest limit attainable through the *outside helps afforded by that land's ideals, influences, and training.* You and I are but sewing-machines. We must turn out what we can; we must do our endeavor and care nothing at all when the unthinking reproach us for not turning out Gobelins.

Y. M. And so we are mere machines! And machines may not boast, nor feel proud of their performance, nor claim personal merit for it, nor applause and praise. It is an infamous doctrine.

O. M. It isn't a doctrine, it is merely a fact.

Y. M. I suppose, then, there is no more merit in being brave than in being a coward?

O. M. *Personal* merit? No. A brave man does not *create* his bravery. He is entitled to no personal credit for possessing it. It is born to him. A baby born with a billion dollars—where is the personal merit in that? A baby born with nothing—where is the personal demerit in that? The one is fawned upon, admired, worshiped, by sycophants, the other is neglected and despised—where is the sense in it?

Y. M. Sometimes a timid man sets himself the task of conquering his cowardice and becoming brave—and succeeds. What do you say to that?

O. M. That it shows the value of *training in right directions over training in wrong ones.* Inestimably valuable is training, influence, education, in right directions—*training one's self-approbation to elevate its ideals.*

Y. M. But as to merit—the personal merit of the victorious coward's project and achievement?

O. M. There isn't any. In the world's view he is a worthier man than he was before, but *he* didn't achieve the change—the merit of it is not his.

Y. M. Whose, then?

O. M. His *make,* and the influences which wrought upon it from the outside.

Y. M. His make?

O. M. To start with, he was *not* utterly and completely a coward, or the influences would have had nothing to work upon. He was not afraid of a cow, though perhaps of a bull: not afraid of a woman, but afraid of a man. There was something to build upon. There was a *seed.* No seed, no plant. Did he make that seed himself, or was it born in him? It was no merit of *his* that the seed was there.

Y. M. Well, anyway, the idea of *cultivating* it, the resolution to cultivate it, was meritorious, and he originated that.

O. M. He did nothing of the kind. It came whence *all* impulses, good or bad, come—from *outside.* If that timid man had lived all his life in a community of human rabbits, had never read of brave deeds, had never heard speak of them, had never heard any one praise them nor express envy of the heroes that had done them, he would have had no more idea of bravery than Adam had of modesty, and it could never by any possibility have occurred to him to *resolve* to become brave. He *could not originate the idea*—it had to come to him from the *outside.* And so, when he heard bravery extolled and cowardice derided, it woke him up. He was ashamed. Perhaps his sweetheart turned up her nose and said, "I am told that you are a coward!" It was not *he* that turned over the new leaf—she did it for him. *He* must not strut around in the merit of it—it is not his.

Y. M. But, anyway, he reared the plant after she watered the seed.

O. M. No. *Outside influences* reared it. At the command—and trembling—he marched out into the field—with other soldiers and in the daytime, not alone and in the dark. He had the *influence of example,* he drew courage from his comrades' courage; he was afraid, and wanted to run, but he did not dare; he was *afraid* to run, with all those soldiers looking on. He was progressing, you see—the moral fear of shame had risen superior to the physical fear of harm. By the end of the campaign experience will have taught him that not *all* who go into battle get hurt—an outside influence which will be helpful to him; and he will also have learned how sweet it is to be praised for courage and be huzza'd at with tear-choked voices as the war-worn regiment marches past the worshiping multitude with flags flying and the drums beating.

After that he will be as securely brave as any veteran in the army—and there will not be a shade nor suggestion of *personal merit* in it anywhere; it will all have come from the *outside.* The Victoria Cross breeds more heroes than—

Y. M. Hang it, where is the sense in his becoming brave if he is to get no credit for it?

O. M. Your question will answer itself presently. It involves an important detail of man's make which we have not yet touched upon.

Y. M. What detail is that?

O. M. The impulse which moves a person to do things—the only impulse that ever moves a person to do a thing.

Y. M. The *only* one! Is there but one?

O. M. That is all. There is only one.

Y. M. Well, certainly that is a strange enough doctrine. What is the sole impulse that ever moves a person to do a thing?

O. M. The impulse to *content his own spirit*—the *necessity* of contenting his own spirit and *winning its approval.*

Y. M. Oh, come, that won't do!

O. M. Why won't it?

Y. M. Because it puts him in the attitude of always looking out for his own comfort and advantage; whereas an unselfish man often does a thing solely for another person's good when it is a positive disadvantage to himself.

O. M. It is a mistake. The act must do *him* good, FIRST; otherwise he will not do it. He may *think* he is doing it solely for the other person's sake, but it is not so; he is contenting his own spirit first—the other person's benefit has to always take *second* place.

Y. M. What a fantastic idea! What becomes of self-sacrifice? Please answer me that.

O. M. What is self-sacrifice?

Y. M. The doing good to another person where no shadow nor suggestion of benefit to one's self can result from it.

II

MAN'S SOLE IMPULSE—THE SECURING OF HIS OWN APPROVAL

Old Man. There have been instances of it—you think?

Young Man. *Instances?* Millions of them!

O. M. You have not jumped to conclusions? You have examined them—critically?

Y. M. They don't need it: the acts themselves reveal the golden impulse back of them.

O. M. For instance?

Y. M. Well, then, for instance. Take the case in the book here. The man lives three miles up-town. It is bitter cold, snowing hard, midnight. He is about to enter the horse-car when a gray and ragged old woman, a touching picture of misery, puts out her lean hand and begs for rescue from hunger and death. The man finds that he has but a quarter in his pocket, but he does not hesitate: he gives it her and trudges home through the storm. There—it is noble, it is beautiful; its grace is marred by no fleck or blemish or suggestion of self-interest.

O. M. What makes you think that?

Y. M. Pray what else could I think? Do you imagine that there is some other way of looking at it?

O. M. Can you put yourself in the man's place and tell me what he felt and what he thought?

Y. M. Easily. The sight of that suffering old face pierced his generous heart with a sharp pain. He could not bear it. He could endure the three-mile walk in the storm, but he could not endure the tortures his conscience would suffer if he turned his back and left that poor old creature to perish. He would not have been able to sleep, for thinking of it.

O. M. What was his state of mind on his way home?

Y. M. It was a state of joy which only the self-sacrificer knows. His heart sang, he was unconscious of the storm,

O. M. He felt well?

Y. M. One cannot doubt it.

O. M. Very well. Now let us add up the details and see how much he got for his twenty-five cents. Let us try to find out the *real* why of his making the investment. In the first place *he* couldn't bear the pain which the old suffering face gave him. So he was thinking of *his* pain—this good man. He must buy a salve for it. If he did not succor the old woman *his* conscience would torture him all the way home. Thinking of *his* pain again. He must buy relief from that. If he didn't relieve the old woman *he* would not get any sleep. He must buy some sleep—still thinking of *himself,* you see. Thus, to sum up, he bought himself free of a sharp pain in his heart, he bought himself free of the tortures of a waiting conscience, he bought a whole night's sleep—all for twenty-five cents! It should make Wall Street ashamed of itself. On his way home his heart was joyful, and it sang—profit on top of profit! The impulse which moved the man to succor the old woman was—*first*—to *content his own spirit;* secondly to relieve *her* sufferings. Is it your opinion that men's acts proceed from one central and unchanging and inalterable impulse, or from a variety of impulses?

Y. M. From a variety, of course—some high and fine and noble, others not. What is your opinion?

O. M. Then there is but *one* law, one source.

Y. M. That both the noblest impulses and the basest proceed from that one source?

O. M. Yes.

Y. M. Will you put that law into words?

O. M. Yes. This is the law, keep it in your mind. *From his cradle to his grave a man never does a single thing which has any* FIRST AND FOREMOST *object but one—to secure peace of mind, spiritual comfort, for* HIMSELF.

Y. M. Come! He never does anything for any one else's comfort, spiritual or physical?

O. M. No. *Except on those distinct terms*—that it shall *first* secure *his own* spiritual comfort. Otherwise he will not do it.

Y. M. It will be easy to expose the falsity of that proposition.

O. M. For instance?

Y. M. Take that noble passion, love of country, patriotism. A man who loves peace and dreads pain, leaves his pleasant home and his weeping family and marches out to manfully expose himself to hunger, cold, wounds, and death. Is that seeking spiritual comfort?

O. M. He loves peace and dreads pain?

Y. M. Yes.

O. M. Then perhaps there is something that he loves *more* than he loves peace—*the approval of his neighbors and the public.* And perhaps there is something which he dreads more than he dreads pain—the *disapproval* of his neighbors and the public. If he is sensitive to shame he will go to the field—not because his spirit will be *entirely* comfortable there, but because it will be more comfortable there than it would be if he remained at home. He will always do the thing which will bring him the *most* mental comfort—for that is *the sole law of his life.* He leaves the weeping family behind; he is sorry to make them uncomfortable, but not sorry enough to sacrifice his *own* comfort to secure theirs.

Y. M. Do you really believe that mere public opinion could force a timid and peaceful man to—

O. M. Go to war? Yes—public opinion can force some men to do *anything.*

Y. M. *Anything?*

O. M. Yes—anything.

Y. M. I don't believe that. Can it force a rightprincipled man to do a wrong thing?

O. M. Yes.

Y. M. Can it force a kind man to do a cruel thing?

O. M. Yes.

Y. M. Give an instance.

O. M. Alexander Hamilton was a conspicuously high-principled man. He regarded dueling as wrong, and as opposed to the teachings of religion— but in deference to *public opinion* he fought a duel. He deeply loved his family, but to buy public approval he treacherously deserted them and threw his life away, ungenerously leaving them to lifelong sorrow in order that he might stand well with a foolish world. In the then condition of the public standards of honor he could not have been comfortable with the stigma upon him of having refused to fight. The teachings of religion, his devotion to his family, his kindness of heart, his high principles, all went for nothing when they stood in the way of his spiritual comfort. A man will do *anything,* no matter what it is, *to secure his spiritual comfort;* and he can neither be forced nor persuaded to any act which has not that goal for its object. Hamilton's act was compelled by the inborn necessity of contenting his own spirit; in this it was like all the other acts of his life, and like all the acts of all men's lives. Do you see where the kernel of the matter lies? A man cannot be comfortable without *his own* approval. He will secure the largest share possible of that, at all costs, all sacrifices.

Y. M. A minute ago you said Hamilton fought that duel to get *public* approval.

O. M. I did. By refusing to fight the duel he would have secured his family's approval and a large share of his own; but the public approval was more valuable in his eyes than all other approvals put together—in the earth or above it; to secure that would furnish him the *most* comfort of mind, the most *self*-approval; so he sacrificed all other values to get it.

Y. M. Some noble souls have refused to fight duels, and have manfully braved the public contempt.

O. M. They acted *according to their make.* They valued their principles and the approval of their families *above* the public approval. They took the thing they valued *most* and let the rest go. They took what would give them the *largest* share of *personal contentment and approval*—a man *always* does. Public opinion cannot force that kind of men to go to the wars. When they go it is for other reasons. Other spirit-contenting reasons.

Y. M. Always spirit-contenting reasons?

O. M. There are no others.

Y. M. When a man sacrifices his life to save a little child from a burning building, what do you call that?

O. M. When he does it, it is the law of *his* make. *He* can't bear to see the child in that peril (a man of a different make *could),* and so he tries to save the child, and loses his life. But he has got what he was after—*his own approval.*

Y. M. What do you call Love, Hate, Charity, Revenge, Humanity, Magnanimity, Forgiveness?

O. M. Different results of the one Master Impulse: the necessity of securing one's self-approval. They wear diverse clothes and are subject

to diverse moods, but in whatsoever ways they masquerade they are the *same person* all the time. To change the figure, the *compulsion* that moves a man—and there is but the one—is the necessity of securing the contentment of his own spirit. When it stops, the man is dead.

Y. M. This is foolishness. Love—

O. M. Why, love is that impulse, that law, in its most uncompromising form. It will squander life and everything else on its object. Not *primarily* for the object's sake, but for *its own*. When its object is happy *it* is happy—and that is what it is unconsciously after.

Y. M. You do not even except the lofty and gracious passion of mother-love?

O. M. No, *it* is the absolute slave of that law. The mother will go naked to clothe her child; she will starve that it may have food; suffer torture to save it from pain; die that it may live. She takes a living *pleasure* in making these sacrifices. *She does it for that reward*—that self-approval, that contentment, that peace, that comfort. *She would do it for your child* IF SHE COULD GET THE SAME PAY.

Y. M. This is an infernal philosophy of yours.

O. M. It isn't a philosophy, it is a fact.

Y. M. Of course you must admit that there are some acts which—

O. M. No. There is *no* act, large or small, fine or mean, which springs from any motive but the one—the necessity of appeasing and contenting one's own spirit.

Y. M. The world's philanthropists—

O. M. I honor them, I uncover my head to them—from habit and training; but *they* could not know comfort or happiness or self-approval if they did not work and spend for the unfortunate. It makes *them* happy to see others happy; and so with money and labor they buy what they are after—*happiness, self-approval.* Why don't misers do the same thing? Because they can get a thousandfold more happiness by *not* doing it. There is no other reason. They follow the law of their make.

Y. M. What do you say of duty for duty's sake?

O. M. That *it does not exist.* Duties are not performed for duty's *sake,* but because their *neglect* would make the man *uncomfortable.* A man performs but *one* duty—the duty of contenting his spirit, the duty of making himself agreeable to himself. If he can most satisfyingly perform this sole and only duty by *helping* his neighbor, he will do it; if he can most satisfyingly perform it by *swindling* his neighbor, he will do that. But he always looks out for Number One—*first;* the effects upon others are a *secondary* matter. Men pretend to self-sacrifices, but this is a thing which, in the ordinary value of the phrase, *does not exist and has not existed*. A man often honestly *thinks* he is sacrificing himself merely and solely for some one else, but he is deceived; his bottom impulse is to content a requirement of his nature and training, and thus acquire peace for his soul.

Y. M. Apparently, then, all men, both good and bad ones, devote their lives to contenting their consciences?

O. M. Yes. That is a good enough name for it: Conscience—that independent Sovereign, that insolent absolute Monarch inside of a man who is the man's Master. There are all kinds of consciences, because there are all kinds of men. You satisfy an assassin's conscience in one way, a philanthropist's in another, a miser's in another, a burglar's in still another. As a *guide* or *incentive* to any authoritatively prescribed line of morals or conduct (leaving *training* out of the account), a man's conscience is totally valueless. I know a kind-hearted Kentuckian whose self-approval was lacking—whose conscience was troubling him, to phrase it with exactness—*because he had neglected to kill a certain man*—a man whom he had never seen. The stranger had killed this man's friend in a fight, this man's Kentucky training made it a duty to kill the stranger for it. He neglected his duty—kept dodging it, shirking it, putting it off, and his unrelenting conscience kept persecuting him for this conduct. At last, to get ease of mind, comfort, self-approval, he hunted up the stranger and took his life. It was an immense act of *self-sacrifice* (as per the usual definition), for he did not want to do it, and he never would have done it if he could have bought a contented spirit and an unworried mind at smaller cost. But we are so made that we will pay *anything* for that contentment—even another man's life.

Y. M. You spoke a moment ago of *trained* consciences. You mean that we are not *born* with consciences competent to guide us aright?

O. M. If we were, children and savages would know right from wrong, and not have to be taught it.

Y. M. But consciences can be *trained?*

O. M. Yes.

Y. M. Of course by parents, teachers, the pulpit, and books.

O. M. Yes—they do their share; they do what they can.

Y. M. And the rest is done by—

O. M. Oh, a million unnoticed influences—for good or bad: influences which work without rest during every waking moment of a man's life, from cradle to grave.

Y. M. You have tabulated these?

O. M. Many of them—yes.

Y. M. Will you read me the result?

O. M. Another time, yes. It would take an hour.

Y. M. A conscience can be trained to shun evil and prefer good?

O. M. Yes.

Y. M. But will prefer it for spirit-contenting reasons only?

O. M. It *can't* be trained to do a thing for any *other* reason. The thing is impossible.

Y. M. There *must* be a genuinely and utterly self-sacrificing act recorded in human history somewhere.

O. M. You are young. You have many years before you. Search one out.

Y. M. It does seem to me that when a man sees a fellow-being struggling in the water and jumps in at the risk of his life to save him—

O. M. Wait. Describe the *man.* Describe the *fellow-being.* State if there is an *audience* present; or if they are *alone.*

Y. M. What have these things to do with the splendid act?

O. M. Very much. Shall we suppose, as a beginning, that the two are alone, in a solitary place, at midnight?

Y. M. If you choose.

O. M. And that the fellow-being is the man's daughter?

Y. M. Well, n-no—make it some one else.

O. M. A filthy, drunken ruffian, then?

Y. M. I see. Circumstances alter cases. I suppose that if there was no audience to observe the act, the man wouldn't perform it.

O. M. But there is here and there a man who *would.* People, for instance, like the man who lost his life trying to save the child from the fire; and the man who gave the needy old woman his twenty-five cents and walked home in the storm—there are here and there men like that who would do it. And why? Because they couldn't *bear* to see a fellow-being struggling in the water and not jump in and help. It would give *them* pain. They would save the fellow-being on that account. *They wouldn't do it otherwise.* They strictly obey the law which I have been insisting upon. You must remember and always distinguish the people who *can't bear* things from the people who *can.* It will throw light upon a number of apparently "self-sacrificing" cases.

Y. M. Oh, dear, it's all so disgusting.

O. M. Yes. And so true.

Y. M. Come—take the good boy who does things he doesn't want to do, in order to gratify his mother.

O. M. He does seven-tenths of the act because it gratifies *him* to gratify his mother. Throw the bulk of advantage the other way and the good boy would not do the act. He *must* obey the iron law. None can escape it.

Y. M. Well, take the case of a bad boy who—

O. M. You needn't mention it, it is a waste of time. It is no matter about the bad boy's act. Whatever it was, he had a spirit-contenting reason for it. Otherwise you have been misinformed, and he didn't do it.

Y. M. It is very exasperating. A while ago you said that a man's conscience is not a born judge of morals and conduct, but has to be taught and trained. Now I think a conscience can get drowsy and lazy, but I don't think it can go wrong; and if you wake it up—

A Little Story

O. M. I will tell you a little story:

Once upon a time an Infidel was guest in the house of a Christian widow whose little boy was ill and near to death. The Infidel often watched by the bedside and entertained the boy with talk, and he used these opportunities to satisfy a strong longing of his nature—that desire which is in us all to better other people's condition by having them think as we think. He was successful. But the dying boy, in his last moments, reproached him and said:

"I believed, and was happy in it; you have taken my belief away, and my comfort. Now I have nothing left, and I die miserable; for the things which you have told me do not take the place of that which I have lost."

And the mother, also, reproached the Infidel, and said:

"My child is forever lost, and my heart is broken. How could you do this cruel thing? We have done you no harm, but only kindness; we made our house your home, you were welcome to all we had, and this is our reward."

The heart of the Infidel was filled with remorse for what he had done, and he said:

"It was wrong—I see it now; but I was only trying to do him good. In my view he was in error; it seemed my duty to teach him the truth"

Then the mother said:

"I had taught him, all his little life, what I believed to be the truth, and in his believing faith both of us were happy. Now he is dead—and lost; and I am miserable. Our faith came down to us through centuries of believing ancestors; what right had you, or any one, to disturb it? Where was your honor, where was your shame?"

Y. M. He was a miscreant, and deserved death!

O. M. He thought so himself, and said so.

Y. M. Ah—you see, *his conscience was awakened!*

O. M. Yes, his Self-Disapproval was. It *pained* him to see the mother suffer. He was sorry he had done a thing which brought *him* pain. It did not occur to him to think of the mother when he was misteaching the boy, for he was absorbed in providing *pleasure* for himself, then. Providing it by satisfying what he believed to be a call of duty.

Y. M. Call it what you please, it is to me a case of *awakened conscience.* That awakened conscience could never get itself into that species of trouble again. A cure like that is a *permanent* cure.

O. M. Pardon—I had not finished the story. We are creatures of *outside influences*—we originate *nothing* within. Whenever we take a new line of thought and drift into a new line of belief and action, the impulse is

always suggested from the *outside.* Remorse so preyed upon the Infidel that it dissolved his harshness toward the boy's religion and made him come to regard it with tolerance, next with kindness, for the boy's sake and the mother's. Finally he found himself examining it. From that moment his progress in his new trend was steady and rapid. He became a believing Christian. And now his remorse for having robbed the dying boy of his faith and his salvation was bitterer than ever. It gave him no rest, no peace. He *must* have rest and peace—it is the law of our nature. There seemed but one way to get it; he must devote himself to saving imperiled souls. He became a missionary. He landed in a pagan country ill and helpless. A native widow took him into her humble home and nursed him back to convalescence. Then her young boy was taken hopelessly ill, and the grateful missionary helped her tend him. Here was his first opportunity to repair a part of the wrong done to the other boy by doing a precious service for this one by undermining his foolish faith in his false gods. He was successful. But the dying boy in his last moments reproached him and said:

"I believed, and was happy in it; you have taken my belief away, and my comfort. Now I have nothing left, and I die miserable; for the things which you have told me do not take the place of that which I have lost."

And the mother, also, reproached the missionary, and said:

"My child is forever lost, and my heart is broken. How could you do this cruel thing? We had done you no harm, but only kindness; we made our house your home, you were welcome to all we had, and this is our reward."

The heart of the missionary was filled with remorse for what he had done, and he said:

"It was wrong—I see it now; but I was only trying to do him good. In my view he was in error; it seemed my duty to teach him the truth...

Then the mother said:

"I had taught him, all his little life, what I believed to be the truth, and in his believing faith both of us were happy. Now he is dead—and lost; and I am miserable. Our faith came down to us through centuries of believing ancestors; what right had you, or any one, to disturb it? Where was your honor, where was your shame?"

The missionary's anguish of remorse and sense of treachery were as bitter and persecuting and unappeasable, now, as they had been in the former case. The story is finished. What is your comment?

Y. M. The man's conscience was a fool! It was morbid. It didn't know right from wrong.

O. M. I am not sorry to hear you say that. If you grant that *one* man's conscience doesn't know right from wrong, it is an admission that there are others like it. This single admission pulls down the whole doctrine of infallibility of judgment in consciences. Meantime there is one thing which I ask you to notice.

Y. M. What is that?

O. M. That in both cases the man's *act* gave him no spiritual discomfort, and that he was quite satisfied with it and got pleasure out of it. But afterward when it resulted in *pain* to *him,* he was sorry. Sorry it had inflicted pain upon the others, *but for no reason under the sun except that their pain gave* HIM *pain.* Our consciences take *no* notice of pain inflicted upon others until it reaches a point where it gives pain to *us*. In *all* cases without exception we are absolutely indifferent to another person's pain until his sufferings make us uncomfortable. Many an infidel would not have been troubled by that Christian mother's distress. Don't you believe that?

Y. M. Yes. You might almost say it of the *average infidel*, I think.

O. M. And many a missionary, sternly fortified by his sense of duty, would not have been troubled by the pagan mother's distress—Jesuit missionaries in Canada in the early French times, for instance; see episodes quoted by Parkman.

Y. M. Well, let us adjourn. Where have we arrived?

O. M. At this. That we (mankind) have ticketed ourselves with a number of qualities to which we have given misleading names. Love, Hate, Charity, Compassion, Avarice, Benevolence, and so on. I mean we attach misleading *meanings* to the names. They are all forms of self-contentment, self-gratification, but the names so disguise them that they distract our attention from the fact. Also we have smuggled a word into the dictionary which ought not to be there at all—Self-Sacrifice. It describes a thing which does not exist. But worst of all, we ignore and never mention the Sole Impulse which dictates and compels a man's every act: the imperious necessity of securing his own approval, in every emergency and at all costs. To it we owe all that we are. It is our breath, our heart, our blood. It is our only spur, our whip, our goad, our only impelling power; we have no other. Without it we should be mere inert images, corpses; no one would do anything, there would be no progress, the world would stand still. We ought to stand reverently uncovered when the name of that stupendous power is uttered.

Y. M. I am not convinced.

O. M. You will be when you think.

III

INSTANCES IN POINT

Old Man. Have you given thought to the Gospel of Self-Approval since we talked?

Young Man. I have.

O. M. It was I that moved you to it. That is to say an *outside influence* moved you to it—not one that originated in your own head. Will you try to keep that in mind and not forget it?

Y. M. Yes. Why?

O. M. Because by and by in one of our talks, I wish to further impress upon you that neither you, nor I, nor any man ever originates a thought in his own head. *The utterer of a thought always utters a second-hand one.*

Y. M. Oh, now—

O. M. Wait. Reserve your remark till we get to that part of our discussion—to-morrow or next day, say. Now, then, have you been considering the proposition that no act is ever born of any but a selfcontenting impulse— (primarily). You have sought. What have you found?

Y. M. I have not been very fortunate. I have examined many fine and apparently self-sacrificing deeds in romances and biographies, but—

O. M. Under searching analysis the ostensible self-sacrifice disappeared? It naturally would.

Y. M. But here in this novel is one which seems to promise. In the Adirondack woods is a wage-earner and lay preacher in the lumber-camps who is of noble character and deeply religious. An earnest and practical laborer in the New York slums comes up there on vacation—he is leader of a section of the University Settlement. Holme, the lumberman, is fired with a desire to throw away his excellent worldly prospects and go down and save souls on the East Side. He counts it happiness to make this sacrifice for the glory of God and for the cause of Christ. He resigns his place, makes the sacrifice cheerfully, and goes to the East Side and preaches Christ and Him crucified every day and every night to little groups of half-civilized foreign paupers who scoff at him. But he rejoices in the scoffings, since he is suffering them in the great cause of Christ. You have so filled my mind with suspicions that I was constantly expecting to find a hidden questionable impulse back of all this, but I am thankful to say I have failed. This man saw his duty, and for *duty's sake* he sacrificed self and assumed the burden it imposed.

O. M. Is that as far as you have read?

Y. M. Yes.

O. M. Let us read further, presently. Meantime, in sacrificing himself—*not* for the glory of God, *primarily,* as *he* imagined, but *first* to content that exacting and inflexible master within him—*did he sacrifice anybody else?*

Y. M. How do you mean?

O. M. He relinquished a lucrative post and got mere food and lodging in place of it. Had he dependants?

Y. M. Well—yes.

O. M. In what way and to what extent did his self-sacrifice affect *them?*

Y. M. He was the support of a superannuated father. He had a young sister with a remarkable voice-he was giving her a musical education, so that her longing to be self-supporting might be gratified. He was furnishing the money to put a young brother through a polytechnic school and satisfy his desire to become a civil engineer.

O. M. The old father's comforts were now curtailed?

Y. M. Quite seriously. Yes.

O. M. The sister's music-lessons had to stop?

Y. M. Yes.

O. M. The young brother's education—well, an extinguishing blight fell upon that happy dream, and he had to go to sawing wood to support the old father, or something like that?

Y. M. It is about what happened. Yes.

O. M. What a handsome job of self-sacrificing he did do! It seems to me that he sacrificed everybody *except* himself. Haven't I told you that no man *ever* sacrifices himself; that there is no instance of it upon record anywhere; and that when a man's Interior Monarch requires a thing of its slave for either its *momentary* or its *permanent* contentment, that thing must and will be furnished and that command obeyed, no matter who may stand in the way and suffer disaster by it? That man *ruined his family* to please and content his Interior Monarch—

Y. M. And help Christ's cause.

O. M. Yes—*secondly.* Not firstly. *He* thought it was firstly.

Y. M. Very well, have it so, if you will. But it could be that he argued that if he saved a hundred souls in New York—

O. M. The sacrifice of the *family* would be justified by that great profit upon the—the—what shall we call it?

Y. M. Investment?

O. M. Hardly. How would *speculation* do? How would *gamble* do? Not a solitary soul-capture was sure. He played for a possible thirty-three-hundred-percent profit. It was *gambling*—with his family for "chips." However, let us see how the game came out. Maybe we can get on the track of the secret original impulse, the *real* impulse, that moved him to so nobly self-sacrifice his family in the Saviour's cause under the superstition that he was sacrificing himself. I will read a chapter or so.... Here we have it! It was bound to expose itself sooner or later. He preached to the East-Side rabble a season, then went back to his old dull, obscure life in the lumber-camps *"hurt to the heart, his pride humbled."* Why? Were not his efforts acceptable to the Saviour, for Whom alone they were made? Dear me, that detail is *lost sight of,* is not even referred to, the fact that it started out as a motive is entirely forgotten! Then what is the trouble? The authoress quite innocently and unconsciously gives the whole business away. The trouble was this: this man merely *preached* to the poor; that is not the University

Settlement's way; it deals in larger and better things than that, and it did not enthuse over that crude Salvation-Army eloquence. It was courteous to Holme—but cool. It did not pet him, did not take him to its bosom. *"Perished were all his dreams of distinction, the praise and grateful approval of—"* Of whom? The Saviour? No; the Saviour is not mentioned. Of whom, then? Of "his *fellow-workers."* Why did he want that? Because the Master inside of him wanted it, and would not be content without it. That emphasized sentence quoted above, reveals the secret we have been seeking, the original impulse, the *real* impulse, which moved the obscure and unappreciated Adirondack lumberman to sacrifice his family and go on that crusade to the East Side which said original impulse was this, to wit: without knowing it *he went there to show a neglected world the large talent that was in him, and rise to distinction.* As I have warned you before, *no* act springs from any but the one law, the one motive. But I pray you, do not accept this law upon my say-so; but diligently examine for yourself. Whenever you read of a self-sacrificing act or hear of one, or of a duty done for *duty's sake,* take it to pieces and look for the *real* motive. It is always there.

Y. M. I do it every day. I cannot help it, now that I have gotten started upon the degrading and exasperating quest. For it is hatefully interesting! —in fact, fascinating is the word. As soon as I come across a golden deed in a book I have to stop and take it apart and examine it, I cannot help myself.

O. M. Have you ever found one that defeated the rule?

Y. M. No—at least, not yet. But take the case of servant-tipping in Europe. You pay the *hotel* for service; you owe the servants *nothing,* yet you pay them besides. Doesn't that defeat it?

O. M. In what way?

Y. M. You are not *obliged* to do it, therefore its source is compassion for their ill paid condition, and—

O. M. Has that custom ever vexed you, annoyed you, irritated you?

Y. M. Well—yes.

O. M. Still you succumbed to it?

Y. M. Of course.

O. M. Why of course?

Y. M. Well, custom is law, in a way, and laws must be submitted to —everybody recognizes it as a *duty*.

O. M. Then you pay the irritating tax for *duty's* sake?

Y. M. I suppose it amounts to that.

O. M. Then the impulse which moves you to submit to the tax is not *all* compassion, charity, benevolence?

Y. M. Well—perhaps not.

O. M. Is *any* of it?

Y. M. I—perhaps I was too hasty in locating its source.

O. M. Perhaps so. In case you ignored the custom would you get prompt and effective service from the servants?

Y. M. Oh, hear yourself talk! Those European servants? Why, you wouldn't get any at all, to speak of.

O. M. Couldn't *that* work as an impulse to move you to pay the tax?

Y. M. I am not denying it.

O. M. Apparently, then, it is a case of for-duty's-sake with a little self-interest added?

Y. M. Yes, it has the look of it. But here is a point: we pay that tax knowing it to be unjust and an extortion; yet we go away with a pain at the heart if we think we have been stingy with the poor fellows; and we heartily wish we were back again, so that we could do the right thing, and *more* than the right thing, the *generous* thing. I think it will be difficult for you to find any thought of self in that impulse.

O. M. I wonder why you should think so. When you find service charged in the *hotel* bill does it annoy you?

Y. M. No.

O. M. Do you ever complain of the amount of it?.

Y. M. No, it would not occur to me.

O. M. The *expense*, then, is not the annoying detail. It is a fixed charge, and you pay it cheerfully, you pay it without a murmur. When you came to pay the servants, how would you like it if each of the men and maids had a fixed charge?

Y. M. Like it? I should rejoice!

O. M. Even if the fixed tax were a shade *more* than you had been in the habit of paying in the form of tips?

Y. M. Indeed, yes!

O. M. Very well, then. As I understand it, it isn't really compassion nor yet duty that moves you to pay the tax, and it isn't the *amount* of the tax that annoys you. Yet *something* annoys you. What is it?

Y. M. Well, the trouble is, you never know *what* to pay, the tax varies so, all over Europe.

O. M. So you have to guess?

Y. M. There is no other way. So you go on thinking and thinking, and calculating and guessing, and consulting with other people and getting their views; and it spoils your sleep nights, and makes you distraught in the daytime, and while you are pretending to look at the sights you are only guessing and guessing and guessing all the time, and being worried and miserable.

O. M. And all about a debt which you don't owe and don't have to pay unless you want to! Strange. What is the purpose of the guessing?

Y. M. To guess out what is right to give them, and not be unfair to any of them.

O. M. It has quite a noble look—taking so much pains and using up

so much valuable time in order to be just and fair to a poor servant to whom you owe nothing, but who needs money and is ill paid.

Y. M. I think, myself, that if there is any ungracious motive back of it, it will be hard to find.

O. M. How do you know when you have not paid a servant fairly?

Y. M. Why, he is silent; does not thank you. Sometimes he gives you a look that makes you ashamed. You are too proud to rectify your mistake there, with people looking, but afterward you keep on wishing and wishing you *had* done it. My, the shame and the pain of it! Sometimes you see, by the signs, that you have hit it *just right,* and you go away mightily satisfied. Sometimes the man is so effusively thankful that you know you have given him a good deal *more* than was necessary.

O. M. *Necessary?* Necessary for what?

Y. M. To content him.

O. M. How do you feel *then?*

Y. M. Repentant.

O. M. It is my belief that you nave *not* been concerning yourself in guessing out his just dues, but only in ciphering out what would *content* him. And I think you had a self-deluding reason for that.

Y. M. What was it?

O. M. If you fell short of what he was expecting and wanting, you would get a look which would *shame you before folk.* That would give you *pain. You*—for you are only working for yourself, not *him*. If you gave him too much you would be *ashamed of yourself* for it, and that would give *you* pain—another case of thinking of *yourself,* protecting yourself, *saving yourself from discomfort.* You never think of the servant once—except to guess out how to get *his approval.* If you get that, you get your *own* approval, and that is the sole and only thing you are after. The Master inside of you is then satisfied, contented, comfortable; there was *no other* thing at stake, as a matter of *first* interest, anywhere in the transaction.

Further Instances

Y. M. Well, to think of it : Self-Sacrifice for others, the grandest thing in man, ruled out! nonexistent!

O. M. Are you accusing me of saying that?

Y. M. Why, certainly.

O. M. I haven't said it.

Y. M. What did you say, then?

O. M. That no man has ever sacrificed himself in the common meaning of that phrase—which is, self-sacrifice for another *alone.* Men make daily sacrifices for others, but it is for their own sake *first*. The act must content their own spirit *first*. The other beneficiaries come second.

Y. M. And the same with duty for duty's sake?

O. M. Yes. No man performs a duty for mere duty's sake; the act must content his spirit *first.* He must feel better for *doing* the duty than he would for shirking it. Otherwise he will not do it.

Y. M. Take the case of the *Berkeley Castle.*

O. M. It was a noble duty, greatly performed. Take it to pieces and examine it, if you like.

Y. M. A British troop-ship crowded with soldiers and their wives and children. She struck a rock and began to sink. There was room in the boats for the women and children only. The colonel lined up his regiment on the deck and said "it is our duty to die, that they may be saved." There was no murmur, no protest. The boats carried away the women and children. When the death-moment was come, the colonel and his officers took their several posts, the men stood at shoulder-arms, and so, as on dress-parade, with their flag flying and the drums beating, they went down, a sacrifice to duty for duty's sake. Can you view it as other than that?

O. M. It was something as fine as that, as exalted as that. Could you have remained in those ranks and gone down to your death in that unflinching way?

Y. M. Could I? No, I could not.

O. M. Think. Imagine yourself there, with that watery doom creeping higher and higher around you.

Y. M. I can imagine it. I feel all the horror of it. I could not have endured it, I could not have remained in my place. I know it.

O. M. Why?

Y. M. There is no why about it: I know myself, and I know I couldn't *do* it.

O. M. But it would be your *duty* to do it.

Y. M. Yes, I know—but I couldn't.

O. M. It was more than a thousand men, yet not one of them flinched. Some of them must have been born with your temperament; if they could do that great duty for duty's *sake,* why not you? Don't you know that you could go out and gather together a thousand clerks and mechanics and put them on that deck and ask them to die for duty's sake, and not two dozen of them would stay in the ranks to the end?

Y. M. Yes, I know that.

O. M. But you *train* them, and put them through a campaign or two; then they would be soldiers; soldiers, with a soldier's pride, a soldier's self-respect, a soldier's ideals. They would have to content a *soldier's* spirit then, not a clerk's, not a mechanic's. They could not content that spirit by shirking a soldier's duty, could they?

Y. M. I suppose not.

O. M. Then they would do the duty not for the *duty's* sake, but for their *own* sake—primarily. The *duty* was *just the same,* and just as imperative, when they were clerks, mechanics, raw recruits, but they wouldn't perform it for that. As clerks and mechanics they had other ideals, another spirit to satisfy, and they satisfied it. They *had* to; it is the law. *Training* is potent. Training toward higher and higher, and ever higher ideals is worth any man's thought and labor and diligence.

Y. M. Consider the man who stands by his duty and goes to the stake rather than be recreant to it.

O. M. It is his make and his training. He has to content the spirit that is in him, though it cost him his life. Another man, just as sincerely religious, but of different temperament, will fail of that duty, though recognizing it as a duty, and grieving to be unequal to it: but he must content the spirit that is in him—he cannot help it. He could not perform that duty for duty's *sake,* for that would not content his spirit, and the contenting of his spirit must be looked to *first*. It takes precedence of all other duties.

Y. M. Take the case of a clergyman of stainless private morals who votes for a thief for public office, on his own party's ticket, and against an honest man on the other ticket.

O. M. He has to content his spirit. He has no public morals; he has no private ones, where his party's prosperity is at stake. He will always be true to his make and training

IV

TRAINING

Young Man. You keep using that word—training. By it do you particularly mean—

Old Man. Study, instruction, lectures, sermons? That is a part of it—but not a large part. I mean *all* the outside influences. There are a million of them. From the cradle to the grave, during all his waking hours, the human being is under training. In the very first rank of his trainers stands *association.* It is his human environment which influences his mind and his feelings, furnishes him his ideals, and sets him on his road and keeps him in it. If he leave that road he will find himself shunned by the people whom he most loves and esteems, and whose approval he most values. He is a chameleon; by the law of his nature he takes the color of his place of resort. The influences about him create his preferences, his aversions, his politics, his tastes, his morals, his religion. He creates none of these things for himself. He *thinks* he does, but that is because he has not examined into the matter. You have seen Presbyterians?

Y. M. Many.

O. M. How did they happen to be Presbyterians and not Congregationalists? And why were the Congregationalists not Baptists, and the Baptists Roman Catholics, and the Roman Catholics Buddhists, and the Buddhists Quakers, and the Quakers Episcopalians, and the Episcopalians Millerites and the Millerites Hindoos, and the Hindoos Atheists, and the Atheists Spiritualists, and the Spiritualists Agnostics, and the Agnostics Methodists, and the Methodists Confucians, and the Confucians Unitarians, and the Unitarians Mohammedans, and the Mohammedans Salvation Warriors, and the Salvation Warriors Zoroastrians, and the Zoroastrians Christian Scientists, and the Christian Scientists Mormons—and so on?

Y. M. You may answer your question yourself.

O. M. That list of sects is not a record of *studies,* searchings, seekings after light; it mainly (and sarcastically) indicates what *association* can do. If you know a man's nationality you can come within a split hair of guessing the complexion of his religion: English—Protestant; American—ditto; Spaniard, Frenchman, Irishman, Italian, South American, Austrian—Roman Catholic; Russian-Greek Catholic; Turk—Mohammedan; and so on. And when you know the man's religious complexion, you know what sort of religious books he reads when he wants some more light, and what sort of books he avoids, lest by accident he get more light than he wants. In America if you know which party-collar a voter wears, you know what his associations are, and how he came by his politics, and which breed of newspaper he reads to get light, and which breed he diligently avoids, and which breed of mass-meetings he attends in order to broaden his political knowledge, and which breed of mass-meetings he doesn't attend, except to refute its doctrines with brickbats. We are always hearing of people who are around *seeking after Truth.* I have never seen a (permanent) specimen. I think he has never lived. But I have seen several entirely sincere people who *thought* they were (permanent) Seekers after Truth. They sought diligently, persistently, carefully, cautiously, profoundly, with perfect honesty and nicely adjusted judgment—until they believed that without doubt or question they had found the Truth. *That was the end of the search.* The man spent the rest of his life hunting up shingles wherewith to protect his Truth from the weather. If he was seeking after political Truth he found it in one or another of the hundred political gospels which govern men in the earth; if he was seeking after the Only True Religion he found it in one or another of the three thousand that are on the market. In any case, when he found the Truth *he sought no further;* but from that day forth, with his soldering-iron in one hand and his bludgeon in the other he tinkered its leaks and reasoned with objectors. There have been innumerable Temporary Seekers after Truth—have you ever heard of a permanent one? In the very nature of man such a

person is impossible. However, to drop back to the text—training: all training is one form or another of *outside influence,* and *associaton* is the largest part of it. A man is never anything but what his outside influences have made him. They train him downward or they train him upward—but they *train* him; they are at work upon him all the time.

Y. M. Then if he happen by the accidents of life to be evilly placed there is no help for him, according to your notions—he must train downward.

O. M. No help for him? No help for this chameleon? It is a mistake. It is in his chameleonship that his greatest good fortune lies. He has only to change his habitat—his *associations.* But the impulse to do it must come from the *outside*—he cannot originate it himself, with that purpose in view. Sometimes a very small and accidental thing can furnish him the initiatory impulse and start him on a new road, with a new ideal. The chance remark of a sweetheart, "I hear that you are a coward," may water a seed that shall sprout and bloom and flourish, and end in producing a surprising fruitage—in the fields of war. The history of man is full of such accidents. The accident of a broken leg brought a profane and ribald soldier under religious influences and furnished him a new ideal. From that accident sprang the Order of the Jesuits, and it has been shaking thrones, changing policies, and doing other tremendous work for two hundred years—and will go on. The chance reading of a book or of a paragraph in a newspaper can start a man on a new track and make him renounce his old associations and seek new ones that are *in sympathy with his new ideal:* and the result, for that man, can be an entire change of his way of life.

Y. M. Are you hinting at a scheme of procedure?

O. M. Not a new one—an old one. Old as mankind.

Y. M. What is it?

O. M. Merely the laying of traps for people. Traps baited with *Initiatory Impulses toward high ideals.* It is what the tract-distributer does. It is what the missionary does. It is what governments ought to do.

Y. M. Don't they?

O. M. In one way they do, in another way they don't. They separate the smallpox patients from the healthy people, but in dealing with crime they put the healthy into the pest-house along with the sick. That is to say, they put the beginners in with the confirmed criminals. This would be well if man were naturally inclined to good, but he isn't, and so *association* makes the beginners worse than they were when they went into captivity. It is putting a very severe punishment upon the comparatively innocent at times. They hang a man—which is a trifling punishment; this breaks the hearts of his family—which is a heavy one. They comfortably jail and feed a wife-beater, and leave his innocent wife and children to starve.

Y. M. Do you believe in the doctrine that man is equipped with an intuitive perception of good and evil?

O. M. Adam hadn't it.

Y. M. But has man acquired it since?

O. M. No. I think he has no intuitions of any kind. He gets *all* his ideas, all his impressions, from the outside. I keep repeating this, in the hope that I may so impress it upon you that you will be interested to observe and examine for yourself and see whether it is true or false.

Y. M. Where did you get your own aggravating notions?

O. M. From the *outside.* I did not invent them. They are gathered from a thousand unknown sources. Mainly *unconsciously* gathered.

Y. M. Don't you believe that God could make an inherently honest man?

O. M. Yes, I know He could. I also know that He never did make one.

Y. M. A wiser observer than you has recorded the fact that "an honest man's the noblest work of God."

O. M. He didn't record a fact, he recorded a falsity. It is windy, and sounds well, but it is not true. God makes a man with honest and dishonest *Possibilities* in him and stops there. The man's *associations* develop the possibilities—the one set or the other. The result is accordingly an honest man or a dishonest one.

Y. M. And the honest one is not entitled to—

O. M. Praise? No. How often must I tell you that? *He* is not the architect of his honesty.

Y. M. Now then, I will ask you where there is any sense in training people to lead virtuous lives. What is gained by it?

O. M. The man himself gets large advantages out of it, and that is the main thing—to *him.* He is not a peril to his neighbors, he is not a damage to them —and so *they* get an advantage out of his virtues. That is the main thing to *them*. It can make this life comparatively comfortable to the parties concerned; the *neglect* of this training can make this life a constant peril and distress to the parties concerned.

Y. M. You have said that training is everything; that training is the man *himself,* for it makes him what he is.

O. M. I said training and *another* thing. Let that other thing pass, for the moment. What were you going to say?

Y. M. We have an old servant. She has been with us twenty-two years. Her service used to be faultless, but now she has become very forgetful. We are all fond of her; we all recognize that she cannot help the infirmity which age has brought her; the rest of the family do not scold her for her remissnesses, but at times I do—I can't seem to control myself. Don't I try? I do try. Now, then, when I was ready to dress, this morning, no clean clothes had been put out. I lost my temper; I lose it easiest and quickest in the early morning. I rang; and immediately began to warn myself not to show temper, and to be careful and speak gently. I safeguarded myself most carefully. I even chose the very words I would use: "You've forgotten the clean clothes, Jane." When she appeared in the door I opened my mouth to say that phrase and out of it, moved by an instant surge of passion which I was not expecting and hadn't time to put under control, came the hot rebuke, "You've forgotten them again!" You say a man always does the thing which will best please his Interior Master. Whence came the impulse to make careful preparation to save the girl the humiliation of a rebuke? Did that come from the Master, who is always primarily concerned about *himself?*

O. M. Unquestionably. There is no other source for any impulse. *Secondarily* you made preparation to save the girl, but *primarily* its object was to save yourself, by contenting the Master.

Y. M. How do you mean?

O. M. Has any member of the family ever implored you to watch your temper and not fly out at the girl?

Y. M. Yes. My mother.

O. M. You love her?

Y. M. Oh, more than that!

O. M. You would always do anything in your power to please her?

Y. M. It is a delight to me to do anything to please her!

O. M. Why? *You would do it for pay, solely* for *profit*. What profit would you expect and certainly receive from the investment?

Y. M. Personally? None. To please *her* is enough.

O. M. It appears, then, that your object, primarily, *wasn't* to save the girl a humiliation, but to *please your mother.* It also appears that to please your mother gives *you* a strong pleasure. Is not that the profit which you get out of the investment? Isn't that the *real* profit and *first* profit?

Y. M. Oh, well? Go on.

O. M. In *all* transactions, the Interior Master looks to it that *you get the first profit.* Otherwise there is no transaction.

Y. M. Well, then, if I was so anxious to get that profit and so intent upon it, why did I throw it away by losing my temper?

O. M. In order to get *another* profit which suddenly superseded it in value.

Y. M. Where was it?

O. M. Ambushed behind your born temperament, and waiting for a chance. Your native warm temper suddenly jumped to the front, and *for the moment* its influence was more powerful than your mother's, and abolished it. In that instance you were eager to flash out a hot rebuke and enjoy it. You did enjoy it, didn't you?

Y. M. For—for a quarter of a second. Yes—I did.

O. M. Very well, it is as I have said: the thing which will give you the *most* pleasure, the most satisfaction, in any moment or *fraction* of a moment, is the thing you will always do. You must content the Master's *latest* whim, whatever it may be.

Y. M. But when the tears came into the old servant's eyes I could have cut my hand off for what I had done.

O. M. Right. You had humiliated *yourself,* you see, you had given yourself *pain.* Nothing is of *first* importance to a man except results which damage *him* or profit him—all the rest is *secondary.* Your Master was displeased with you, although you had obeyed him. He required a prompt *repentance,* you obeyed again; you *had* to—there is never any escape from his commands. He is a hard master and fickle; he changes his mind in the fraction of a second, but you must be ready to obey, and you will obey, *always.* If he requires repentance, to content him, you will always furnish it. He must be nursed, petted, coddled, and kept contented, let the terms be what they may.

Y. M. Training! Oh, what is the use of it? Didn't I, and didn't my mother try to train me up to where I would no longer fly out at that girl?

O. M. Have you never managed to keep back a scolding?

Y. M. Oh, certainly—many times.

O. M. More times this year than last?

Y. M. Yes, a good many more.

O. M. More times last year than the year before?

Y. M. Yes.

O. M. There is a large improvement, then, in the two years?

Y. M. Yes, undoubtedly.

O. M. Then your question is answered. You see there *is* use in training. Keep on. Keep faithfully on. You are doing well.

Y. M. Will my reform reach perfection?

O. M. It will. Up to *your* limit.

Y. M. My limit? What do you mean by that?

O. M. You remember that you said that I said training was *everything.* I corrected you, and said "training and *another* thing." That other thing is *temperament*—that is, the disposition you were born with. *You can't eradicate your disposition nor any rag of it*—you can only put a pressure on it and keep it down and quiet. You have a warm temper?

Y. M. Yes.

O. M. You will never get rid of it; but by watching it you can keep it down nearly all the time. *Its presence is your limit.* Your reform will never quite reach perfection, for your temper will beat you now and then, but you will come near enough. You have made valuable progress and can make more. There *is* use in training. Immense use. Presently you will reach a new stage of development, then your progress will be easier; will proceed on a simpler basis, anyway.

Y. M. Explain.

O. M. You keep back your scoldings now, to please *yourself* by pleasing your *mother;* presently the mere triumphing over your temper will delight your vanity and confer a more delicious pleasure and satisfaction upon you than even the approbation of your *mother* confers upon you now. You will then labor for yourself directly and at *first hand,* not by the roundabout way through your mother. It simplifies the matter, and it also strengthens the impulse.

Y. M. Ah, dear! But I shan't ever reach the point where I will spare the girl for *her* sake *primarily,* not mine?

O. M. Why—yes. In heaven.

Y. M. *(After a reflective pause.)* Temperament. Well, I see one must allow for temperament. It is a large factor, sure enough. My mother is thoughtful, and not hot-tempered. When I was dressed I went to her room; she was not there; I called, she answered from the bathroom. I heard the water running. I inquired. She answered, without temper, that Jane had forgotten her bath, and she was preparing it herself. I offered to ring, but she said, "No, don't do that; it would only distress her to be confronted with her lapse, and would be a rebuke; she doesn't deserve that—she is not to blame for the tricks her memory serves her." I say—has my mother an Interior Master?—and where was he?

O. M. He was there. There, and looking out for his own peace and pleasure and contentment. The girl's distress would have pained *your mother.* Otherwise the girl would have been rung up, distress and all. I know women who would have gotten a No. 1 *pleasure* out of ringing Jane up—and so they would infallibly have pushed the button and obeyed the law of their make and training, which are the servants of their Interior Masters. It is quite likely that a part of your mother's forbearance came from training. The *good* kind of training—whose best and highest function is to see to it that every time it confers a satisfaction upon its pupil a benefit shall fall at second hand upon others.

Y. M. If you were going to condense into an admonition your plan for the general betterment of the race's condition, how would you word it?

Admonition

O. M. Diligently train your ideals *upward* and *still upward* toward a summit where you will find your chiefest pleasure in conduct which, while contenting you, will be sure to confer benefits upon your neighbor and the community.

Y. M. Is that a new gospel?

O. M. No.

Y. M. It has been taught before?

O. M. For ten thousand years.
Y. M. By whom?
O. M. All the great religions all the great gospels.
Y. M. Then there is nothing new about it?
O. M. Oh yes, there is. It is candidly stated, this time. That has not been done before.
Y. M. How do you mean?
O. M. Haven't I put *you* FIRST, and your neighbor and the community *afterward?*
Y. M. Well, yes, that is a difference, it is true.
O. M. The difference between straight speaking and crooked; the difference between frankness and shuffling.
Y. M. Explain.
O. M. The others offer you a hundred bribes to be good, thus conceding that the Master inside of you must be conciliated and contented first, and that you will do nothing at *first hand* but for his sake; then they turn square around and require you to do good for *others'* sake *chiefly;* and to do your duty for duty's *sake,* chiefly; and to do acts of *self-sacrifice.* Thus at the outset we all stand upon the same ground —recognition of the supreme and absolute Monarch that resides in man, and we all grovel before him and appeal to him; then those others dodge and shuffle, and face around and unfrankly and inconsistently and illogically change the form of their appeal and direct its persuasions to man's *second-place* powers and to powers which have *no existence* in him, thus advancing them to *first* place; whereas in my Admonition I stick logically and consistently to the original position: I place the Interior Master's requirements *first,* and keep them there.
Y. M. If we grant, for the sake of argument, that your scheme and the other schemes aim at and produce the same result—*right living*—has yours an advantage over the others?
O. M. One, yes—a large one. It has no concealments, no deceptions. When a man leads a right and valuable life under it he is not deceived as to the *real* chief motive which impels him to it—in those other cases he is.
Y. M. Is that an advantage? Is it an advantage to live a lofty life for a mean reason? In the other cases he lives the lofty life under the *impression* that he is living it for a lofty reason. Is not that an advantage?
O. M. Perhaps so. The same advantage he might get out of thinking himself a duke, and living a duke's life and parading in ducal fuss and feathers, when he wasn't a duke at all, and could find it out if he would only examine the herald's records.
Y. M. But anyway, he is obliged to do a duke's part; he puts his hand in his pocket and does his benevolences on as big a scale as he can stand, and that benefits the community.

O. M. He could do that without being a duke.
Y. M. But would he?
O. M. Don't you see where you are arriving?
Y. M. Where?
O. M. At the standpoint of the other schemes: That it is good morals to let an ignorant duke do showy benevolences for his pride's sake, a pretty low motive, and go on doing them unwarned, lest if he were made acquainted with the actual motive which prompted them he might shut up his purse and cease to be good?
Y. M. But isn't it best to leave him in ignorance, as long as he *thinks* he is doing good for others' sake?
O. M. Perhaps so. It is the position of the other schemes. They think humbug is good enough morals when the dividend on it is good deeds and handsome conduct.
Y. M. It is my opinion that under your scheme of a man's doing a good deed for his *own* sake first-off, instead of first for the *good deed's* sake, no man would ever do one.
O. M. Have you committed a benevolence lately?
Y. M. Yes. This morning.
O. M. Give the particulars.
Y. M. The cabin of the old negro woman who used to nurse me when I was a child and who saved my life once at the risk of her own, was burned last night, and she came mourning this morning, and pleading for money to build another one.
O. M. You furnished it?
Y. M. Certainly.
O. M. You were glad you had the money?
Y. M. Money? I hadn't. I sold my horse.
O. M. You were glad you had the horse?
Y. M. Of course I was; for if I hadn't had the horse I should have been incapable, and my *mother* would have captured the chance to set old Sally up.
O. M. You were cordially glad you were not caught out and incapable?
Y. M. Oh, I just was!
O. M. Now, then—
Y. M. Stop where you are! I know your whole catalogue of questions, and I could answer every one of them without your wasting the time to ask them; but I will summarize the whole thing in a single remark: I did the charity knowing it was because the act would give *me* a splendid pleasure, and because old Sally's moving gratitude and delight would give *me* another one; and because the reflection that she would be happy now and out of her trouble would fill *me* full of happiness. I did the whole thing with my eyes open and recognizing and realizing

that I was looking out for *my* share of the profits *first*. Now then, I have confessed. Go on.

O. M. I haven't anything to offer; you have covered the whole ground. Could you have been any *more* strongly moved to help Sally out of her trouble —could you have done the deed any more eagerly—if you had been under the delusion that you were doing it for *her* sake and profit only?

Y. M. No! Nothing in the world could have made the impulse which moved me more powerful, more masterful, more thoroughly irresistible. I played the limit!

O. M. Very well. You begin to suspect—and I claim to *know* that when a man is a shade *more strongly moved* to do *one* of two things or of two dozen things than he is to do any one of the *others,* he will infallibly do that *one* thing, be it good or be it evil; and if it be good, not all the beguilements of all the casuistries can increase the strength of the impulse by a single shade or add a shade to the comfort and contentment he will get out of the act.

Y. M. Then you believe that such tendency toward doing good as is in men's hearts would not be diminished by the removal of the delusion that good deeds are done primarily for the sake of No. 2 instead of for the sake of No. 1?

O. M. That is what I fully believe.

Y. M. Doesn't it somehow seem to take from the dignity of the deed?

O. M. If there is dignity in falsity, it does. It removes that.

Y. M. What is left for the moralist to do?

O. M. Teach unreservedly what he already teaches with one side of his mouth and takes back with the other: Do right *for your own sake,* and be happy in knowing that your *neighbor* will certainly share in the benefits resulting.

Y. M. Repeat your Admonition.

O. M. *Diligently train your ideals upward and still upward toward a summit where you will find your chiefest pleasure in conduct which, while contenting you, will be sure to confer benefits upon your neighbor and the community.*

Y. M. One's *every* act proceeds from *exterior influences*, you think?

O. M. Yes.

Y. M. If I conclude to rob a person, I am not the *originator* of the idea, but it comes in from the *outside?* I see him handling money—for instance—and *that* moves me to the crime?

O. M. That, by itself? Oh, certainly not. It is merely the *latest* outside influence of a procession of preparatory influences stretching back over a period of years. No *single* outside influence can make a man do a thing which is at war with his training. The most it can do is to start his mind on a new tract and open it to the reception of *new* influences —as

in the case of Ignatius Loyola. In time these influences can train him to a point where it will be consonant with his new character to yield to the *final* influence and do that thing. I will put the case in a form which will make my theory clear to you, I think. Here are two ingots of virgin gold. They shall represent a couple of characters which have been refined and perfected in the virtues by years of diligent right training. Suppose you wanted to break down these strong and well-compacted characters—what influence would you bring to bear upon the ingots?

Y. M. Work it out yourself. Proceed.

O. M. Suppose I turn upon one of them a steamjet during a long succession of hours. Will there be a result?

Y. M. None that I know of.

O. M. Why?

Y. M. A steam-jet cannot break down such a substance.

O. M. Very well. The steam is an *outside influence*, but it is ineffective because the gold *takes no interest in it.* The ingot remains as it was. Suppose we add to the steam some quicksilver in a vaporized condition, and turn the jet upon the ingot, will there be an instantaneous result?

Y. M. No.

O. M. The *quicksilver* is an outside influence which gold (by its peculiar nature—say *temperament, disposition) cannot be indifferent to*. It stirs the interest of the gold, although we do not perceive it; but a *single* application of the influence works no damage. Let us continue the application in a steady stream, and call each minute a year. By the end of ten or twenty minutes—ten or twenty years—the little ingot is sodden with quicksilver, its virtues are gone, its character is degraded. At last it is ready to yield to a temptation which it would have taken no notice of, ten or twenty years ago. We will apply that temptation in the form of a pressure of my finger. You note the result?

Y. M. Yes; the ingot has crumbled to sand. I understand, now. It is not the *single* outside influence that does the work, but only the *last* one of a long and disintegrating accumulation of them. I see, now, how my *single* impulse to rob the man is not the one that makes me do it, but only the *last* one of a preparatory series. You might illustrate it with a parable.

A Parable

O. M. I will. There was once a pair of New England boys—twins. They were alike in good dispositions, fleckless morals, and personal appearance. They were the models of the Sunday-school. At fifteen George had an opportunity to go as cabin-boy in a whale-ship, and sailed away for the Pacific. Henry remained at home in the village. At eighteen George was a sailor before the mast, and Henry was teacher of the advanced Bible class. At twenty-two George, through fighting-habits and drinking-habits acquired at sea and in the sailor boarding-houses of the European and Oriental ports, was a common rough in Hong-Kong, and out of a job; and Henry was superintendent of the Sunday-school. At twenty-six George was a wanderer, a tramp, and Henry was pastor of the village church. Then George came home, and was Henry's guest. One evening a man passed by and turned down the lane, and Henry said, with a pathetic smile, "Without intending me a discomfort, that man is always keeping me reminded of my pinching poverty, for he carries heaps of money about him, and goes by here every evening of his life." That *outside influence*—that remark—was enough for George, but *it* was not the one that made him ambush the man and rob him, it merely represented the eleven years' accumulation of such influences, and gave birth to the act for which their long gestation had made preparation. It had never entered the head of Henry to rob the man—his ingot had been subjected to clean steam only; but George's had been subjected to vaporized quicksilver.

V

MORE ABOUT THE MACHINE

Note.—When Mrs. W. asks how can a millionaire give a single dollar to colleges and museums while one human being is destitute of bread, she has answered her question herself. Her feeling for the poor shows that she has a standard of benevolence; therefore she has conceded the millionaire's privilege of having a standard; since she evidently requires him to adopt her standard, she is by that act requiring herself to adopt his. The human being always looks down when he is examining another person's standard; he never finds one that he has to examine by looking up.

The Man-Machine Again

Young Man. You really think man is a mere machine?

Old Man. I do.

Y. M. And that his mind works automatically and is independent of his control—carries on thought on its own hook?

O. M. Yes. It is diligently at work, unceasingly at work, during every waking moment. Have you never tossed about all night, imploring, beseeching, commanding your mind to stop work and let you go to sleep?—you who perhaps imagine that your mind is your servant and must obey your orders, think what you tell it to think, and stop when you tell it to stop. When it chooses to work, there is no way to keep it still for an instant. The brightest man would not be able to supply it with subjects if he had to hunt them up. If it needed the man's help it would wait for him to give it work when he wakes in the morning.

Y. M. Maybe it does.

O. M. No, it begins right away, before the man gets wide enough awake to give it a suggestion. He may go to sleep saying, "The moment I wake I will think upon such and such a subject," but he will fail. His mind will be too quick for him; by the time he has become nearly enough awake to be half conscious, he will find that it is already at work upon another subject. Make the experiment and see.

Y. M. At any rate, he can make it stick to a subject if he wants to.

O. M. Not if it finds another that suits it better. As a rule it will listen to neither a dull speaker nor a bright one. It refuses all persuasion. The dull speaker wearies it and sends it far away in idle dreams; the bright speaker throws out stimulating ideas which it goes chasing after and is at once unconscious of him and his talk. You cannot keep your mind from wandering, if it wants to; it is master, not you.

After an Interval of Days

O. M. Now, dreams—but we will examine that later. Meantime, did you try commanding your mind to wait for orders from you, and not do any thinking on its own hook?

Y. M. Yes, I commanded it to stand ready to take orders when I should wake in the morning.

O. M. Did it obey?

Y. M. No. It went to thinking of something of its own initiation, without waiting for me. Also—as you suggested—at night I appointed a theme for it to begin on in the morning, and commanded it to begin on that one and no other.

O. M. Did it obey?

Y. M. No.

O. M. How many times did you try the experiment?

Y. M. Ten.

O. M. How many successes did you score?

Y. M. Not one.

O. M. It is as I have said: the mind is independent of the man. He has no control over it; it does as it pleases. It will take up a subject in spite of him; it will stick to it in spite of him; it will throw it aside in spite of him. It is entirely independent of him.

Y. M. Go on. Illustrate.

O. M. Do you know chess?

Y. M. I learned it a week ago.

O. M. Did your mind go on playing the game all night that first night?

Y. M. Don't mention it!

O. M. It was eagerly, unsatisfiably interested; it rioted in the combinations; you implored it to drop the game and let you get some sleep?

Y. M. Yes. It wouldn't listen; it played right along. It wore me out and I got up haggard and wretched in the morning.

O. M. At some time or other you have been captivated by a ridiculous rhyme-jingle?

Y. M. Indeed, yes!

"I saw Esau kissing Kate,
And she saw I saw Esau;
I saw Esau, he saw Kate,
And she saw—"

And so on. My mind went mad with joy over it. It repeated it all day and all night for a week in spite of all I could do to stop it, and it seemed to me that I must surely go crazy.

O. M. And the new popular song?

Y. M. Oh yes! "In the Swee-eet By and By"; etc. Yes, the new popular song with the taking melody sings through one's head day and night, asleep and awake, till one is a wreck. There is no getting the mind to let it alone.

O. M. Yes, asleep as well as awake. The mind is quite independent. It is master. You have nothing to do with it. It is so apart from you that it can conduct its affairs, sing its songs, play its chess, weave its complex and ingeniously constructed dreams, while you sleep. It has no use for your help, no use for your guidance, and never uses either, whether you be asleep or awake. You have imagined that you could originate a thought in your mind, and you have sincerely believed you could do it.

Y. M. Yes, I have had that idea.

O. M. Yet you can't originate a dream-thought for it to work out, and get it accepted?

Y. M. No.

O. M. And you can't dictate its procedure after it has originated a dream-thought for itself?

Y. M. No. No one can do it. Do you think the waking mind and the dream mind are the same machine?

O. M. There is argument for it. We have wild and fantastic day-thoughts? Things that are dreamlike?

Y. M. Yes—like Mr. Wells's man who invented a drug that made him invisible; and like the Arabian tales of the Thousand Nights.

O. M. And there are dreams that are rational, simple, consistent, and unfantastic?

Y. M. Yes. I have dreams that are like that. Dreams that are just like real life; dreams in which there are several persons with distinctly differentiated characters—inventions of my mind and yet strangers to me: a vulgar person; a refined one; a wise person; a fool; a cruel person; a kind and compassionate one; a quarrelsome person; a peacemaker; old persons and young; beautiful girls and homely ones. They talk in character, each preserves his own characteristics. There are vivid fights, vivid and biting insults, vivid love-passages; there are tragedies and comedies, there are griefs that go to one's heart, there are sayings and doings that make you laugh: indeed, the whole thing is exactly like real life.

O. M. Your dreaming mind originates the scheme, consistently and artistically develops it, and carries the little drama creditably through all without help or suggestion from you?

Y. M. Yes.

O. M. It is argument that it could do the like awake without help or suggestion from you—and I think it does. It is argument that it is the same old mind in both cases, and never needs your help. I think the mind is purely a machine, a thoroughly independent machine, an automatic machine. Have you tried the other experiment which I suggested to you?

Y. M. Which one?

O. M. The one which was to determine how much influence you have over your mind—if any.

Y. M. Yes, and got more or less entertainment out of it. I did as you ordered: I placed two texts before my eyes—one a dull one and barren of interest, the other one full of interest, inflamed with it, white-hot with it. I commanded my mind to busy itself solely with the dull one.

O. M. Did it obey?

Y. M. Well, no, it didn't. It busied itself with the other one.

O. M. Did you try hard to make it obey?

Y. M. Yes, I did my honest best.

O. M. What was the text which it refused to be interested in or think about?

Y. M. It was this question: If A owes B a dollar and a half, and B owes C two and three-quarters, and C owes A thirty-five cents, and D and A together owe E and B three-sixteenths of—of—I don't remember the rest, now, but anyway it was wholly uninteresting, and I could not force my mind to stick to it even half a minute at a time; it kept flying off to the other text.

O. M. What was the other text?

Y. M. It is no matter about that.

O. M. But what was it?

Y. M. A photograph.

O. M. Your own?

Y. M. No. It was hers.

O. M. You really made an honest good test. Did you make a second trial?

Y. M. Yes. I commanded my mind to interest itself in the morning paper's report of the pork-market, and at the same time I reminded it of an experience of mine of sixteen years ago. It refused to consider the pork and gave its whole blazing interest to that ancient incident.

O. M. What was the incident?

Y. M. An armed desperado slapped my face in the presence of twenty spectators. It makes me wild and murderous every time I think of it.

O. M. Good tests, both; very good tests. Did you try my other suggestion?

Y. M. The one which was to prove to me that if I would leave my mind to its own devices it would find things to think about without any of my help, and thus convince me that it was a machine, an automatic machine, set in motion by exterior influences, and as independent of me as it could be if it were in some one else's skull? Is that the one?

O. M. Yes.

Y. M. I tried it. I was shaving. I had slept well, and my mind was very lively, even gay and frisky. It was reveling in a fantastic and joyful episode of my remote boyhood which had suddenly flashed up in my memory—moved to this by the spectacle of a yellow cat picking its way carefully along the top of the garden wall. The color of this cat brought the bygone cat before me, and I saw her walking along the side-step of the pulpit; saw her walk on to a large sheet of sticky fly-paper and get all her feet involved; saw her struggle and fall down, helpless and dissatisfied, more and more urgent, more and more unreconciled, more and more mutely profane; saw the silent congregation quivering like jelly, and the tears running down their faces. I saw it all. The sight of the tears whisked my mind to a far distant and a sadder scene—in Terra del Fuego—and with Darwin's eyes I saw a naked great savage hurl his little boy against the rocks for a trifling fault; saw the poor mother gather up her dying child and hug it to her breast and weep, uttering no word. Did my mind stop to mourn with that nude black sister of mine? No—it was far away from that scene in an instant, and was busying itself with an ever-recurring and disagreeable dream of mine. In this dream I always find myself, stripped to my shirt, cringing and dodging about in the midst of a great drawing-room throng of finely dressed ladies and gentlemen, and wondering how I got there. And so on and so on, picture after picture, incident after incident, a drifting panorama of ever-changing, ever-dissolving views manufactured by my mind without any help from me—why, it would take me two hours to merely name the multitude of things my mind tallied off and photographed in fifteen minutes, let alone describe them to you.

O. M. A man's mind, left free, has no use for his help. But there is one way whereby he can get its help when he desires it.

Y. M. What is that way?

O. M. When your mind is racing along from subject to subject and strikes an inspiring one, open your mouth and begin talking upon that matter—or take your pen and use that. It will interest your mind and concentrate it, and it will pursue the subject with satisfaction. It will take full charge, and furnish the words itself.

Y. M. But don't I tell it what to say?

O. M. There are certainly occasions when you haven't time. The words leap out before you know what is coming.

Y. M. For instance?

O. M. Well, take a "flash of wit"—repartee. Flash is the right word. It is out instantly. There is no time to arrange the words. There is no thinking, no reflecting. Where there is a wit-mechanism it is automatic in its action and needs no help. Where the wit-mechanism is lacking, no amount of study and reflection can manufacture the product.

Y. M. You really think a man originates nothing, creates nothing.

The Thinking-Process

O. M. I do. Men perceive, and their brain-machines automatically combine the things perceived. That is all.

Y. M. The steam-engine?

O. M. It takes fifty men a hundred years to invent it. One meaning of invent is discover. I use the word in that sense. Little by little they discover and apply the multitude of details that go to make the perfect engine. Watt noticed that confined steam was strong enough to lift the lid of the teapot. He didn't create the idea, he merely discovered the fact; the cat had noticed it a hundred times. From the teapot he evolved the cylinder—from the displaced lid he evolved the piston-rod. To attach something to the piston-rod to be moved by it, was a simple matter—crank and wheel. And so there was a working engine.[1]

One by one, improvements were discovered by men who used their eyes, not their creating powers—for they hadn't any—and now, after a hundred years the patient contributions of fifty or a hundred observers stand compacted in the wonderful machine which drives the ocean liner.

Y. M. A Shakespearian play?

O. M. The process is the same. The first actor was a savage. He reproduced in his theatrical war-dances, scalp-dances, and so on, incidents which he had seen in real life. A more advanced civilization produced more incidents, more episodes; the actor and the story-teller borrowed them. And so the drama grew, little by little, stage by stage. It is made up of the facts of life, not creations. It took centuries to develop the

[1]The Marquess of Worcester had done all of this more than a century earlier.

Greek drama. It borrowed from preceding ages; it lent to the ages that came after. Men observe and combine, that is all. So does a rat.

Y. M. How?

O. M. He observes a smell, he infers a cheese, he seeks and finds. The astronomer observes this and that; adds his this and that to the this-and-thats of a hundred predecessors, infers an invisible planet, seeks it and finds it. The rat gets into a trap; gets out with trouble; infers that cheese in traps lacks value, and meddles with that trap no more. The astronomer is very proud of his achievement, the rat is proud of his. Yet both are machines; they have done machine work, they have originated nothing, they have no right to be vain; the whole credit belongs to their Maker. They are entitled to no honors, no praises, no monuments when they die, no remembrance. One is a complex and elaborate machine, the other a simple and limited machine, but they are alike in principle, function, and process, and neither of them works otherwise than automatically, and neither of them may righteously claim a *personal* superiority or a personal dignity above the other.

Y. M. In earned personal dignity, then, and in personal merit for what he does, it follows of necessity that he is on the same level as a rat?

O. M. His brother the rat; yes, that is how it seems to me. Neither of them being entitled to any personal merit for what he does, it follows of necessity that neither of them has a right to arrogate to himself (personally created) superiorities over his brother.

Y. M. Are you determined to go on believing in these insanities? Would you go on believing in them in the face of able arguments backed by collated facts and instances?

O. M. I have been a humble, earnest, and sincere Truth-Seeker.

Y. M. Very well?

O. M. The humble, earnest, and sincere Truth-Seeker is always convertible by such means.

Y. M. I am thankful to God to hear you say this, for now I know that your conversion—

O. M. Wait. You misunderstand. I said I have *been* a Truth-Seeker.

Y. M. Well?

O. M. I am not that now. Have you forgotten? I told you that there are none but temporary Truth-Seekers; that a permanent one is a human impossibility; that as soon as the Seeker finds what he is thoroughly convinced is the Truth, he seeks no further, but gives the rest of his days to hunting junk to patch it and caulk it and prop it with, and make it weather-proof and keep it from caving in on him. Hence the Presbyterian remains a Presbyterian, the Mohammedan a Mohammedan, the Spiritualist a Spiritualist, the Democrat a Democrat, the Republican a Republican, the Monarchist a Monarchist; and if a humble, earnest, and sincere Seeker after Truth should find it in the proposition that the moon is made of green cheese nothing could ever budge him from that position; for he is nothing but an automatic machine, and must obey the laws of his construction.

Y. M. And so—

O. M. Having found the Truth; perceiving that beyond question man has but one moving impulse—the contenting of his own spirit—and is merely a machine and entitled to no personal merit for anything he does, it is not humanly possible for me to seek further. The rest of my days will be spent in patching and painting and puttying and caulking my priceless possession and in looking the other way when an imploring argument or a damaging fact approaches.

VI

INSTINCT AND THOUGHT

Young Man. It is odious. Those drunken theories of yours, advanced a while ago—concerning the rat and all that—strip Man bare of all his dignities, grandeurs, sublimities.

Old Man. He hasn't any to strip—they are shams, stolen clothes. He claims credits which belong solely to his Maker.

Y. M. But you have no right to put him on a level with a rat.

O. M. I don't—morally. That would not be fair to the rat. The rat is well above him, there.

Y. M. Are you joking?

O. M. No, I am not.

Y. M. Then what do you mean?

O. M. That comes under the head of the Moral Sense. It is a large question. Let us finish with what we are about now, before we take it up.

Y. M. Very well. You have seemed to concede that you place Man and the rat on *a* level. What is it? The intellectual?

O. M. In form—not in degree.

Y. M. Explain.

O. M. I think that the rat's mind and the man's mind are the same machine, but of unequal capacities—like yours and Edison's; like the African pygmy's and Homer's; like the Bushman's and Bismarck's.

Y. M. How are you going to make that out, when the lower animals have no mental quality but instinct, while man possesses reason?

O. M. What is instinct?

Y. M. It is merely unthinking and mechanical exercise of inherited habit.

O. M. What originated the habit?

Y. M. The first animal started it, its descendants have inherited it.

O. M. How did the first one come to start it?

Y. M. I don't know; but it didn't *think* it out.

O. M. How do you know it didn't?

Y. M. Well—I have a right to suppose it didn't, anyway.

O. M. I don't believe you have. What is thought?

Y. M. I know what you call it: the mechanical and automatic putting together of impressions received from outside, and drawing an inference from them.

O. M. Very good. Now my idea of the meaningless term "instinct" is, that it is merely *petrified thought;* solidified and made inanimate by habit; thought which was once alive and awake, but is become unconscious—walks in its sleep, so to speak.

Y. M. Illustrate it.

O. M. Take a herd of cows, feeding in a pasture. Their heads are all turned in one direction. They do that instinctively; they gain nothing by it, they have no reason for it, they don't know why they do it. It is an inherited habit which was originally thought—that is to say, observation of an exterior fact, and a valuable inference drawn from that observation and confirmed by experience. The original wild ox noticed that with the wind in his favor he could smell his enemy in time to escape; then he inferred that it was worth while to keep his nose to the wind. That is the process which man calls reasoning. Man's thought-machine works just like the 'other animals', but it is a better one and more Edisonian. Man, in the ox's place, would go further, reason wider: he would face part of the herd the other way and protect both front and rear.

Y. M. Did you say the term instinct is meaningless?

O. M. I think it is a bastard word. I think it confuses us; for as a rule it applies itself to habits and impulses which had a far-off origin in thought, and now and then breaks the rule and applies itself to habits which can hardly claim a thought-origin.

Y. M. Give an instance.

O. M. Well, in putting on trousers a man always inserts the same old leg first—never the other one. There is no advantage in that, and no sense in it. All men do it, yet no man thought it out and adopted it of set purpose, I imagine. But it is a habit which is transmitted, no doubt, and will continue to be transmitted.

Y. M. Can you prove that the habit exists?

O. M. You can prove it, if you doubt. If you will take a man to a clothing-store and watch him try on a dozen pairs of trousers, you will see.

Y. M. The cow illustration is not—

O. M. Sufficient to show that a dumb animal's mental machine is just

the same as a man's and its reasoning processes the same? I will illustrate further. If you should hand Mr. Edison a box which you caused to fly open by some concealed device he would infer a spring, and would hunt for it and find it. Now an uncle of mine had an old horse who used to get into the closed lot where the corncrib was and dishonestly take the corn. I got the punishment myself, as it was supposed that I had heedlessly failed to insert the wooden pin which kept the gate closed. These persistent punishments fatigued me; they also caused me to infer the existence of a culprit, somewhere; so I hid myself and watched the gate. Presently the horse came and pulled the pin out with his teeth and went in. Nobody taught him that; he had observed—then thought it out for himself. His process did not differ from Edison's; he put this and that together and drew an inference—and the peg, too; but I made him sweat for it.

Y. M. It has something of the seeming of thought about it. Still it is not very elaborate. Enlarge.

O. M. Suppose that Edison has been enjoying some one's hospitalities. He comes again by and by, and the house is vacant. He infers that his host has moved. A while afterward, in another town, he sees the man enter a house; he infers that that is the new home, and follows to inquire. Here, now, is the experience of a gull, as related by a naturalist. The scene is a Scotch fishing village where the gulls were kindly treated. This particular gull visited a cottage; was fed; came next day and was fed again; came into the house, next time, and ate with the family; kept on doing this almost daily, thereafter. But, once the gull was away on a journey for a few days, and when it returned the house was vacant. Its friends had removed to a village three miles distant. Several months later it saw the head of the family on the street there, followed him home, entered the house without excuse or apology, and became a daily guest again. Gulls do not rank high mentally, but this one had memory and the reasoning faculty, you see, and applied them Edisonially.

Y. M. Yet it was not an Edison and couldn't be developed into one.

O. M. Perhaps not. Could you?

Y. M. That is neither here nor there. Go on.

O. M. If Edison were in trouble and a stranger helped him out of it and next day he got into the same difficulty again, he would infer the wise thing to do in case he knew the stranger's address. Here is a case of a bird and a stranger as related by a naturalist. An Englishman saw a bird flying around about his dog's head, down in the grounds, and uttering cries of distress. He went there to see about it. The dog had a young bird in his mouth—unhurt. The gentleman rescued it and put it on a bush and brought the dog away. Early the next morning the mother bird came for the gentleman, who was sitting on his veranda, and by its maneuvers persuaded him to follow it to a distant part of the grounds—flying a little way in front of him and waiting for him to catch up, and so on; and keeping to the winding path, too, instead of flying the near way across lots. The distance covered was four hundred yards. The same dog was the culprit; he had the young bird again, and once more he had to give it up. Now the mother bird had reasoned it all out: since the stranger had helped her once, she inferred that he would do it again; she knew where to find him, and she went upon her errand with confidence. Her mental processes were what Edison's would have been. She put this and that together—and that is all that thought *is*—and out of them built her logical arrangement of inferences. Edison couldn't have done it any better himself.

Y. M. Do you believe that many of the dumb animals can think?

O. M. Yes—the elephant, the monkey, the horse, the dog, the parrot, the macaw, the mocking-bird, and many others. The elephant whose mate fell into a pit, and who dumped dirt and rubbish into the pit till the bottom was raised high enough to enable the captive to step out, was equipped with the reasoning quality. I conceive that all animals that can learn things through teaching and drilling have to know how to observe, and put this and that together and draw an inference—the process of thinking. Could you teach an idiot the manual of arms, and to advance, retreat, and go through complex field maneuvers at the word of command?

Y. M. Not if he were a thorough idiot.

O. M. Well, canary-birds can learn all that; dogs and elephants learn all sorts of wonderful things. They must surely be able to notice, and to put things together, and say to themselves, " I get the idea, now: when I do so and so, as per order, I am praised and fed; when I do differently I am punished." Fleas can be taught nearly anything that a Congressman can.

Y. M. Granting, then, that dumb animals are able to think upon a low plane, is there any that can think upon a high one? Is there one that is well up toward man?

O. M. Yes. As a thinker and planner the ant is the equal of any savage race of men; as a self-educated specialist in several arts she is the superior of any savage race of men; and in one or two high mental qualities she is above the reach of any man, savage or civilized!

Y. M. Oh, come! you are abolishing the intellectual frontier which separates man and beast.

O. M. I beg your pardon. One cannot abolish what does not exist.

Y. M. You are not in earnest, I hope. You cannot mean to seriously say there is no such frontier.

O. M. I do say it seriously. The instances of the horse, the gull, the mother bird, and the elephant show that those creatures put their this's

and thats together just as Edison would have done it and drew the same inferences that he would have drawn. Their mental machinery was just like his, also its manner of working. Their equipment was as inferior to his in elaboration as a Waterbury is inferior to the Strasburg clock, but that is the only difference—there is no frontier.

Y. M. It looks exasperatingly true; and is distinctly offensive. It elevates the dumb beasts to—to—

O. M. Let us drop that lying phrase, and call them the Unrevealed Creatures; so far as we can know, there is no such thing as a dumb beast.

Y. M. On what grounds do you make that assertion?

O. M. On quite simple ones. "Dumb" beast suggests an animal that has no thought-machinery, no understanding, no speech, no way of communicating what is in its mind. We know that a hen *has* speech. We cannot understand everything she says, but we easily learn two or three of her phrases. We know when she is saying, "I have laid an egg"; we know when she is saying to the chicks, "Run here, dears, I've found a worm"; we know what she is saying when she voices a warning: "Quick! hurry! gather yourselves under mamma, there's a hawk coming!" We understand the cat when she stretches herself out, purring with affection and contentment and lifts up a soft voice and says, "Come, kitties, supper's ready"; we understand her when she goes mourning about and says, " Where can they be? They are lost. Won't you help me hunt for them?" and we understand the disreputable Tom when he challenges at midnight from his shed, "You come over here, you product of immoral commerce, and I'll make your fur fly!" We understand a few of a dog's phrases and we learn to understand a few of the remarks and gestures of any bird or other animal that we domesticate and observe. The clearness and exactness of the few of the hen's speeches which we understand is argument that she can communicate to her kind a hundred things which we cannot comprehend—in a word, that she can converse. And this argument is also applicable in the case of others of the great army of the Unrevealed. It is just like man's vanity and impertinence to call an animal dumb because it is dumb to his dull perceptions. Now as to the ant—

Y. M. Yes, go back to the ant, the creature that. —as you seem to think—sweeps away the last vestige of an intellectual frontier between man and the Unrevealed.

O. M. That is what she surely does. In all his history the aboriginal Australian never thought out a house for himself and built it. The ant is an amazing architect. She is a wee little creature, but she builds a strong and enduring house eight feet high—a house which is as large in proportion to her size as is the largest capitol or cathedral in the world compared to man's size. No savage race has produced architects who could approach the ant in genius or culture. No civilized race has produced architects who could plan a house better for the uses proposed than can hers. Her house contains a throne-room; nurseries for her young; granaries; apartments for her soldiers, her workers, etc.; and they and the multifarious halls and corridors which communicate with them are arranged and distributed with an educated and experienced eye for convenience and adaptability.

Y. M. That could be mere instinct.

O. M. It would elevate the savage if he had it. But let us look further before we decide. The ant has soldiers—battalions, regiments, armies; and they have their appointed captains and generals, who lead them to battle.

Y. M. That could be instinct, too.

O. M. We will look still further. The ant has a system of government; it is well planned, elaborate, and is well carried on.

Y. M. Instinct again.

O. M. She has crowds of slaves, and is a hard and unjust employer of forced labor.

Y. M. Instinct.

O. M. She has cows, and milks them.

Y. M. Instinct, of course.

O. M. In Texas she lays out a farm twelve feet square, plants it, weeds it, cultivates it, gathers the crop and stores it away.

Y. M. Instinct, all the same.

O. M. The ant discriminates between friend and stranger. Sir John Lubbock took ants from two different nests, made them drunk with whisky and laid them, unconscious, by one of the nests, near some water. Ants from the nest came and examined and discussed these disgraced creatures, then carried the friends home and threw the strangers overboard. Sir John repeated the experiment a number of times. For a time the sober ants did as they had done at first—carried their friends home and threw the strangers overboard. But finally they lost patience, seeing that their reformatory efforts went for nothing, and threw both friends and strangers overboard. Come—is this instinct, or is it thoughtful and intelligent discussion of a thing new—absolutely new—to their experience; with a verdict arrived at, sentence passed, and judgment executed? Is it instinct?—thought petrified by ages of habit—or isn't it brand-new thought, inspired by the new occasion, the new circumstances?

Y. M. I have to concede it. It was not a result of habit; it has all the look of reflection, thought, putting this and that together, as you phrase it. I believe it was thought.

O. M. I will give you another instance of thought. Franklin had a cup of sugar on a table in his room. The ants got at it. He tried several

preventives; the ants rose superior to them. Finally he contrived one which shut off access—probably set the table's legs in pans of water, or drew a circle of tar around the cup, I don't remember. At any rate, he watched to see what they would do. They tried various schemes—failures, every one. The ants were badly puzzled. Finally they held a consultation, discussed the problem, arrived at a decision—and this time they beat that great philosopher. They formed in procession, crossed the floor, climbed the wall, marched across the ceiling to a point just over the cup, then one by one they let go and fell down into it! Was that instinct—thought petrified by ages of inherited habit?

Y. M. No, I don't believe it was. I believe it was a newly reasoned scheme to meet a new emergency.

O. M. Very well. You have conceded the reasoning power in two instances. I come now to a mental detail wherein the ant is a long way the superior of any human being. Sir John Lubbock proved by many experiments that an ant knows a stranger ant of her own species in a moment, even when the stranger is disguised—with paint. Also he proved that an ant knows every individual in her hive of five hundred thousand souls. Also, after a year's absence of one of the five hundred thousand she will straightway recognize the returned absentee and grace the recognition with an affectionate welcome. How are these recognitions made? Not by color, for painted ants were recognized. Not by smell, for ants that had been dipped in chloroform were recognized. Not by speech and not by antennae signs nor contacts, for the drunken and motionless ants were recognized and the friend discriminated from the stranger. The ants were all of the same species, therefore the friends had to be recognized by form and feature—friends who formed part of a hive of five hundred thousand! Has any man a memory for form and feature approaching that?

Y. M. Certainly not.

O. M. Franklin's ants and Lubbock's ants show fine capacities of putting this and that together in new and untried emergencies and deducting smart conclusions from the combinations— a man's mental process exactly. With memory to help, man preserves his observations and reasonings, reflects upon them, adds to them, recombines, and so proceeds, stage by stage, to far results—from the teakettle to the ocean greyhound's complex engine; from personal labor to slave labor; from wigwam to palace; from the capricious chase to agriculture and stored food; from nomadic life to stable government and concentrated authority; from incoherent hordes to massed armies. The ant has observation, the reasoning faculty, and the preserving adjunct of a prodigious memory; she has duplicated man's development and the essential features of his civilization, and you call it all instinct!

Y. M. Perhaps I lacked the reasoning faculty myself.

O. M. Well, don't tell anybody, and don't do it again.

Y. M. We have come a good way. As a result—as I understand it—I am required to concede that there is absolutely no intellectual frontier separating Man and the Unrevealed Creatures?

O. M. That is what you are required to concede. There is no such frontier—there is no way to get around that. Man has a finer and more capable machine in him than those others, but it is the same machine and works in the same way. And neither he nor those others can command the machine—it is strictly automatic, independent of control, works when it pleases, and when it doesn't please, it can't be forced.

Y. M. Then man and the other animals are all alike, as to mental machinery, and there isn't any difference of any stupendous magnitude between them, except in quality, not in kind.

O. M. That is about the state of it—intellectuality. There are pronounced limitations on both sides. We can't learn to understand much of their language, but the dog, the elephant, etc., learn to understand a very great deal of ours. To that extent they are our superiors. On the other hand, they can't learn reading, writing, etc., nor any of our fine and high things, and there we have a large advantage over them.

Y. M. Very well, let them have what they've got, and welcome; there is still a wall, and a lofty one. They haven't got the Moral Sense; we have it, and it lifts us immeasurably above them.

O. M. What makes you think that?

Y. M. Now look here—let us call a halt. I have stood the other infamies and insanities and that is enough; I am not going to have man and the other animals put on the same level morally.

O. M. I wasn't going to hoist man up to that.

Y. M. This is too much! I think it is not right to jest about such things.

O. M. I am not jesting, I am merely reflecting a plain and simple truth—and without uncharitableness. The fact that man knows right from wrong proves his *intellectual* superiority to the other creatures; but the fact that he can *do* wrong proves his *moral* inferiority to any creature that *cannot*. It is my belief that this position is not assailable.

Free Will

Y. M. What is your opinion regarding Free Will?

O. M. That there is no such thing. Did the man possess it who gave the old woman his last shilling and trudged home in the storm?

Y. M. He had the choice between succoring the old woman and leaving her to suffer. Isn't it so?

O. M. Yes, there was a choice to be made, between bodily comfort on the one hand and the comfort of the spirit on the other. The body made a strong appeal, of course—the body would be quite sure to do that;

the spirit made a counter appeal. A choice had to be made between the two appeals, and was made. Who or what determined that choice?

Y. M. Any one but you would say that the man determined it, and that in doing it he exercised Free Will.

O. M. We are constantly assured that every man is endowed with Free Will, and that he can and must exercise it where he is offered a choice between good conduct and less-good conduct. Yet we clearly saw that in that man's case he really had no Free Will: his temperament, his training, and the daily influences which had molded him and made him what he was, *compelled* him to rescue the old woman and thus save *himself*—save himself from spiritual pain, from unendurable wretchedness. He did not make the choice, it was made *for* him by forces which he could not control. Free Will has always existed in *words,* but it stops there, I think—stops short of *fact*. I would not use those words—Free Will—but others.

Y. M. What others?

O. M. Free Choice.

Y. M. What is the difference?

O. M. The one implies untrammeled power to *act* as you please, the other implies nothing beyond a mere *mental process:* the critical ability to determine which of two things is nearest right and just.

Y. M. Make the difference clear, please.

O. M. The mind can freely *select, choose, point out* the right and just one—its function stops there. It can go no further in the matter. It has no authority to say that the right one shall be acted upon and the wrong one discarded. That authority is in other hands.

Y. M. The man's?

O. M. In the machine which stands for him. In his born disposition and the character which has been built around it by training and environment.

Y. M. It will act upon the right one of the two?

O. M. It will do as it pleases in the matter. George Washington's machine would act upon the right one; Pizarro's mind would know which was the right one and which the wrong, but the Master inside of Pizarro would act upon the wrong one.

Y. M. Then as I understand it a bad man's mental machinery calmly and judicially points out which of two things is right and just—

O. M. Yes, and his *moral* machinery will freely act upon the one or the other, according to its make, and be quite indifferent to the *mind's* feelings concerning the matter—that is, *would* be, if the mind had any feelings; which it hasn't. It is merely a thermometer: it registers the heat and the cold, and cares not a farthing about either.

Y. M. Then we must not claim that if a man *knows* which of two things is right he is absolutely *bound* to do that thing?

O. M. His temperament and training will decide what he shall do, and he will do it; he cannot help himself, he has no authority over the matter. Wasn't it right for David to go out and slay Goliath?

Y. M. Yes.

O. M. Then it would have been equally *right* for any one else to do it?

Y. M. Certainly.

O. M. Then it would have been *right* for a born coward to attempt it?

Y. M. It would—yes.

O. M. You know that no born coward ever would have attempted it, don't you?

Y. M. Yes.

O. M. You know that a born coward's make and temperament would be an absolute and insurmountable bar to his ever essaying such a thing, don't you?

Y. M. Yes, I know it.

O. M. He clearly perceives that it would be *right* to try it?

Y. M. Yes.

O. M. His mind has Free Choice in determining that it would be *right* to try it?

Y. M. Yes.

O. M. Then if by reason of his inborn cowardice he simply can *not* essay it, what becomes of his Free Will? Where is his Free Will? Why claim that he has Free Will when the plain facts show that he hasn't? Why contend that because he and David *see* the right alike, both must *act* alike? Why impose the same laws upon goat and lion?

Y. M. There is really no such thing as Free Will?

O. M. It is what I think. There is *Will.* But it has nothing to do with *intellectual perceptions of right and wrong,* and is not under their command. David's temperament and training had Will, and it was a compulsory force; David had to obey its decrees, he had no choice. The coward's temperament and training possess Will, and *it* is compulsory; it commands him to avoid danger, and he obeys, he has no choice. But neither the Davids nor the cowards possess Free Will—will that may do the right or do the wrong, as their *mental* verdict shall decide.

Not Two Values, but Only One

Y. M. There is one thing which bothers me: I can't tell where you draw the line between *material* covetousness and *spiritual* covetousness.

O. M. I don't draw any.

Y. M. How do you mean?

O. M. There is no such thing as *material* covetousness. All covetousness is spiritual.

Y. M. *All* longings, desires, ambitions *spiritual,* never material?

O. M. Yes. The Master in you requires that in *all* cases you shall content his *spirit*—that alone. He never requires anything else, he never interests himself in any other matter.

Y. M. Ah, come! When he covets somebody's money—isn't that rather distinctly material and gross?

O. M. No. The money is merely a symbol—it represents in visible and concrete form a *spiritual desire*. Any so-called material thing that you want is merely a symbol: you want it not for *itself,* but because it will content your spirit for the moment.

Y. M. Please particularize.

O. M. Very well. Maybe the thing longed for is a new hat. You get it and your vanity is pleased, your spirit contented. Suppose your friends deride the hat, make fun of it: at once it loses its value; you are ashamed of it, you put it out of your sight, you never want to see it again.

Y. M. I think I see. Go on.

O. M. It is the same hat, isn't it? It is in no way altered. But it wasn't the *hat* you wanted, but only what it stood for—a something to please and content your *spirit.* When it failed of that, the whole of its value was gone. There are no *material* values; there are only spiritual ones. You will hunt in vain for a material value that is *actual, real*—there is no such thing. The only value it possesses, for even a moment, is the spiritual value back of it: remove that and it is at once worthless—like the hat.

Y. M. Can you extend that to money?

O. M. Yes. It is merely a symbol, it has no *material* value; you think you desire it for its own sake, but it is not so. You desire it for the spiritual content it will bring; if it fail of that, you discover that its value is gone. There is that pathetic tale of the man who labored like a slave, unresting, unsatisfied, until he had accumulated a fortune, and was happy over it, jubilant about it; then in a single week a pestilence swept away all whom he held dear and left him desolate. His money's value was gone. He realized that his joy in it came not from the money itself, but from the spiritual contentment he got out of his family's enjoyment of the pleasures and delights it lavished upon them. Money has no *material* value; if you remove its spiritual value nothing is left but dross. It is so with all things, little or big, majestic or trivial—there are no exceptions. Crowns, scepters, pennies, paste jewels, village notoriety, world-wide fame—they are all the same, they have no *material* value: while they content the *spirit* they are precious, when this fails they are worthless.

A Difficult Question

Y. M. You keep me confused and perplexed all the time by your elusive terminology. Sometimes you divide a man up into two or three separate personalities, each with authorities, jurisdictions, and responsibilities of its own, and when he is in that condition I can't grasp him. Now when *I* speak of a man, he is *the whole thing in one,* and easy to hold and contemplate.

O. M. That is pleasant and convenient, if true. When you speak of "my body" who is the "my"?

Y. M. It is the "me."

O. M. The body is a property, then, and the Me owns it. Who is the Me?

Y. M. The Me is *the whole thing;* it is a common property; an undivided ownership, vested in the whole entity.

O. M. If the Me admires a rainbow, is it the whole Me that admires it, including the hair, hands, heels, and all?

Y. M. Certainly not. It is my *mind* that admires it.

O. M. So *you* divide the Me yourself. Everybody does; everybody must. What, then, definitely, is the Me?

Y. M. I think it must consist of just those two parts—the body and the mind.

O. M. You think so? If you say "I believe the world is round," who is the "I" that is speaking?

Y. M. The mind.

O. M. If you say "I grieve for the loss of my father," who is the "I"?

Y. M. The mind.

O. M. Is the mind exercising an intellectual function when it examines and accepts the evidence that the world is round?

Y. M. Yes.

O. M. Is it exercising an Intellectual function when it grieves for the loss of your father?

Y. M. No. That is not cerebration, brain-work, it is a matter of *feeling.*

O. M. Then its source is not in your mind, but in your *moral* territory?

Y. M. I have to grant it.

O. M. Is your mind a part of your *physical* equipment?

Y. M. No. It is independent of it; it is spiritual

O. M. Being spiritual, it cannot be affected by physical influences?

Y. M. No.

O. M. Does the mind remain sober when the body is drunk?

Y. M. Well—no.

O. M. There *is* a physical effect present, then?

Y. M. It looks like it.

O. M. A cracked skull has resulted in a crazy mind. Why should that happen if the mind is spiritual, and *independent* of physical influences?

Y. M. Well—I don't know.

O. M. When you have a pain in your foot, how do you know it?

Y. M. I feel it.

O. M. But you do not feel it until a nerve reports the hurt to the brain. Yet the brain is the seat of the mind, is it not?

Y. M. I think so.

O. M. But isn't spiritual enough to learn what is happening in the outskirts without the help of the *physical* messenger? You perceive that the question of who or what the Me is, is not a simple one at all. You say "I admire the rainbow," and "I believe the world is round," and in these cases we find that the Me is not all speaking, but only the *mental* part. You say "I grieve," and again the Me is not all speaking, but only the *moral* part. You say the mind is wholly spiritual; then you say "I have a pain" and find that this time the Me is mental *and* spiritual combined. We all use the "I" in this indeterminate fashion, there is no help for it. We imagine a Master and King over what you call The Whole Thing, and we speak of him as " I, " but when we try to define him we find we cannot do it. The intellect and the feelings can act quite *independently* of each other; we recognize that, and we look around for a Ruler who is master over both, and can serve as a *definite and indisputable "I,"* and enable us to know what we mean and who or what we are talking about when we use that pronoun, but we have to give it up and confess that we cannot find him. To me, Man is a machine, made up of many mechanisms, the moral and mental ones acting automatically in accordance with the impulses of an interior Master who is built out of born-temperament and an accumulation of multitudinous outside influences and trainings; a machine whose *one* function is to secure the spiritual contentment of the Master, be his desires good or be they evil; a machine whose Will is absolute and must be obeyed, and always *is* obeyed.

Y. M. Maybe the Me is the Soul?

O. M. Maybe it is. What is the Soul?

Y. M. I don't know.

O. M. Neither does any one else.

The Master Passion

Y. M. What is the Master?—or, in common speech, the Conscience? Explain it.

O. M. It is that mysterious autocrat, lodged in a man, which compels the man to content its desires. It may be called the Master Passion—the hunger for Self-Approval.

Y. M. Where is its seat?

O. M. In man's moral constitution.

Y. M. Are its commands for the man's good?

O. M. It is indifferent to the man's good; it never concerns itself about anything but the satisfying of its own desires. It can be *trained* to prefer things which will be for the man's good, but it will prefer them only because they will content *it* better than other things would.

Y. M. Then even when it is trained to high ideals it is still looking out for its own contentment, and not for the man's good?

O. M. True. Trained or untrained, it cares nothing for the man's good, and never concerns itself about it.

Y. M. It seems to be an *immoral* force seated in the man's moral constitution?

O. M. It is a *colorless* force seated in the man's moral constitution. Let us call it an instinct—a blind, unreasoning instinct, which cannot and does not distinguish between good morals and bad ones, and cares nothing for results to the man provided its own contentment be secured; and it will *always* secure that.

Y. M. It seeks money, and it probably considers that that is an advantage for the man?

O. M. It is not always seeking money, it is not always seeking power, nor office, nor any other *material* advantage. In *all* cases it seeks a *spiritual* contentment, let the *means* be what they may. Its desires are determined by the man's temperament—and it is lord over that. Temperament, Conscience, Susceptibility, Spiritual Appetite, are, in fact, the same thing. Have you ever heard of a person who cared nothing for money?

Y. M. Yes. A scholar who would not leave his garret and his books to take a place in a business house at a large salary.

O. M. He had to satisfy his master—that is to say, his temperament, his Spiritual Appetite—and it preferred the books to money. Are there other cases?

Y. M. Yes, the hermit.

O. M. It is a good instance. The hermit endures solitude, hunger, cold, and manifold perils, to content his autocrat, who prefers these things, and prayer and contemplation, to money or to any show or luxury that money can buy. Are there others?

Y. M. Yes. The artist, the poet, the scientist.

O. M. Their autocrat prefers the deep pleasures of these occupations, either well paid or ill paid, to any others in the market, at any price. You *realize* that the Master Passion—the contentment of the spirit concerns itself with many things besides so called material advantage, material prosperity, cash, and all that?

Y. M. I think I must concede it.

O. M. I believe you must. There are perhaps as many Temperaments that would refuse the burdens and vexations and distinctions of public office as there are that hunger after them. The one set of Temperaments seek the contentment of the spirit, and that alone; and this is exactly the case with the other set. Neither set seeks anything *but* the contentment of the spirit. If the one is sordid, both are sordid; and equally so, since the end in view is precisely the same in both cases. And in both cases Temperament decides the preference—and Temperament is *born,* not made.

Conclusion

O. M. You have been taking a holiday?

Y. M. Yes; a mountain tramp covering a week. Are you ready to talk?

O. M. Quite ready. What shall we begin with?

Y. M. Well, lying abed resting up, two days and nights, I have thought over all these talks, and passed them carefully in review. With this result: that... that... are you intending to publish your notions about Man some day?

O. M. Now and then, in these past twenty years, the Master inside of me has half-intended to order me to set them to paper and publish them. Do I have to tell you why the order has remained unissued, or can you explain so simple a thing without my help?

Y. M. By your doctrine, it is simplicity itself: outside influences moved your interior Master to give the order; stronger outside influences deterred him. Without the outside influences, neither of these impulses could ever have been born, since a person's brain is incapable of originating an idea within itself.

O. M. Correct. Go on.

Y. M. The matter of publishing or withholding is still in your Master's hands. If some day an outside influence shall determine him to publish, he will give the order, and it will be obeyed.

O. M. That is correct. Well?

Y. M. Upon reflection I have arrived at the conviction that the publication of your doctrines would be harmful. Do you pardon me?

O. M. Pardon *you?* You have done nothing. You are an instrument—a speaking-trumpet. Speaking-trumpets are not responsible for what is said through them. Outside influences—in the form of lifelong teachings, trainings, notions, prejudices, and other second-hand importations—have persuaded the Master within you that the publication of these doctrines would be harmful. Very well, this is quite natural, and was to be expected; in fact, was inevitable. Go on; for the sake of ease and convenience, stick to habit: speak in the first person, and tell me what your Master thinks about it.

Y. M. Well, to begin: it is a desolating doctrine; it is not inspiring, enthusing, uplifting. It takes the glory out of man, it takes the pride out of him, it takes the heroism out of him, it denies him all personal credit, all applause; it not only degrades him to a machine, but allows him no control over the machine; makes a mere coffee-mill of him, and neither permits him to supply the coffee nor turn the crank, his sole and piteously humble function being to grind coarse or fine, according to his make, outside impulses doing all the rest.

O. M. It is correctly stated. Tell me—what do men admire most in each other?

Y. M. Intellect, courage, majesty of build, beauty of countenance, charity, benevolence, magnanimity, kindliness, heroism, and—and—

O. M. I would not go any further. These are *elementals.* Virtue, fortitude, holiness, truthfulness, loyalty, high ideals—these, and all the related qualities that are named in the dictionary, are *made of the elementals,* by blendings, combinations, and shadings of the elementals, just as one makes green by blending blue and yellow, and makes several shades and tints of red by modifying the elemental red. There are seven elemental colors; they are all in the rainbow; out of them we manufacture and name fifty shades of them. You have named the elementals of the human rainbow, and also one *blend* —heroism, which is made out of courage and magnanimity. Very well, then; which of these elements does the possessor of it manufacture for himself? Is it intellect?

Y. M. No.

O. M. Why?

Y. M. He is born with it.

O. M. Is it courage?

Y. M. No. He is born with it.

O. M. Is it majesty of build, beauty of countenance?

Y. M. No. They are birthrights.

O. M. Take those others—the elemental moral qualities—charity, benevolence, magnanimity, kindliness; fruitful seeds, out of which spring, through cultivation by outside influences, all the manifold blends and combinations of virtues named in the dictionaries: does man manufacture any one of those seeds, or are they all born in him?

Y. M. Born in him.

O. M. Who manufactures them, then?

Y. M. God.

O. M. Where does the credit of it belong?

Y. M. To God.

O. M. And the glory of which you spoke, and the applause?

Y. M. To God.

O M. Then it is *you* who degrade man. You make him claim glory, praise, flattery, for every valuable thing he possesses—*borrowed* finery, the whole of it; no rag of it earned by himself, not a detail of it produced by his own labor. *You* make man a humbug; have I done worse by him?

Y. M. You have made a machine of him.

O. M. Who devised that cunning and beautiful mechanism, a man's hand?

Y. M. God.

O. M. Who devised the law by which it automatically hammers out of a piano an elaborate piece of music, without error, while the man is thinking about something else, or talking to a friend?

Y. M. God.

O. M. Who devised the blood? Who devised the wonderful machinery which automatically drives its renewing and refreshing streams through the body, day and night, without assistance or advice from the man? Who devised the man's mind, whose machinery works automatically, interests itself in what it pleases, regardless of his will or desire, labors all night when it likes, deaf to his appeals for mercy? God devised all these things. *I* have not made man a machine, God made him a machine. I am merely calling attention to the fact, nothing more. Is it wrong to call attention to the fact? Is it a crime?

Y. M. I think it is wrong to *expose* a fact when harm can come of it.

O. M. Go on.

Y. M. Look at the matter as it stands now. Man has been taught that he is the supreme marvel of the Creation; he believes it; in all the ages he has never doubted it, whether he was a naked savage, or clothed in purple and fine linen, and civilized. This has made his heart buoyant, his life cheery. His pride in himself, his sincere admiration of himself, his joy in what he supposed were his own and unassisted achievements, and his exultation over the praise and applause which they evoked—these have exalted him, enthused him, ambitioned him to higher and higher flights; in a word, made his life worth the living. But by your scheme, all this is abolished; he is degraded to a machine, he is a nobody, his noble prides wither to mere vanities; let him strive as he may, he can never be any better than his humblest and stupidest neighbor; he would never be cheerful again, his life would not be worth the living.

O. M. You really think that?

Y. M. I certainly do.

O. M. Have you ever seen me uncheerful, unhappy?

Y. M. No.

O. M. Well, I believe these things. Why have they not made me unhappy?

Y. M. Oh, well—temperament, of course! You never let *that* escape from your scheme.

O. M. That is correct. If a man is born with an unhappy temperament, nothing can make him happy; if he is born with a happy temperament, nothing can make him unhappy.

Y. M. What—not even a degrading and heartchilling system of beliefs?

O. M. Beliefs? Mere beliefs? Mere convictions? They are powerless. They strive in vain against inborn temperament.

Y. M. I can't believe that, and I don't.

O. M. Now you are speaking hastily. It shows that you have not studiously examined the facts. Of all your intimates, which one is the happiest? Isn't it Burgess?

Y. M. Easily.

O. M. And which one is the unhappiest? Henry Adams?

Y. M. Without a question!

O. M. I know them well. They are extremes, abnormals; their temperaments are as opposite as the poles. Their life-histories are about alike—but look at the results! Their ages are about the same around about fifty. Burgess has always been buoyant, hopeful, happy; Adams has always been cheerless, hopeless, despondent. As young fellows both tried country journalism—and failed. Burgess didn't seem to mind it; Adams couldn't smile, he could only mourn and groan over what had happened and torture himself with vain regrets for not having done so and so instead of so and so—*then* he would have succeeded. They tried the law—and failed. Burgess remained happy—because he couldn't help it. Adams was wretched—because he couldn't help it. From that day to this, those two men have gone on trying things and failing: Burgess has come out happy and cheerful every time; Adams the reverse. And we do absolutely know that these men's inborn temperaments have remained unchanged through all the vicissitudes of their material affairs. Let us see how it is with their immaterials. Both have been zealous Democrats; both have been zealous Republicans; both have been zealous Mugwumps. Burgess has always found happiness and Adams unhappiness in these several political beliefs and in their migrations out of them. Both of these men have been Presbyterians, Universalists, Methodists, Catholics—then Presbyterians again, then Methodists again. Burgess has always found rest in these excursions, and Adams unrest. They are trying Christian Science, now, with the customary result, the inevitable result. No political or religious belief can make Burgess unhappy or the other man happy. I assure you it is purely a matter of temperament. Beliefs are *acquirements,* temperaments are *born;* beliefs are subject to change, nothing whatever can change temperament.

Y. M. You have instanced extreme temperaments.

O. M. Yes. The half-dozen others are modifications of the extremes. But the law is the same. Where the temperament is two-thirds happy, or two thirds unhappy, no political or religious beliefs can change the proportions. The vast majority of temperaments are pretty equally balanced; the intensities are absent, and this enables a nation to learn to accommodate itself to its political and religious circumstances and like

them, be satisfied with them, at last prefer them. Nations do not *think,* they only *feel.* They get their feelings at second hand through their temperaments, not their brains. A nation can be brought—by force of circumstances, not argument—to reconcile itself to *any kind of government or religion that can be devised;* in time it will fit itself to the required conditions; later, it will prefer them and will fiercely fight for them. As instances, you have all history: the Greeks, the Romans, the Persians, the Egyptians, the Russians, the Germans, the French, the English, the Spaniards, the Americans, the South Americans, the Japanese, the Chinese, the Hindoos, the Turks—a thousand wild and tame religions, every kind of government that can be thought of, from tiger to housecat, each nation *knowing* it has the only true religion and the only sane system of government, each despising all the others, each an ass and not suspecting it, each proud of its fancied supremacy, each perfectly sure it is the pet of God, each with undoubting confidence summoning Him to take command in time of war, each surprised when He goes over to the enemy, but by habit able to excuse it and resume compliments—in a word, the whole human race content, always content, persistently content, indestructibly content, happy, thankful, proud, *no matter what its religion is, nor whether its master be tiger or house-cat.* Am I stating facts? You know I am. Is the human race cheerful? You know it is. Considering what it can stand, and be happy, you do me too much honor when you think that *I* can place before it a system of plain cold facts that can take the cheerfulness out of it. Nothing can do that. Everything has been tried. Without success. I beg you not to be troubled.

What Are Autonomy & Responsibility?

Rick Marken
(Life Learning Associates, 10459 Holman Ave., Los Angeles, CA 90024)

Bill Powers
(73 Ridge Place, CR 510, Durango, CO 81301)

Powers: I use the term "autonomy" in a way that has a rather complicated meaning, with a basis farther back than current behavioral interactions with the current environment. I'll slip in my hypotheses about the role of control without marking them; I'm sure others will be able to tell what is hypothesis from what is generally accepted "fact" in this story.

Start with DNA. While the surface appearance is that genetic characteristics are transmitted via the DNA molecule, in fact a lot more passes from generation to generation than just DNA. Much cellular material, as in mitochondria, is passed along with the DNA through the mother's egg; in the lowest orders, the cellular material simply splits up during the reproductive divisions. Immediately after a new individual is launched, the DNA is in an environment that is continuous with the previous environment, at least locally.

So the biochemical control systems whose reference signals are carried in DNA can operate right across the boundary between generations. These control systems, finding themselves isolated, begin again building the control systems that build the control systems that build the control systems that constitute the adult organism. The entire interior milieu is regenerated, with whatever changes that occurred during the division of the genetic material. The continuity proceeds, as people have long suspected, through the mother.

One of the final products of this process is a set of intrinsic reference signals. These reference signals are the basis of reorganization or learning, through which the new organism establishes control in the environments it first and subsequently encounters. The intrinsic reference signals represent the target states of some as yet poorly defined set of variables critical to the survival of the individual. There is no reason to think that the reference signals are identically set from one person to the next, or even that they are all of the same kind. Each individual differs in details of organization at all levels from DNA through cellular and organ structures through gross bodily structure through neural circuitry. And the mix of intrinsic reference signals will differ from individual to individual.

Intrinsic reference signals are part of a system, probably distributed rather than lumped in one place, that controls for zero intrinsic error. The means of control is blind variation of the organization of the nervous system and the biochemical control systems. Reorganization is driven by intrinsic error, and it ceases when intrinsic error drops below some threshold. As a result, the organism acquires control systems that can maintain perceived aspects of the external world at learned reference levels by means of motor behavior (and at the biochemical level, changes in such things as strength, speed, organ size and activity, and so forth). The criterion for acquiring any behavioral control system, and for setting its reference signal to any specific value, is that intrinsic error be maintained at the lowest possible level.

Thus the overriding concern of the reorganizing system, and the purpose for which it causes any behavioral organization to appear, is to control its own basic physical state; to maintain its component variables at endogenously determined reference levels. It neither knows about, nor cares about, nor *can* care about any processes external to the body. Everything it causes to be done by way of interacting with an external world is done for the purpose of controlling an internal state. It is therefore completely and absolutely autonomous in its purposes.

It does, to be sure, have a history. But this is not so much a history of antecedent events as it is a history of gradually changing organization. The reorganizing system of one generation is continuous with the reorganizing systems of previous generations: it is the same system, evolving. At the center of this system are reference signals that have not changed in billions of years, having survived even speciation. Reorganizations that preserve these basic reference signals have led to the development of instrumental reference signals and associated control systems, and those have led to still more elaborate control mechanisms, and so on to the various physical forms that life has ultimately adopted—as a means of preserving the fundamental function, which is to control. And to control is the ultimate meaning of being autonomous.

If the criterion for stopping reorganization is bringing intrinsic variables to their respective reference levels, it follows that only those behavioral control systems will survive reorganization that do entail actual effects of the right kinds on the intrinsic variables. The effect of any given behavioral act is not determined by the organism: it is determined by the nature of the surrounding world (including the behavioral organization of other organisms in that world). So reorganization

can't cease until the actual effects on intrinsic states, via that external world, are correct for maintaining zero intrinsic error.

Thus the organism learns first what variables are critical to perceive in that external world, and second what specific states of those variables are critical to maintain. This process of learning has been going on through geological time, with the appearance of control structures of greater and greater generality, and what we recognize as higher and higher levels of control. As each new level of control appeared, new and more important aspects of the environment became perceivable and came under control by the organism. The actions of the organism adapted themselves to the environment in more and more subtle ways.

The means of action did not change nearly as much as the neural control systems that use actions to control ever more complex variables. A human being and a monkey share nearly identical means of motor action. Both have hands at the ends of jointed limbs; but the human being can accomplish things with its hands that a monkey cannot. This is not because of having an opposable thumb, but because of having higher levels of control. Human beings can do more even with their thumbs cut off than a monkey can do with ten digits.

So we arrive finally at the question of autonomy in the individual human being. Autonomy is clearly not freedom from physical constraints (which include, in the final analysis, social constraints). The environment, not the organism, dictates the effects of any given action. But the environment does not dictate the desired consequences of any action. It is the organism that chooses those consequences and learns how it must act in order to produce them.

In a hierarchical control system, built, I presume, level by level over the eons and recapitulated in the individual, the lower systems give up their autonomy to the higher systems that manipulate their reference signals. At whatever level is currently the highest, the reference signals are set from within the organism by the process of reorganization; the purpose of choosing a particular setting is to maintain intrinsic error as close to zero as possible—as the purpose has always been. In order to bring the highest level of perception into a match with this autonomously set reference signal, the highest control systems must, as usual, be altered to produce actions which are among those that will have the required effects. Now those actions are determined by properties of the existing lower levels as well as by the characteristics of the world external to the organism.

The organism can't choose what properties the external world will have, no matter what the level of perception. Once its lower levels have been built and brought into mutual harmony, the organism has less than a completely free choice even as to the kinds of actions it can produce (without starting again from scratch, which is probably no longer possible in the adult organism, in the time remaining to it). So the particular behavioral organizations that appear in the adult are shaped by the properties of the world around it and by the properties of its own already-acquired lower levels of control.

However, the highest levels of reference signals remain autonomous and are changed only in service of maintaining the individual organism's mix of intrinsic variables at their unique mix of reference settings. The external world has no influence over that basic requirement. Intrinsic error remains the organism's sole criterion for judging the value of any aspect of its experiences. This is true of all organisms from the amoeba to the human being.

If the highest levels of reference signals are autonomously determined, then the next-to-highest levels of reference signals are varied so as to prevent the environment (as perceived through all the lower levels) from making the highest perceptions depart materially from their reference settings. This means that the next-to-highest levels of perception will also be shaped to meet the requirements of the highest reference signals.

But the next-to-highest reference signals will be determined by what the environment requires, for the highest perceptual signals in general contain effects of uncontrolled elements. To make the net result match what the highest system requires, the reference signals for the controllable parts of the next-to-highest system must be varied, and those variations must be matched to the properties of the lower systems and the external world. The organism can't choose the settings freely, because only certain settings will result in the required perceptions. There might be many alternative settings that will produce the required perceptions, but there is freedom to choose only among those alternatives, given that the highest reference signals are to be satisfied. All other alternatives are ruled out by properties of the external world.

The general picture is that the environment determines behavior, while the autonomous organism determines consequences of behavior. Given the intended consequences, the environment sets the limits as to what lower-level actions can in fact bring those consequences about.

So we can see where autonomy begins and ends. It is the organism that selects consequences that keep its intrinsic errors as close to zero as possible. It is the environment—including other organisms—that determines what actions must be produced in order that those consequences be brought about and maintained. The external world sets the stage on which existence is played out. But the reorganizing system writes the play.

Reorganization occurs precisely when interactions with the organism's niche lead to loss of control—that is, when the current regularities in the interactions are insufficient to preserve control.

The reorganizing system is effective because the changes it institutes do *not* depend in any regular way on the current organization or the current niche. The whole point is to break out of the conflict or the circle or the failure—the local minimum—by trying something *new.* So the idea of tracing the current organization backward, while all right in a general sense, is wrong if it implies any predictable course of development. Reorganization breaks the cause-effect chain.

Reorganization—that is, the actual output effect of the process—is independent (save for the frequency with which reorganizations occur) of any prior causes. The outcome of reorganization, to be sure, has to be such that intrinsic error is corrected; if it's not, reorganization simply continues. But there are uncountable ways of reorganizing that would result in correcting intrinsic error, so that result is not a constraint on any particular act of reorganization. In fact, one episode of reorganization is just as likely to make matters worse as it is to make them better (unless, of course, there is an unsuspected systematic component in it). The statistics of reorganization are very different from the statistics of stochastic—but on the average systematic—causation. Reorganization will work even when the changes it produces show no trend at all in any direction.

Also, we mustn't forget that what makes reorganization effective are not the individual reorganizational events, but the selection effects that terminate reorganization. All that is required is the existence of something that can say, "There! That feels better." Or, of course, something that says, "Oh, no! Reorganize!" In fact, we could accept a mechanical randomness generator that actually does the reorganizing acts, and limit free will to the single act of triggering a reorganization. The "awareness" part of free will would then superimpose its judgments of what is acceptable and what is not on automatic judgments about such things as body temperature and state of nourishment. Thus free will could select for outcomes acceptable to it simply by causing reorganizations until the result is acceptable. The grounds for acceptability need have nothing to do with the niche.

And even the reorganizing system is just the product of a deeper control process, at the core of which lies a tiny and unimaginably ancient spark of purpose that makes life different from everything else.

Marken: I would probably have handled the Branch Davidian situation in Waco, Texas, about the same as it was actually handled, if I had the same high-level goals as the participants. The initial "confrontation" occurred (as I recall) because the ATF had the goal of regulating firearms—so they went to the Davidian compound to confiscate firearms there. Four ATF people got killed in the process. So FBI agents (as I recall) surrounded the place and tried to get the people out, because they had the goal of arresting and bringing to trial people suspected of killing other people. There was obviously a conflict of goals—the FBI wanted the Davidians to get out to stand trial; the Davidians wanted to stay in, probably so as not to stand trial but possibly for other reasons, like waiting for the apocalypse.

There were no "mistakes" made in Waco—everybody was trying to control for perceptions that were important to them (for higher-order reasons) and doing the best they could. Saying that one group or the other had the "wrong" goals (I've heard people say that the FBI shouldn't have wanted to flush out the Davidians with gas) seems pretty non-PCT to me; people set their goals to satisfy higher-order references, and the particular settings of these lower-order goals depend on disturbances and other higher-level goals as much as on the higher-order goal itself.

It seems to me that legal guilt requires that two things be demonstrated about the "offending" result (such as burning up 100 people in a house): it must be shown that the result was intended, and that the person who produced the result knew that it was "wrong." PCT shows that some results are intended (controlled results), and some are not (uncontrolled side-effects of the outputs that produce control). Both adults and children can produce results intentionally, so there is no "guilty" distinction here. HPCT suggests that the "wrongness" of a result is probably not perceived until you get to the category level—where you can perceive many different lower-order perceptual results as "wrong" and other lower-order perceptual results as "right." I think that it's possible that people don't completely flesh out the categorical distinctions between right and wrong (as defined in the context of interactions with other people, of course) until they are well into their teens. I think society recognizes this fact and, because of it, treats juveniles (who have intentionally produced "wrong" results) differently than adults who have intentionally produced the same results. Since the development of the hierarchy of perception is likely to occur at quite different rates in different people, the line between innocent children (who intentionally do "wrong") and guilty adults will always be fuzzy, legally.

The "wrongness" of references (in a PCT sense) can only be defined in terms of the higher-order goals that they are set to satisfy. In this sense, setting wrong reference levels just means that you have not yet learned to control the higher-level variable whose value is influenced by the setting of the lower level reference. When you are in control, then, by definition, your settings of the lower-level references that influence the controlled variable are, indeed, always right—because they result in control. Whether or not you, as an observer, think that these reference settings are right or wrong is quite another story. But,

again, the wrongness of the other person's reference settings *for you* depends on your own reference settings for the same perceptual variables. Wrongness is always defined with respect to the references of the observer.

I think that PCT shows that the legal conditions for guilt are real aspects of human nature (in contrast to the behaviorist position attributing all behavior to the environment). What PCT doesn't tell you is whether the results that are produced intentionally are *really* ("objectively") wrong. I think PCT can help us get away from the hopeless quagmire of arguing about which results are *really* right and which are *really* wrong and reframe ethics in terms of control. If it helps people control, it's right; if it prevents them from controlling, it's wrong. So control is right; conflict is wrong.

Powers: We are *not* responsible for our actions. What we are responsible for (that is, are the cause of) are the goals that require the actions. Once you've picked a goal, a perception to maintain in some specific state, from then on your actions relative to that goal are determined by disturbances. If there's any goal that you can't reorganize, then from that level down the environment controls your actions with respect to that goal.

So once David Koresh had settled on his goal of never surrendering, and once the government coordinator had settled on the goal of getting Koresh to surrender, each side's actions were determined by the disturbances from the other side. All that kept the situation from escalating to its ultimate conclusion immediately was internal conflict: Koresh and his followers did not relish dying, and the government coordinator did not relish killing. But the fixed goals eventually had their way. This is what beliefs accomplish: they set fixed goals, and as a result leave the environment and other people in charge of one's actions.

Marken: But how do we decide which goals are "picked" and which are responses to disturbance? If we take the hierarchical model seriously, not even the highest-level goals in the hierarchy are "picked" arbitrarily—their selection is constrained by intrinsic goals which are "picked" by evolution. So all goals are ultimately varied as a means of compensating for disturbances to the intrinsic goals. "Responsibility" is in the doghouse in the sense that a hierarchical arrangement of control systems has no way of "picking" goals at *any* level of the hierarchy.

A particular control system can be considered responsible for producing the particular perception demanded by its reference input; but it is not responsible for how it produces this reference perception (because that is determined mainly by environmental disturbances to the controlled perception); nor is it responsible for the particular level of the reference input being sent to it. So a system controlling the level of honesty perceived in its relations with other people is responsible only for maintaining this perception (degree of honesty) at the level specified by the reference input (say, "very honest"). It is not responsible for *how* it maintains that perception (sometimes it might require telling a person that he or she is being very nice and sometimes it might require telling the same person that he or she is being a jerk). Nor is it responsible for the *level* of honesty that is specified by the reference input—be it "very honest," "moderately honest," or "deceitful."

This is not a very "strong" kind of responsibility. I think we typically use the word "responsibility" to suggest that people are "choosing" what we think of as good or bad goals. But goals (the settings of the reference inputs) are determined by higher-order systems as a means of controlling perceptions (and, hence, resisting disturbances). So there is really no "choice" in goal selection; the goal for the "honesty" control system, for example, is determined by disturbances to higher-order variables (such as "political success"), not by the control system itself.

Since control systems are not really responsible (in the strong sense —meaning "choosing their own goals") for their goal-setting, it does not seem to me to make much sense to judge control-system goals in terms of "conventional morality"—in which some goals are good, others are bad, and the behavior of the system is judged on the basis of its selection of good vs. bad goals. Conventional morality assumes that the system is responsible—i.e., that each goal-attaining component of the system has chosen its goals on its own. In a control hierarchy, each goal-attaining component (individual control system) does *not* select its goals—it simply *achieves* them.

So how do we judge the "goodness" or "badness" of control systems? I have already suggested a way; control systems should be judged only in terms of how well they control (and, I should add, how well they allow *other* control systems, of the same type, to control; I don't think this addendum is really necessary because an individual would not control well for long if it were busy screwing up fellow controllers; but it is possible, in the short run, to control well by interfering with others, so I'll leave it in).

A control system that selects goals well is a control system that *controls* well; when you control well, it feels great. I think the experience of being in control like this is what religious people call *grace*. A control hierarchy can't achieve this kind of grace if it cannot select goals well. A control system cannot select goals well if it is in internal *or* external *conflict*; and it apparently cannot select goals well if some goals are fixed by belief. Both of these problems seemed to contribute to the tragedy in Waco.

The Control Systems Group is a membership organization which supports the understanding of cybernetic control systems in organisms and their environments: *living control systems.* Academicians, clinicians, and other professionals in several disciplines, including biology, psychology, social work, economics, education, engineering, and philosophy, are members of the Group. Annual meetings have been held since 1985. The CSG Business Office is located at 73 Ridge Pl., CR 510, Durango, CO 81301; phone (303) 247-7986.

The CSG logo shows the generic structure of cybernetic control systems. A Comparator (C) computes the difference between a reference signal (represented by the arrow coming from above) and the output signal from Sensory (S) computation. The resulting difference signal is the input to the Gain generator (G). Disturbances (represented by the black box) alter the Gain generator output on the way to Sensory computation, where the negative-feedback km/ is closed.

Back cover

Closed Loop

Journal of Living Control Systems

Summer 1994

Volume 4

Number 3

Front cover

Closed Loop

Journal of Living Control Systems

Summer 1994 Volume 4 Number 3

Edited by Greg Williams, 460 Black Lick Rd., Gravel Switch, KY 40328

CONTENTS

Back issues of *Closed Loop* are available from Mary Powers, 73 Ridge Place, CR 510, Durango, CO 81301, phone (303)247-7986.

CSG Notes

Mary Powers
(73 Ridge Pl., CR 510, Durango, CO 81301)

This is the last issue of *Closed Loop* for now. Perhaps it will be revived in the future, but at this time the contributors we hoped to have are directing their energies elsewhere: to CSGnet or to papers for journals in their own fields.

On behalf of the membership, I want to thank Greg Williams for his time and effort over the past four years as editor and publisher. He worked for very little more than the satisfaction of producing an attractive, professional publication, and we owe him our appreciation and thanks.

My main concern with ending the publication of *Closed Loop* is losing touch with those members for whom this has been the only contact with the group. I will send conference information to everyone on the mailing list, and meanwhile I hope that they can find their way onto the information superhighway.

CSGnet, the electronic mail network for individuals interested in control theory as applied to living systems (perceptual control theory, PCT), is a lively forum for sharing ideas, asking questions, and learning more about the theory. There are no sign-up or connect-time charges for participation on CSGnet. The Internet address is csg-l@vmd.cso.uiuc.edu and csg-l@uiucvmd is the Bitnet address. Messages sent to CSGnet via these addresses are automatically forwarded to over 100 participants on five continents, as well as to hundreds of NetNews (Usenet) sites where CSGnet can be found as the newsgroup bit.listserv.csg-I.

CSGnet can also be accessed via CompuServe, AT&T Mail, MCI Mail, or any other computer communication service with a gateway to Internet or Bitnet. For more information about subscribing to CSGnet, contact Gary Cziko, the network manager, at g-cziko@uiuc.edu or phone him at (217) 333-8577.

The officers of CSG agree that members are entitled to a refund of that portion of dues received which was intended to cover the cost of producing and mailing Closed Loop. If you wish to receive a refund of $25, please contact me by e-mail at powers_w%flc@vaxf.colorado.edu or by writing to the address above. Alternatively, if you would like to contribute your refund to the fund for student conference expenses, then do nothing. If you sent in dues in the early summer in lieu of coming to the conference, *or* registered for the conference (whether or not you were able to come), *or* sent in dues after the conference and up to the present time, you are entitled to a refund.

If you have not gotten around to renewing for 1995, dues are now $20 (students, $5). For this, you will get conference announcements, perhaps other mailings about CSG doings, and the satisfaction of supporting CSG and helping students whose interest in PCT we want to encourage. Depending on our bank balance, we might also support some research in a small way.

Following are some notes on recent and upcoming activities of the Control Systems Group:

CSG Conference, Gregynog, Wales, June, 1994: This was the first European conference of the CSG, with participants from the U.K., the U.S., Germany, and France. A *Proceedings* is available from Marcos Rodriguez, Computer Science Dept., University of Wales, Aberystwyth, Wales, U.K., for 15 pounds. While interesting, it does not capture the flavor of the conference itself, which included several presentations not included in the *Proceedings,* and which primarily served (in my opinion) to align the various reports more closely to the PCT perspective. I will not attempt to summarize these reports (participants can send their own summaries to this column) except to remark that PCT provided a common ground for a disparate crew of researchers (in robotics, computer science, experimental psychology, sociology, and enzyme systems—a rather different mix from CSG conferences in the U.S.) to exchange views and learn from one another. Emerging from the conference was a proposal to try to engage the prosthetics community in the design of control systems-oriented devices. An effort is being made to convert videos of this conference to U.S.-compatible tapes, and information about these will be published when available.

CSG Conference, Durango CO, July, 1994: This was our 10th conference. Again, I will not attempt to summarize. Videotapes of this meeting are available from Dag Forssell, 23903 Via Flamenco, Valencia, CA 91355 ($35 for three six-hour tapes). A unique feature of this conference was the presentation by three educators of their experiences using a PCT approach to create a peaceful and orderly environment for learning in the classroom. Much of what they had to say is in Ed Ford's book, *Discipline in the Home and School* (Brandt Publishing, 10209 N. 56th St, Scottsdale, AZ 85253). That's a whip-cracking title, but the reality of the approach is to teach children who raise hell in the classroom, and who prevent other children from learning, to understand their own goals and learn alternative, less obstructive ways of achieving them.

Ongoing research: Tom Bourbon is currently developing a research program to make quantitative measurements of performance in spinal injury

patients before and after corrective surgery. Current assessments of patients' abilities before and after surgery are qualitative and subjective.

Ongoing research: Bruce Abbott, Bill Powers, and Rick Marken (and any others who are interested) are beginning a project to reanalyze operant conditioning (the matching law, etc.) in terms of PCT. This will be ongoing on CSGnet.

Conference: Hugh Petrie is arranging a presentation by CSG members at the American Educational Research Association in April in San Francisco. Participants will include Gary Cziko, Ed Ford, and Bill Powers.

Conference: The annual meeting of the CSG will be at Fort Lewis College, Durango, CO, July 19-23, 1995.

Publishing: CSG Book Publishing has transferred publishing and distribution of *Living Control Systems* and *Living Control Systems II* by William T. Powers, *Introduction to Modern Psychology* edited by Richard J. Robertson and William T. Powers, and *Mind Readings* by Richard S. Marken, to New View Publications, P.O. Box 3021, Chapel Hill, NC 27515-3021. Contact New View for ordering information.

Catalog of CSG Archive Materials

Greg Williams
(460 Black Lick Rd., Gravel Switch, KY 40328)

This lists holdings as of November 1994. Multiple copies of items have not been noted, and originals have not been differentiated from photocopies. In some cases, dating was made by inference from various types of clues, such as dates on computer files from which printouts were made. Some papers published in CONTINUING THE CONVERSATION are actually housed in the Gregory Bateson Archive at the same address as the CSG Archive (given above).

Any and all additions and corrections are invited. In particular, copies of papers presented and distributed at the 1994 CSG meeting would be most welcome.

Many individuals have provided materials and support for the Archive. I especially want to thank Tom Bourbon, Bill Powers, and Mary Powers for giving financial contributions as well as donating materials. Robert K. Clark, Gary Cziko, Ed Ford, David Goldstein, Phil Lewin, Clark McPhail, Richard J. Robertson, Philip Runkel, Stuart Umpleby, and Michael Yocum donated hard-to-find materials.

I plan to keep this catalog up-to-date as a computer database allowing quick searching for names and titles. For a copy of the latest version of the database (for IBM-PC-compatibles running MS-DOS), send $10 to me at the address given above.

NOT ITEMIZED

Computer disks with CSG-related files, 1 box.

Correspondence and other materials about CSG, 1 box.

"Little man" computer program materials, including papers, computer disks, and correspondence between William T. POWERS and Greg WILLIAMS, 1 box.

Photographs and slides of CSG members and events.

Unpublished papers and correspondence of William T. POWERS, 1 box.

NO DATE

BAUM, Cindy S., “Full circle,” unpublished term paper for seminar taught by William T. POWERS at Northwestern University, Evanston, Illinois, 1971-1972, 6 + 5 pages. (With letter from Bill POWERS to Greg WILLIAMS dated October 8, 1989, and letter from Greg WILLIAMS to Cindy STACKHOUSE dated February 5, 1990, including copies of redone figures for this paper.)

BRADY, Pat, MANIA AND ITS RESIDUAL FALLOUT, unpublished manuscript, 183 pages. (With letter from Pat BRADY to William T. POWERS dated June 21, 1989.)

Button, imprinted with “I’M IN CONTROL[,] INTERGRAPH Distributed Publishing System.”

Cap, imprinted with “Control Theorists Do It With LOOP GAIN,” blue and white.

CLARK, Robert K., Richard D. CHESSICK, and Robert L. McFARLAND, “High speed data processing—Compromises and considerations,” unpublished paper, 6 pages. (Given by principal author to Greg WILLIAMS in December 1988.)

Corporate seal of the Control Systems Group, Illinois.

CUNNINGHAM, W. B., M. M. TAYLOR, J. R. GABRIEL, J. H. DISCENZA, and T. B. BAINES, “Optimal decision support[:] A perception based paradigm for decision centered information fusion,” unpublished paper, paginated by section. (Probably 1993.)

“Executive decision maker,” with “Reorganize” as a choice.

FORD, Edward E., LOVE GUARANTEED WITH ED FORD, VHS Videotape. (Sent by Ed FORD to Greg WILLIAMS on February 29, 1992.)

HYLAND, Michael E., and William T. POWERS, “Qualitative control of human motivation: Explaining the dynamics of behavior,” unpublished paper, [2] + 48 + [1] pages. (Probably 1985, possibly 1984. With related materials.)

HYLAND, Michael E., and William T. POWERS, “Reciprocal determinism and systems analysis,” unpublished paper, [2] + 20 + [1] pages. (Probably 1983. With related materials.)

NEWSLETTER OF THE CONTROL THEORY CHAPTER, AMERICAN SOCIETY FOR CYBERNETICS, [1(1)], 2 pages. (Probably 1985.)

POWERS, Bill, “Of angels and abstractions,” unpublished paper, 2 pages.

POWERS, William T., “A cognitive control system,” unpublished paper, [5] + 18 pages. (Given by author to Greg WILLIAMS, August 1987.)

POWERS, William T., “A feedback theory of human behavior. III. The negentropy function—The process of organization,” unpublished paper, 8 pages. (Page numbers: 8415].)

POWERS, William T., “A new psychology,” unpublished paper, 18 pages. (With slip of paper showing address of editor of SATURDAY REVIEW OF SCIENCE.)

POWERS, William T., “A non-functional analysis,” unpublished paper, 2 pages. (Given by author to Greg WILLIAMS, August 1987.)

POWERS, William T., “A reply to Fowler and Turvey,” unpublished paper, 14 pages.

POWERS, William T., “A step toward an integrated science,” unpublished paper, [9] pages.

POWERS, William T., “An approach to the nature of experience,” unpublished paper, 13 pages.

POWERS, William T., “Autopoeisis [sic], recursion, and control theory,” unpublished paper, 16 pages.

POWERS, William T., CONTROL THEORY[:] A NEW LOOK AT HUMAN BEHAVIOR, 4 VHS Videotapes, Control System Group, Scottsdale, Arizona.

POWERS, William T., “Control theory and the market,” unpublished paper, 2 pages.

POWERS, William T., “CT psychology—A new view of organisms,” 74 pages. (Probably 1983.)

[POWERS, William T.,] “Cybernetics: A new science of life,” unpublished paper, [8] pages.

POWERS, William T., “Deriving closed-loop transfer functions for a behavioral model, and vice versa,” 12 pages. (Figure 1 missing.)

POWERS, William T., “Feedback concepts applied to a rat experiment,” unpublished paper, 13 + [3] pages. (Probably written ca. 1960.)

POWERS, William T., “Feedback exercises,” quiz for seminar taught by author at Northwestern University, Evanston, Illinois, 1971-1972, [2] pages. (With letter from Bill POWERS to Greg WILLIAMS dated October 8, 1989.)

[POWERS, William T.,] “Feedback interpretation of the process of concept formation,” unpublished paper, 9 pages. (With earlier draft.)

[POWERS, William T., “Feedback theory and principles of rationality,”] unpublished paper, 7 pages. (Title on tab of file folder containing paper.)

POWERS, William T., “Human agency in social cognitive theory: A comment,” unpublished paper, 6 pages. (Sent by author to Greg WILLIAMS in October 1989.)

[POWERS, William T.,] “Introduction,” unpublished paper, 48 pages. (With earlier draft.)

POWERS, William T., Letter to the editor of SCIENCE, unpublished paper, [2] pages.

POWERS, William T., “On B. E Skinner, non-scientist,” unpublished paper, [3] pages.

[POWERS, William T.,] “Operant and other conditioning,” unpublished paper, [1] page.

POWERS, William T., “Operant conditioning,” unpublished paper, 19 pages.

[POWERS, William T.,] “Perception and the laws of physics: An enquiry,” unpublished paper, 10 pages.

[POWERS, William T.,] “Signals, variables, and functions in bio-feedback systems,” unpublished paper, 3 pages. (With pages 3 and 4 only of “fb in phys & biochem.”)

POWERS, William T., “Sketch of new version of paper,” unpublished paper, [10] pages.

POWERS, William T., “The coming ascendency of control theory,” unpublished paper, [26] pages.

POWERS, William T., “The epistemology of the control system model of human behavior,” unpublished paper, 8 pages.

POWERS, William T., “The illusion of control,” unpublished paper, 15 pages.

[POWERS, William T.,] “The Journal of Unpublishable Ideas,” unpublished paper, 2 pages.

POWERS, William T., “The meaning of control theory for behavioral science,” unpublished paper, 66 pages.

[POWERS, William T.,] Untitled handwritten manuscript on control-system theory and therapy,” 66 pages.

[POWERS, William T., and Edward E. FORD,] Transcriptions of telephone conversations, paginated by topic. (Transcribed by Edward E. FORD. Some topics are dated 1986 or 1988.)

POWERS, W. T., “Feedback concepts for practical application,” unpublished paper, 9 pages.

POWERS, W. T., “Feedback theory and concept-attainment,” unpublished paper, 12 pages. [With alternate version of page 12.)

POWERS, W. T., “Some implications of feedback theory concerning behavior,” unpublished paper, 10 + [2] pages. (Two copies, Diagram 1 missing in one copy; probably 1958. With correspondence between author and Richard L. SOLOMON.)

POWERS, W. T., R. L. McFARLAND, and R. K. CLARK, “Feedback theory and psychoanalytic theory,” unpublished paper, 17 pages. (Probably written in late 1950s.)

ROBERTSON, Richard J., and William T. POWERS, INTRODUCTION TO MODERN PSYCHOLOGY[:] GENERAL PSYCHOLOGY: THE CONTROL THEORY VIEW, [Richard J. Robertson], [Chicago], [ii] + 195 pages. (Probably first published in 1985, according to Richard J. Robertson.)

[ROBERTSON, Richard J.,] INTRODUCTION TO MODERN PSYCHOLOGY[:] MANUAL OF OBJECTIVE CHAPTER EXAMS & TUTORIAL ANSWERS, paginated by chapter. (Probably 1990.)

Rubber bands.

Shirt, blue with yellow "Archivist[:] CONTROL SYSTEMS GROUP" and white CSG logo embroidery, size M. (Given by Ed FORD to Greg WILLIAMS.)

Shirt, blue with white "CONTROL SYSTEMS GROUP" and CSG logo embroidery, size L. (Given by Ed FORD to Greg WILLIAMS.)

Watch, blue CSG logo on white face. (Given by Ed FORD to Greg WILLIAMS.)

1948

DEWEY, John, "The reflex arc concept in psychology," in Wayne DENNIS, editor, READINGS IN THE HISTORY OF PSYCHOLOGY, AppletonCentury-Crofts, New York, 355-365.

1951

POWERS, William T., "A synthesis of processing techniques: Function-processing," Second Impression, Dianetics Processing and Research Foundation, Chicago, October 10, [9] pages.

POWERS, William T., and Gerald N. KNOWLTON, "Logical development of dianetics," Second Edition, Dianetics Processing & Research Foundation, Chicago, October 5, [2] + 38 pages.

1952

POWERS, William T., "Principles and techniques of self-processing," Dianetic Processing & Research Foundation, Chicago, September 10, [1] + 22 + [2] pages.

POWERS, William T., Gerald N. KNOWLTON, and David H. GREENE, "Development of awareness of present time," First Edition, Dianetic Processing & Research Foundation, Chicago, January 20, iv + 62 pages.

1955

SLACK, Charles W., "Feedback theory and the reflex arc concept," PSYCHOLOGICAL REVIEW 62(4), date?, 263-267.

1957

McFARLAND, R. L., W. T. POWERS, and R. K. CLARK, "A general feedback theory of human behavior: A prospectus," THE AMERICAN PSYCHOLOGIST 12(7), July, 462. (Abstract.)

McFARLAND, R. L., W. T. POWERS, and R. K. CLARK, "A general feedback theory of human behavior: A prospectus," unpublished paper, 20 pages. ("April 57" in pencil on first page.)

POWERS, W. T., R. L. McFARLAND, and R. K. CLARK, "A general feedback theory of human behavior," UNIVERSITY OF CHICAGO COUNSELING CENTER DISCUSSION PAPERS 3(18), 21 pages.

POWERS, W. T., R. L. McFARLAND, and R. K. CLARK, "A general feedback theory of human behavior," unpublished paper, 50 + [1] pages. (Date given as 1957 in "A preliminary report on a clinical rating scale to measure participation in group psychotherapy derived from a hierarchical feedback model" by R. L. McFARLAND, W. T. POWERS, and R. K. CLARK, 1958.)

1958

CLARK, R. K., W. T. POWERS, J. ARBIT, C. VAN BUSKIRK, and R. L. McFARLAND, "Symposium[:] Application of a general feedback theory to complex learning, concept formation and personality assessment," presented at American Psychological Association meeting, Washington, D.C., August 30, paginated by section.

McFARLAND, R. L., W. T. POWERS, and R. K. CLARK, "A preliminary report on a clinical rating scale to measure participation in group psychotherapy derived from a hierarchical feedback model," unpublished paper, 9 pages. (Dated by inference from references.)

1959

BERNATOWICZ, A. J., "Teleology in science teaching," SCIENCE 128(3336), December 5, 1402-1405.

BORING, Edwin G., William T. POWERS, Anthony STANDEN, George Gaylord SIMPSON, Perry R. STOUT, and Albert J. BERNATOWICZ, "On eschewing teleology," SCIENCE 129(3349), March 6, 608, 610, 671-672, 674-676. (Comments and rejoinder on "Teleology in science teaching" by A. J. BERNATOWICZ.)

1960

CLARK, R. K., "A general theory of human behavior from the viewpoint of physical science," paper presented to Northwestern University Society for Neuroelectrokinetics, December, 1960, 22 pages. (Stamped on front page: "Reprinted from Newsletter Research Psychology Vol. IV 1962 pp 41-61.")

MOWRER, 0. Hobart, LEARNING THEORY AND THE SYMBOLIC PROCESSES, John Wiley & Sons, New York and London, 473 pages. (Chapter 7 includes a discussion of "A General Feedback Theory of Human Behavior" by W. T. POWERS, R. L. McFARLAND, and R. K. CLARK, quoting from a letter from R. K. CLARK to 0. H. MOWRER.)

POWERS, W., "Application of feedback concepts to a rat experiment," ABSTRACTS[:] FOURTH ANNUAL MEETING OF THE BIOPHYSICAL SOCIETY, Philadelphia, February 24-26, 23-24. (Abstract.)

POWERS, W. T., R. K. CLARK, and R. L. McFARLAND, "A general feedback theory of human behavior: Part I," PERCEPTUAL AND MOTOR SKILLS 11(1), August, 71-88.

POWERS, W. T., R. K. CLARK, and R. L. McFARLAND, "A general feedback theory of human behavior: Part I," in Ludwig VON BERTALANFFY and Anatol RAPOPORT, editors, GENERAL SYSTEMS[:] YEARBOOK OF THE SOCIETY FOR GENERAL SYSTEMS RESEARCH 5, Society for General Systems Research, Ann Arbor, Michigan, 63-73.

POWERS, W. T., R. K. CLARK, and R. L. McFARLAND, "A general feedback theory of human behavior: Part II," PERCEPTUAL AND MOTOR SKILLS 11(3), December, 309-323.

POWERS, W. T., R. K. CLARK, and R. L. McFARLAND, "A general feedback theory of human behavior: Part II," in Ludwig VON BERTALANFFY and Anatol RAPOPORT, editors, GENERAL SYSTEMS[:] YEARBOOK OF THE SOCIETY FOR GENERAL SYSTEMS RESEARCH 5, Society for General Systems Research, Ann Arbor, Michigan, 75-83.

SYMPOSIUM IN NEURAL MECHANISMS, INFORMATION THEORY AND BEHAVIOR, March 10-11, unpaginated. (Program lists "A General Feedback Theory of Human Behavior" by Robert L. McFARLAND, Robert K. CLARK, and William POWERS, March 11.)

1961

CLARK, R. K., Robert L. McFARLAND, and Richard J. ROBERTSON, "Symposium: New diagnostic and therapeutic procedures derived from general feedback theory," presented at American Psychological Association meeting, New York, September, paginated by section.

1962

CLARK, Robert K., and Robert L. McFARLAND, "How can the scientist help the psychotherapist?" unpublished paper presented at American Academy of Psychotherapists meeting, location?, October, 12 pages.

1963

CLARK, Robert K., and Robert L. McFARLAND, "Systems concept of stimulus," PERCEPTUAL AND MOTOR SKILLS 17(1), date?, 99-102.

GRODINS, Fred S., CONTROL THEORY AND BIOLOGICAL SYSTEMS, Columbia University Press, New York and London, xii + 205 pages.

1964

CLARK, R. K., R. L. McFARLAND, and M. BASSAN, "Integrated data collecting and processing systems in psychophysiology," ANNALS OF THE NEW YORK ACADEMY OF SCIENCES 115(number?), July 31, 905914.

1966

POWERS, W. T., R. K. CLARK, and R. I. [sic] McFARLAND, "A general feedback theory of human behavior," in Alfred G. SMITH, editor, COMMUNICATION AND CULTURE[:] READINGS IN THE CODES OF HUMAN INTERACTION, Holt, Rinehart and Winston, New York, et al., 1966, 333-343.

1969

POWERS, William T., “A feedback theory approach to cancer,” unpublished paper, March 24, 19 pages.

POWERS, William T., “A new look at feedback,” unpublished paper, 16 + [1] pages. (With related correspondence among author and periodical editors.)

[POWERS, William T.,] “The fundamental ambiguity of perception,” unpublished paper presented at Foundations of Science seminar, [Northwestern University, Evanston, Illinois,] October 30, 11 pages. (With seminar program and mailing list, [2] pages.)

1970

ALADDIN INDUSTRIES, INC., Advertisement for “Trippples” game, SCIENTIFIC AMERICAN volume?(number?), date?, 130. (“Trippples,” marketed by Aladdin Industries of Nashville, Tennessee, was invented by William T. POWERS.)

1971

POWERS, William T., “The Wiener feedback model—A strategic error,” unpublished paper, August 13, 21 + [9] pages. (With earlier draft and correspondence between William T. POWERS and BEHAVIORAL SCIENCE editors.)

POWERS, William T., “A feedback model for behavior: Application to a rat experiment,” BEHAVIORAL SCIENCE 16(6), November, 558-563. (With brief biography of William T. POWERS on page 580.)

TAYLOR, M. M., “Information for action: Perceptual principles related to remote sensing,” transcription of early (October 29, 1971) version of paper published in PROCEEDINGS OF THE FIRST CANADIAN SYMPOSIUM ON REMOTE SENSING, February 1992, 14 pages. (Distributed by author at CSG meeting, Durango, Colorado, July-August, 1993.)

1972

POWERS, William T., BEHAVIOR: THE CONTROL OF PERCEPTION, [William T. Powers], [Northbrook, Illinois], paginated by chapter.

1973

POWERS, William T., BEHAVIOR: THE CONTROL OF PERCEPTION, Aldine Publishing Company, Chicago, xiv + 296 pages.

POWERS, William T., “Feedback: Beyond behaviorism,” SCIENCE 179(4071), January 26, 351-356.

BAUM, William M., Hayne W. REESE, and William T. POWERS, “Behaviorism and feedback control,” SCIENCE 181(4105), September 21, 1114, 1116, 1118-1120.

1974

BOHANNAN, Paul, William POWERS, and Mark SCHOEPFLE, “Systems conflict in the learning alliance,” in Lindley J. STILES, editor, THEORIES FOR TEACHING, Dodd, Mead & Company, New York, 76-96.

HOLDEN, Constance, “Butner: Experimental U.S. prison holds promise, stirs trepidation,” SCIENCE 185(4149), August 2, 423-426.

LINDSAY, Robert K., “Behaving man,” SCIENCE 184(number?), April 26, 455-457. (Review of BEHAVIOR: THE CONTROL OF PERCEPTION by William T. POWERS, 1973.)

POWERS, William T., “Some cybernetics and some psychology,” ASC FORUM 6(4), Winter, 4-9. (With brief biography of William T. POWERS on page 31.)

POWERS, William T., “Applied epistemology,” in Charles D. SMOCK and Ernst VON GLASERSFELD, editors, EPISTEMOLOGY AND EDUCATION[:] THE IMPLICATIONS OF RADICAL CONSTRUCTIVISM FOR KNOWLEDGE ACQUISITION, Mathemagenic Activities Program Follow Through Research Report (14), University of Georgia, Athens, 84-98.

POWERS, William T., and B. F. SKINNER, “Behavior reinforcement,” SCIENCE 186(4166), November 29, 782. (Comment on “Behavior modification” by B. F SKINNER and rejoinder.)

SKINNER, B. E, “Behavior modification,” SCIENCE 186(4154), September 6, 813. (Comment on “Butner: Experimental U.S. prison holds promise, stirs trepidation” by Constance HOLDEN.)

1975

POWERS, William T., Review of THE LOGIC OF SOCIAL SYSTEMS by Alfred Kuhn, CONTEMPORARY SOCIOLOGY 4(2), March, 92-94.

1976

HERSHBERGER, Wayne A., "Afference copy, the closed-loop analogue of von Hoist's efference copy," unpublished paper presented at American Psychological Association meeting, Washington, D.C., September, 8 + [5] pages. (Distributed by author at CSG meeting, Haimowoods, Wisconsin, October 1987.)

KATZ, Stuart, "An analysis of the theory of knowledge in Power's [sic] model of brain," ASC CYBERNETICS FORUM 8(1 & 2), Spring/Summer, 41-49.

POWERS, William T., "A reply to Katz' analysis," ASC CYBERNETICS FORUM 8(3 & 4), Fall/Winter, 143-146.

POWERS, William T., "Control-system theory and performance objectives," JOURNAL OF PSYCHOLINGUISTIC RESEARCH 5(3), July, 285297.

POWERS, William T., "The cybernetic revolution in psychology," ASC CYBERNETICS FORUM 8(3 & 4), Fall/Winter, 72-86.

[POWERS, William T.,] Draft of unpublished book, paginated by chapter.

1977

GIBBONS, Hugh, A THEORY OF DEMOCRACY, Franklin Pierce Law Center, Concord, New Hampshire, x + 263 pages. (With letter from author to William POWERS dated September 7, 1977.)

POWERS, William T., "Rückkopplungsprinzipien in der Organisation von Verhalten," in Hans Zeier, editor, PAWLOW UND DIE FOLGEN[:] VON DER KLASSISCHEN KONDITIONIERUNG BIS ZUR VERHALTENSTHERAPIE, DIE PSYCHOLOGIE DES 20. JAHRHUNDERTS 4, Kindler Verlag AG, Zurich, 573-613. (Also copies of English manuscript, prior to editorial changes, and galley proof.)

1978

POWERS, William T., "Quantitative analysis of purposive systems: Some spadework at the foundations of scientific psychology," PSYCHOLOGY REVIEW 85(5), September, 417-435.

1979

POWERS, William T., "A cybernetic model for research in human development," in Mark N. OZER, editor, A CYBERNETIC APPROACH TO THE ASSESSMENT OF CHILDREN: TOWARD A MORE HUMANE USE OF HUMAN BEINGS, Westview Press, Boulder, Colorado, 11-66.

POWERS, William T., "Cause/effect metaphors versus control theory," THE BEHAVIORAL AND BRAIN SCIENCES 2(1), March, 115.

POWERS, William T., "Control theory for sociology," unpublished paper presented at conference on Experimental Advances in the Social Psychology of G. H. Mead, Department of Sociology, Kent State University, 1979, 34 + [1] pages. (With a differently formatted version, 17 pages.)

POWERS, William T., "Cybernetics and the assessment of children," ASC CYBERNETICS FORUM 9(1 & 2), Spring/Summer, 57-65.

POWERS, William T., "Degrees of freedom in social interactions," in Klaus Krippendorff, editor, COMMUNICATION AND CONTROL IN SOCIETY, Gordon and Breach Science Publishers, New York, London, and Paris, 267-278.

POWERS, William T., "The nature of robots[:] Part 1: Defining behavior," BYTE 4(6), June, 132, 134, 136, 138, 140-141, 144.

POWERS, William T., "The nature of robots[:] Part 2: Simulated control system," BYTE 4(7), July, 134-136, 138, 140, 142, 144, 146, 148-150, 152.

POWERS, William T., "The nature of robots[:] Part 3: A closer look at human behavior," BYTE 4(8), August, 94-96, 98, 100, 102-104, 106-108, 110112, 114, 116.

POWERS, William T., "The nature of robots[:] Part 4: Looking for controlled variables," BYTE 4(8), September, 96, 98-102, 104, 106-110, 112.

RICHARDS, John, and Ernst VON GLASERSFELD, “The control of perception and the construction of reality[:] Epistemological aspects of the feedback-control system,” DIALECTICA 33(1), date?, 37-58.

1980

BOURBON, W. Tom, “Anticipatory regulation: A raincoat does not feedforward make,” THE BEHAVIORAL AND BRAIN SCIENCES 3(number?), date?, 465-466.

EDWARDS, Lauren, A MULTIDIMENSIONAL STUDY OF AROUSAL CHANGES DURING GAME-PLAYING, M.A. Thesis, Stephen F. Austin State University, Nachadoches, Texas, May, iii + 63 pages. (Thesis Director: Tom BOURBON.)

MARKEN, Richard, “The cause of control movements in a tracking task,” PERCEPTUAL AND MOTOR SKILLS 51(3), December, 755-758.

POWERS, William T., “A systems approach to consciousness,” in Julian M. DAVIDSON and Richard J. DAVIDSON, editors, THE PSYCHOLOGY OF CONSCIOUSNESS, Plenum Press, New York and London, 217-242.

POWERS, William T., “Control-theory psychology and social organizations: Background for a theory of the corruption of social indicators when used for social decision making,” unpublished paper, March 1980, 34 pages. (With other versions.)

POWERS, William T., “Pylyshyn and perception,” THE BEHAVIORAL AND BRAIN SCIENCES 3(1), March, 148-149. (Comment on “Computation and cognition: Issues in the foundations of cognitive science” by Zenon W. PYLYSHYN.)

PYLYSHYN, Zenon W., “Computation and cognition: Issues in the foundations of cognitive science,” THE BEHAVIORAL AND BRAIN SCIENCES 3(1), March, 111-169.

1981

CARVER, Charles S., and Michael F. SCHEIER, ATTENTION AND SELF-REGULATION: A CONTROL-THEORY APPROACH TO HUMAN BEHAVIOR, Springer-Verlag, New York, Heidelberg, and Berlin, xviii + 403 pages.

GLASSER, William, STATIONS OF THE MIND: NEW DIRECTIONS FOR REALITY THERAPY, Harper & Row, Publishers, New York, et al., xviii + 288 pages. (On dust jacket cover: “In consultation with William T. POWERS.”)

MARKEN, Richard S., “Closed-loop behavior: Human performance as control of input,” unpublished paper presented at Western Psychological Association meeting, Los Angeles, month?, [12] pages. (Filed with MIND READINGS by Richard S. MARKEN, 1992.)

PETRIE, Hugh G., THE DILEMMA OF ENQUIRY AND LEARNING, University of Chicago Press, Chicago and London, ? pages. (Entire book not in Archive.)

POWERS, William T., “Maximization, or control?” THE BEHAVIORAL AND BRAIN SCIENCES 4(3), September, 400-401.

1982

CARVER, Charles S., and Michael F. SCHEIER, “Control theory: A useful conceptual framework for personality-social, clinical, and health psychology,” PSYCHOLOGICAL BULLETIN 92(1), date?, 111-135.

JEFFREY, D. Balfour, and Laurence H. BERGER, “A self-environmental systems model and its implications for behavior change,” in Kirk R. BLANKSTEIN and Janet POLIVY, editors, SELF-CONTROL AND SELF-MODIFICATION OF EMOTIONAL BEHAVIOR (Advances in the Study of Communication and Affect, Volume 7), Plenum Press, New York and London, 29-69.

MARKEN, Richard, “Intentional and accidental behavior: A control theory analysis,” PSYCHOLOGICAL REPORTS 50(2), April, 647-650.

McCORD, David M., “Control theory: Overview and empirical demonstration,” ALABAMA STUDIES IN PSYCHOLOGY 1(1), Fall, 7-16.

MURPHY, Philip Edward, AN ANALYSIS OF A TRACKING TASK FROM THE PERSPECTIVE OF CONTROL THEORY, M.A. Thesis, Stephen F Austin State University, Nachadoches, Texas, August, 146 pages. (Thesis Director: W. Tom BOURBON.)

PLOOIJ, Frans X., THE BEHAVIORAL DEVELOPMENT OF FREE-LIVING CHIMPANZEE BABIES AND INFANTS, Ablex Publishing Corp., Norwood, New Jersey, xvi + 207 pages.

1983

FORD, Edward E., CHOOSING TO LOVE: A NEW WAY TO RESPOND, Winston Press, Minneapolis, [xii] + 149 pages.

MARKEN, Richard, "'Mind reading': A look at changing intentions," PSYCHOLOGICAL REPORTS 53(1), August, 267-270.

McCORD, David M., CONTROL THEORY: A NEW LOOK AT PURPOSIVE BEHAVIOR, Ph.D. Dissertation, University of Alabama, ix + 166 pages.

POWERS, William T., "A bucket of beans: Applications of control theory to the understanding of human nature and the design of social systems," unpublished paper, October 7, 11 pages.

ROBERTSON, Richard J., "Theoretical and clinical views of reorganization in human systems," unpublished paper presented at ASC meeting, San Jose?, California, 10 pages.

1984

GLASSER, William, TAKE EFFECTIVE CONTROL OF YOUR LIFE, Harper & Row, Publishers, New York et al., ? pages. (Entire book not in Archive. Contents identical to CONTROL THEORY[:] A NEW EXPLANATION OF HOW WE CONTROL OUR LIVES by William GLASSER, 1985.)

HENDRICKSON, Jeffrey John, THE SACCADIC EYE MOVEMENT CONTROL SYSTEM, M.A. Thesis, Stephen F. Austin State University, Nachadoches, Texas, May, viii + 106 pages. (Thesis Director: Walter T. BOURBON.)

[POWERS, William T.,] "Essay on control of social variables," unpublished paper, [4] pages.

POWERS, William T., "Interactionism and control theory," in Joseph R. ROYCE and Leendert P. MOS, editors, ANNALS OF THEORETICAL PSYCHOLOGY 2, Plenum Press, New York and London, 355-358.

ROBERTSON, R. J., "Control theory," in Raymond J. CORSINI, editor, ENCYCLOPEDIA OF PSYCHOLOGY, Volume 1, John Wiley & Sons, New York, et al., 288-291. (With "Bibliography" from Volume 4.)

SOLDANI, James C., and Edward E. FORD, MONEY ISN'T ENOUGH[:] MANAGING PEOPLE EFFECTIVELY USING CONTROL SYSTEM THEORY, revised edition, Ford-Soldani, Inc., Scottsdale, Arizona, paginated by chapter.

1985

GLASSER, William, CONTROL THEORY[:] A NEW EXPLANATION OF HOW WE CONTROL OUR LIVES, Harper & Row, Publishers, New York et al., xvi + 240 pages.

GOSSEN, Diane Chelsom, "Taking effective control of teaching effective control," unpublished paper, April, 17 pages. (Distributed by author at CSG meeting, Haimowoods, Wisconsin, October 1987.)

MARKEN, Richard S., "Selection of consequences: Adaptive behavior from random reinforcement," PSYCHOLOGICAL REPORTS 56(number?), date?, 379-383.

NEWSLETTER OF THE CONTROL THEORY CHAPTER, AMERICAN SOCIETY FOR CYBERNETICS, 1(2), June, 3 pages.

POWERS, William T., "The arm," unpublished paper, May 29, [7] pages.

VARELA, Francisco, "Steps to a cybernetics of autonomy," in Robert Trappl, editor, POWER, AUTONOMY, UTOPIA: NEW APPROACHES TOWARD COMPLEX SYSTEMS, Plenum Press, New York and London, 117-122.

VON GLASERSFELD, Ernst, "Steps in the construction of 'others' and 'reality': A study in self-regulation," in Robert TRAPPL, editor, POWER, AUTONOMY, UTOPIA: NEW APPROACHES TOWARD COMPLEX SYSTEMS, Plenum Press, New York and London, 107-116.

1986

BOURBON, W. Tom, "A case of different intentions concerning intentionality," unpublished paper distributed at CSG meeting, Haimowoods, Wisconsin, August, [4] pages.

BOURBON, W. Tom, "Auditory event-related potentials: A comparison of self-initiated and automated stimuli," unpublished paper distributed at CSG meeting, Haimowoods, Wisconsin, August, [6] pages.

CENTER FOR SYSTEMS RESEARCH WORKING PAPER (86-1), February, University of Alberta, Edmonton, 19 pages.

CENTER FOR SYSTEMS RESEARCH WORKING PAPER (86-4), November, University of Alberta, Edmonton, 17 pages. ("Interaction: The control of perception: Commentary on a paper by J. M. White" by William T. POWERS, "A biographical note: Dr. [sic] William T. Powers," "Some early and pertinent Publications" [by W. T. POWERS], "Interaction: The control of perception: Reply to comments by W. T. Powers" by James M. WHITE, 13-15, "A biographical note: Dr. James (Jim) Martin White," "Some recent publications" [by J. M. WHITE].)

CONTROL SYSTEM GROUP NEWSLETTER, October 15, 4 pages.

"Control Systems Group '86 [meeting]," 15 audiocassettes. (Recorded by Greg WILLIAMS.)

FEEDBACK: THE CONTROL SYSTEM GROUP NEWSLETTER, March 5, 3 pages.

FEEDBACK: THE CONTROL SYSTEM GROUP NEWSLETTER, May 1, 2 pages.

FORD, Edward E., "Control system psychology," 1 page.

FORD, Edward E., "COUNSELING," card printed on both sides. FORD, Edward E., "LOVE," card printed on both sides.

FORD, Edward E., "The role of feelings in control system psychology," unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, August, 3 pages.

FORD, Edward E., "The role of feelings in control systems theory," CONTINUING THE CONVERSATION: A NEWSLETTER ON THE IDEAS OF GREGORY BATESON (7), Winter, 11-13.

GLASSER, William, "The basic concepts of reality therapy," Institute For Reality Therapy, Canoga Park, California, [4] pages.

GOLDFARB, Lawrence William, "A moving conversation," CONTINUING THE CONVERSATION: A NEWSLETTER ON THE IDEAS OF GREGORY BATESON (7), Winter, 13-14.

GOLDSTEIN, David M., "Control systems in a clinical setting," CONTINUING THE CONVERSATION: A NEWSLETTER ON THE IDEAS OF GREGORY BATESON (7), Winter, 10-11.

GOLDSTEIN, David M., "Q methodology, control system theory and psychotherapy," unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, August, 4 + [4] pages. ("August 20, 1986.")

HERSHBERGER, Wayne A., "An approach through the looking-glass," ANIMAL LEARNING & BEHAVIOR 14(4), date?, 443-451. (Also have preprint given at CSG meeting, Haimowoods, Wisconsin, August 1986.)

HERSHBERGER, Wayne, Memorandum to Control Theory Group regarding "Some behaviors are neither elicited nor emitted," unpublished paper distributed at CSG meeting, Haimowoods, Wisconsin, August, dated August 13, 1986, [1] + 3 pages. (Includes draft of "Some behaviors are neither elicited nor emitted" by Wayne Hershberger, [1985].)

HERSHBERGER, Wayne, "Saccadic eye movements and the perception of visual direction," unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, August, 43 + [4] pages.

JORDAN, Jerome S., "The behavioral illusion," unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, August, [1] + 22 pages.

MARKEN, Richard S., "Human factors and human nature: Is psychological theory really necessary?" unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, August, 5 pages.

MARKEN, Richard S., "Human factors and human nature: Is psychological theory really necessary?" HUMAN FACTORS SOCIETY BULLETIN 29(9), September, 1-2.

MARKEN, Richard S., "Perceptual organization of behavior: A hierarchical control model of coordinated action," JOURNAL OF EXPERIMENTAL PSYCHOLOGY: HUMAN PERCEPTION AND PERFORMANCE 12(3), August, 267-276.

MARKEN, Richard S., "The nature of behavior: Control as fact and theory," unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, August, [1] + 26 pages. (On cover page: "1 August, 1986.")

McPHAIL, Clark, and Ronald T. WOHLSTEIN, “Collective locomotion as collective behavior,” AMERICAN SOCIOLOGICAL REVIEW 51(4), August, 447-463.

MEUTSCH, Dietrich, “Mental models in literary discourse: Towards the integration of linguistic and psychological levels of description,” POETICS 15(3), June, 307-331.

PAVLOSKI, Raymond, “Control systems and cardiovascular reactivity: Toward an understanding of cardiac-somatic interactions,” CONTINUING THE CONVERSATION: A NEWSLETTER ON THE IDEAS OF GREGORY BATESON (7), Winter, 8-10.

PAVLOSKI, Raymond P., “A control theory approach to person-environment transactions characterized by cardiovascular reactivity,” unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, August, 59 + [2] pages.

POWERS, William T., “An agenda for the Control Theory Group,” unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, August, 3 pages.

POWERS, W. T., “Control theory, constructivism, and autopoeisis [sic],” unpublished paper, August, 11 pages.

POWERS, William T., “Intentionality: No mystery,” THE BEHAVIORAL AND BRAIN SCIENCES 9(1), March, 152-153.

POWERS, W. T., “On purpose,” CONTINUING THE CONVERSATION: A NEWSLETTER ON THE IDEAS OF GREGORY BATESON (7), Winter, 1-3.

POWERS, William T., “Purposive behavior (V3),” unpublished paper, 18 + [7] pages.

POWERS, W. T., “Purposive behavior,” unpublished paper, 7 + *[6]* pages. (With letter from author to Control Theory Group members dated April 20, 1986.)

POWERS, W. T., “Principles of biological autonomy: A limited and biased commentary illustrating some contrasts with control theory,” unpublished paper distributed by author at CSG meeting, Haimowoods, Wisconsin, August, 14 pages.

POWERS, William, and Ernst VON GLASERSFELD, “A control conversation,” CONTINUING THE CONVERSATION: A NEWSLETTER ON THE IDEAS OF GREGORY BATESON (7), Winter, 3-4.

RICHARDS, Larry, “Impressions of recent conferences,” CONTINUING THE CONVERSATION: A NEWSLETTER ON THE IDEAS OF GREGORY BATESON (6), Fall, 7-8.

RICHARDS, Larry, and Bill POWERS, “Letters,” CONTINUING THE CONVERSATION: A NEWSLETTER ON THE IDEAS OF GREGORY BATESON (7), Winter, 4-6.

RICHARDS, Laurence D., “Regarding Bateson and Powers,” CONTINUING THE CONVERSATION: A NEWSLETTER ON THE IDEAS OF GREGORY BATESON (7), Winter, 19-20.

ROBERTSON, Richard J., “CST and self image,” CONTINUING THE CONVERSATION: A NEWSLETTER ON THE IDEAS OF GREGORY BATESON (7), Winter, 11.

ROBERTSON, Richard J., “Research program — Human Systems Laboratory, Northeastern Illinois University,” unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, August, 22 pages.

ROBERTSON, Richard J., SETTING YOUR OWN EGO-STAT: HOW TO CONTROL YOUR LIFE INSTEAD OF YOURSELF, Richard J. Robertson, Chicago, [iv] + 198 + 6 pages.

ROBERTSON, Richard J., and David M. GOLDSTEIN, “The self as control system,” unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, August, 24 + 2 pages.

ROSS, Douglas A., “Control systems and psychological applications,” CONTINUING THE CONVERSATION: A NEWSLETTER ON THE IDEAS OF GREGORY BATESON (7), Winter, 6-8.

RUNKEL, Philip J., GENERALIZING, unpublished manuscript, July, 113 pages.

RUNKEL, Philip J., INSIDE AND OUTSIDE, unpublished manuscript, June, 394 pages.

WILLIAMS, Greg, “Control theory & the I-it boundary,” unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, August, [2] pages. (“8/18/86.”)

WILLIAMS, Gregory, compiler, “A basic bibliography of control systems theory,” CONTINUING THE CONVERSATION: A NEWSLETTER ON THE IDEAS OF GREGORY BATESON (7), Winter, 18-19.

[WILLIAMS, Gregory,] “‘Control systems theory,⁻ CONTINUING THE CONVERSATION: A NEWSLETTER ON THE IDEAS OF GREGORY BATESON (7), Winter, 1.

[WILLIAMS, Gregory,] “The Control Systems Group,” CONTINUING THE CONVERSATION: A NEWSLETTER ON THE IDEAS OF GREGORY BATESON (7), Winter, 20.

WILLIAMS, W. D., “Control theory and the Smithian economics,” CONTINUING THE CONVERSATION: A NEWSLETTER ON THE IDEAS OF GREGORY BATESON (7), Winter, 14-17.

[WILLIAMS, William D.,] “The Giffen effect, a control theory analysis: A comparison of two choice theoretic constructs of economic behavior,” unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, August, [29] pages.

1987

AITKEN, Michael Luke, “Re: CC number 7,” CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (9), Summer, 8-9.

BOURBON, Tom, “Ugliness: A reply to Yocum, White, and others,” CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (11), Winter, 14-16.

BOURBON, Tom, Bill POWERS, Mary POWERS, Bill WILLIAMS, and Greg WILLIAMS, “CST oral history,” August 20, 2 audiocassettes. (With related materials.)

BOURBON, W. Tom, “A case of different intentions concerning intentionality,” BEHAVIORAL AND BRAIN SCIENCES 10(number?), date?, 755-756. (Commentary on “Intentionality and information processing: An alternative model for cognitive science” by Kenneth M. SAYRE.)

BOURBON, W. Tom, “Some thoughts on the occult and the supernatural in cybernetics, behaviorism, and cognitive science,” CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (8), Spring, 4-5.

CAMPBELL, Thomas M., AN EXAMINATION OF THE INFLUENCE OF CONTROL THEORY AND CONCURRENT ORGANIZATIONAL CHANGE ON THE REALITY THERAPY MOVEMENT, Ed.D. Dissertation, University of Tulsa, Oklahoma, vii + 67 pages. (Chairperson: Dr. Richard L. HALL.)

CONTROL SYSTEM GROUP NEWSLETTER, July 24, 4 pages.

“Control Systems Group ‘87 [meeting],” 17 audiocassettes. (Recorded by Greg WILLIAMS.)

DELPRATO, Dennis J., “A sketch of three versions of control,” unpublished paper presented at CSG Meeting, Haimowoods, Wisconsin, October, 5 pages.

DELPRATO, Dennis J., “Control system theory and the evolution of scientific thinking,” unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, October, 12 pages.

FENNELL, Geraldine, “Two projects: Representing action, planning ‘control,⁻ CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (10), Fall, 16-17.

FORD, Donald H., HUMANS AS SELF-CONSTRUCTING LIVING SYSTEMS: A DEVELOPMENTAL PERSPECTIVE ON BEHAVIOR AND PERSONALITY, Lawrence Erlbaum Associates, Hillsdale, New Jersey, xiv + 787 pages. (Entire book not in Archive.)

FORD, Edward E., “Control system psychology,” 1 page.

FORD, Edward E., LOVE GUARANTEED[:] A BETTER MARRIAGE IN EIGHT WEEKS, Harper & Row, Publishers, San Francisco, et al., xx + 124 pages.

FORD, Martin E., and Donald H. FORD, editors, HUMANS AS SELF-CONSTRUCTING LIVING SYSTEMS: PUTTING THE FRAMEWORK TO WORK, Lawrence Erlbaum Associates, Hillsdale, New Jersey, xiv + 409 pages. (Entire book not in Archive.)

GOLDSTEIN, David M., "Q-methodology and control systems theory: Applied to individual therapy with Tom D.," unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, October, 8 pages.

GOLDSTEIN, David M., Bill POWERS, and David SAUNDERS, "Transfer function analysis of pursuit tracking in adults," unpublished paper, October 3, 5 pages.

GOOD, E. Perry, IN PURSUIT OF HAPPINESS[:] KNOWING WHAT YOU WANT[,] GETTING WHAT YOU NEED, New View Publications, Chapel Hill, North Carolina, kJ + 199 pages. (With order form and undated note from author to Greg WILLIAMS.)

GOUGH, Pauline B., "The key to improving schools: An interview with William Glasser," PHI DELTA KAPPAN 68(9), May, 656-662.

HERSHBERGER, Wayne, "Saccadic eye movements and the perception of visual direction," PERCEPTION & PSYCHOPHYSICS 41(1), date?, 3544.

HERSHBERGER, Wayne A., "Of course there can be an empirical science of volitional action," unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, October, 9 pages.

HERSHBERGER, Wayne A., "Psychology as a conative science," unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, October, 9 pages.

HERSHBERGER, Wayne A., "Some behaviors are neither elicited nor emitted," AMERICAN PSYCHOLOGIST volume?(number?), June, 605606.

HOLLENBECK, John R., and Charles R. WILLIAMS, "Goal importance, self-focus, and the goal-setting process," JOURNAL OF APPLIED PSYCHOLOGY 72(2), date?, 204-211.

HYLAND, Michael, "Some heretical thoughts about control theory," CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (9), Summer, 5-6.

HYLAND, Michael E., "Control theory interpretation of psychological mechanisms of depression: Comparison and integration of several theories," PSYCHOLOGICAL BULLETIN 102(1), date?, 109-121.

LORD, Robert G., and Paul J. HANGES, "A control system model of organizational motivation: Theoretical development and applied implications," BEHAVIORAL SCIENCE 32(3), July, 161-178.

MARKEN, Richard S., "The nature of behavior: Control as fact and theory," unpublished paper, October 6, 25 pages.

MARKEN, Richard S., and William T. POWERS, "Random-walk chemotaxis: A comparison of reinforcement and control-theoretic models of operant behavior," unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, October, 14 + [4] pages.

MARKEN, Richard S., and William T. POWERS, "Hierarchical control in human performance," unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, October, 23 + [7] pages.

MATSUI, Tamao, Takashi KAKUYAMA, and Mary Lou Uy ONGLATCO, "Effects of goals and feedback on performance in groups," JOURNAL OF APPLIED PSYCHOLOGY 72(3), date?, 407-415.

McPHAIL, Clark, "Social behavior in public places: From clusters to arcs to rings," unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, October, 35 pages. (Also handout which accompanied presentation, [11] pages.)

MILLER, Russell, "Farce and fear in Scientology's private navy," [LONDON] SUNDAY TIMES, November 15, 51-52. (About L. Ron HUBBARD.)

MILLER, Russell, "Messiah at the manor," [LONDON] SUNDAY TIMES, November 8, pages? (About L. Ron HUBBARD.)

MILLER, Russell, "Science fictions," [LONDON] SUNDAY TIMES, November 1, pages? (About L. Ron HUBBARD.)

NEILL, John R., "LOVE GUARANTEED or 'the job loving,'" CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (9), Summer, 7-8.

PITMAN, Roger K., "Behavioral output CAN be controlled (in a sense)," CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (8), Spring, 2-3.

PITMAN, Roger K., "Who (or what?) is in control here?" CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (9), Summer, 2.

[POWERS,] Bill, "On violence and aggression," unpublished paper, July 21, [2] pages. (Part of letter to Frans PLOOIJ.)

POWERS, Bill, and Ranulph GLANVILLE, "87 ASC conference (Urbana)," 2 audiocassettes.

POWERS, Mary, "'Half a loaf?'?" CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (9), Summer, 2-3.

POWERS, Mary A., "To Michael Yocum:" CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (10), Fall, 15.

POWERS, William T., "A cognitive control system," unpublished paper, [20] pages.

POWERS, William T., "A non-functional analysis[:] Commentary on 'A framework for the functional analysis of behavior' by Houston, A. I. and McNamara, J. M.," unpublished paper distributed by author at CSG meeting, Haimowoods, Wisconsin, October, 3 pages.

POWERS, William T., "Control theory and cybernetics," CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (11), Winter, 13-14. (With draft versions.)

POWERS, William T., "Control theory: The road to utopia," unpublished paper presented at ASC meeting, Urbana, Illinois, December, 5 pages.

POWERS, William T., "From the constructed viewpoint of a control theorist," unpublished paper, 4 pages.

POWERS, William T., "Is control theory just another point of view?" CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (9), Summer, 1-2.

POWERS, William T., "Maybe complaining will help," unpublished paper written for, but not presented at, ASC meeting, Urbana, Illinois, December, 3 pages.

POWERS, [William T.,] "Safari Resort," 4 audiocassettes.

RICHARDS, Larry, "Moving from linear feedback circuits to topological recursion in social systems," unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, October, 14 pages.

RICHARDS, Larry, "Reflections on the St. Gallen Conference," CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (9), Summer, 9-10.

RICHARDS, Laurence D., "Engineering cybernetics and cybernetic engineering," in 1987 ASEE ANNUAL CONFERENCE PROCEEDINGS, 5 pages.

RICHARDS, Laurence D., "Future directions in expert systems with reference to engineering management," unpublished paper distributed by author at CSG meeting, Haimowoods, Wisconsin, October, 11 pages.

[ROBERTSON, Dick, and Greg WILLIAMS,] "Dick Robertson's recollections on CST history (recorded on 10-10-87 near Kenosha, WI), 2 audiocassettes.

ROBERTSON, Richard J., David M. GOLDSTEIN, Michael MERMEL, and Melanie MUSGRAVE, "Testing the self as a control system: Theoretical and methodological issues," unpublished paper presented by first two authors at CSG meeting, Haimowoods, Wisconsin, October, 17 pages.

RUNKEL, Philip J., "Dear Michael Yocum," CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (9), Summer, 3.

RUNKEL, Philip J., "Feedback," CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (10), Fall, 17.

SAYRE, Kenneth M., "Various Senses of 'Intentional System,'" BEHAVIORAL AND BRAIN SCIENCES 10(number?), date?, 760-765. (Includes reply to "A case of different intentions concerning intentionality" by W. Tom BOURBON.)

[TUCKER, Charles,] Handout accompanying presentation on control theory analyses at CSG meeting, Haimowoods, Wisconsin, October, [5] pages.

[TUCKER, Charles,) "Milgram's studies of obedience," handout accompanying presentation at CSG meeting, Haimowoods, Wisconsin, October, 5 pages + [2] loose sheets.

VALLACHER, Robin R., and Daniel M. WEGNER, "What do people think they're doing? Action identification and human behavior," PSYCHOLOGICAL REVIEW 94(1), date?, 3-15.

VAN DE RIJT-PLOOIJ, Hedwig H. C., and Frans X. PLOOIJ, "Growing independence, conflict and learning in mother-infant relations in free-ranging chimpanzees," BEHAVIOUR 101(1-3), May, 1-86.

VARELA, Francisco J., "Perception and the origin of cognition: A cartography of current ideas," unpublished paper, 34 pages.

WEATHERS, Tom, Jr., "A larger view of control theory: In pursuit of the ULOM," CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (9), Summer, 6-7.

WHITE, Dan, "Bateson's cybernetics of liberation and the bureaucracy of the sacred," CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (10), Fall, 11-15.

[WILLIAMS, Gregory,] "Control Systems Group 1987 meeting tapes," CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (11), Winter, 18.

[WILLIAMS, Gregory,] "CST videotapes and workshop," CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (8), Spring, 6.

[WILLIAMS, Gregory,] "Erratum," CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (10), Fall, 15.

[WILLIAMS, Gregory,] "From the editor," CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (9), Summer, 1.

WILLIAMS, W. D., "A further note on the Smithian economics," CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (11), Winter, 16.

[WILLIAMS, William D.,] "Erratum," CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (8), Spring, 8.

YOCUM, Michael, "Half a Loaf?" CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (8), Spring, 3.

YOCUM, Michael, "In lieu of a reply to the Powerses," CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (9), Summer, 4-5.

YOCUM, Michael, "Reply to Philip Runkel," CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (9), Summer, 3-4.

1988

BASCH, Michael Franz, UNDERSTANDING PSYCHOTHERAPY[:] THE SCIENCE BEHIND THE ART, Basic Books, Inc., Publishers, New York, xviii + 329 pages.

BEER, Stafford, "The utility of cybernetics," CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (12), Spring, 2-3.

BOURBON, W. Thomas, "Modeling control by dyads," unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, September-October, [10] pages.

BOURBON, W. Thomas, and William T. POWERS, "Purposive Behavior: A tutorial with data," unpublished paper, 25 + [10] pages.

BOURBON, W. Thomas, and William T. POWERS, "The experimental analysis of control behavior," unpublished paper distributed by principal author at CSG meeting, Haimowoods, Wisconsin, September-October, 14 + [4] pages. (Also other draft versions.)

CHELSOM Consultants, "Control theory workshops[:] 1988-1989," [2] pages. (Distributed at CSG meeting, Haimowoods, Wisconsin, September-October, by Diane GOSSEN.)

CHONG, Eric, COOPERATION INTERPRETED IN TERMS OF CONTROL SYSTEM THEORY, M.A. Thesis, Stephen E Austin State University, August, vi + 141 pages. (Thesis Director: W. T. BOURBON.)

CONTROL SYSTEM GROUP NEWSLETTER, [February], 3 pages.

CONTROL SYSTEM GROUP NEWSLETTER, March 3, 4 pages.

CONTROL SYSTEMS GROUP [NEWSLETTER], [June], [4] pages.

CONTROL SYSTEM GROUP NEWSLETTER, [November], 10 pages.

"CSG meeting '88," 15 audiocassettes. (Recorded by Greg WILLIAMS.)

DENNIS, Brent G., "Reorganization and the left brain interpreter: A connection?" unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, September-October, [2] pages.

DELPRATO, Dennis J., and Anna M. WISE, "Some preliminary considerations on the application of control system theory to stereotyped behavior," unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, September-October, 4 + [2] pages.

"Fall '88 ASC conference", 12 Audiocassettes. (From Rodney DONALDSON.)

[Fall '88 ASC conference], 38 Audiocassettes. (Recorded by Greg WILLIAMS.)

FENNELL, Geraldine, "What is special about behavioral cybernetics?" CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS, Summer, 10-13.

FLEISCHAKER, Gail Raney, AUTOPOIESIS: SYSTEM LOGIC AND ORIGINS OF LIFE, Ph.D. Dissertation, Boston University, April, xii + 122 pages.

FLEISCHAKER, Gail Raney, "Autopoiesis: The status of its system logic," BIOSYSTEMS 22(1), 37-49.

FORD, Edward E., "Control theory," 1 page. (Five different versions, later ones titled "A diagram of control theory.")

FORD, Edward E., "COUNSELING," card printed on both sides.

FORD, Edward E., Information on Reality Therapy course offerings, [1] page.

GOLDSTEIN, David M., "Life perceptions profile Q-sort: A control systems counseling assessment tool," unpublished paper presented at

CSG meeting, Haimowoods, Wisconsin, September-October, 13 + [5] + 3-page "LPPQ answer sheet" and 1-page "LPPQ data sheet."

GOSSEN, D., "How am I causing myself this problem by the picture I have in my head?" unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, September-October, [2] pages. (Numbered "18" and "19.")

HERSHBERGER, Wayne A., "Psychology as a conative science," preprint from AMERICAN PSYCHOLOGIST distributed at CSG meeting, Haimowoods, Wisconsin, September-October, [1] + i + 4 pages.

HERSHBERGER, Wayne, and SKINNER, B. F., "Some overt behavior is neither elicited nor emitted," in A. Charles CATANIA and Stevan HARNAD, editors, THE SELECTION OF BEHAVIOR[:] THE OPERANT BEHAVIORISM OF B. E SKINNER: COMMENTS AND CONSEQUENCES, Cambridge University Press, Cambridge, et al., 107-109.

HYLAND, Michael, "Setting the record straight," CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (12), Spring, 16.

HYLAND, Michael E., "Motivational control theory: An integrative framework," JOURNAL OF PERSONALITY AND SOCIAL PSYCHOLOGY 55(4), date?, 642-651.

MATURANA, Humberto, "The notions of cybernetics," CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (12), Spring, 7.

MARKEN, Richard S., "Degrees of freedom in behavior," unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, September-October, 29 + [8] pages.

MARKEN, Richard S., "REPLY: And yet they seek goals," CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (14), Fall, 4-5.

MARKEN, Richard S., "The nature of behavior: Control as fact and theory," BEHAVIORAL SCIENCE 33(number?), date?, 196-206.

McCORD, David M., "An [sic] technical note on programming disturbances," unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, September-October 1988, 4 pages.

MIDGLEY, Bryan D., "On behaviorists, control theorists, and straw men," CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (14), Fall, 3-4.

MIDGLEY, Bryan D., "REJOINDER: Straw men... again," CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (14), Fall, 6.

POWERS, Mary A., "Cybernetics: Is usefulness important? Yes! Does 'understanding' convey its meaning better than 'control'? No!" CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (13), Summer, 3-4.

POWERS, Mary A., "On the ambiguity of action," CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (14), Fall, 6-7.

POWERS, William T., "A comparison of Maturana's and Powers' models," unpublished paper distributed by author at ASC conference, Felton, California, October, 14 pages. (Also earlier [August] draft.)

POWERS, William T., "A comparison of Powers' and von Glasersfeld's epistemologies," unpublished paper distributed by author at ASC conference, Felton, California, October, 8 pages.

POWERS, William T., "A manifesto for control theorists," unpublished paper presented at CSG meeting, September-October, [4] pages.

POWERS, William T., "A nonfunctional analysis of behavior," THE BEHAVIORAL AND BRAIN SCIENCES 11(1), March, 143-144.

POWERS, William T., "An outline of control theory" in CONFERENCE WORKBOOK FOR "TEXTS IN CYBERNETIC THEORY"[:] AN IN-DEPTH EXPLORATION OF THE THOUGHT OF HUMBERTO R. MATURANA, WILLIAM T. POWERS, AND ERNST VON GLASERSFELD, American Society for Cybernetics, Felton, California, October 18-23, 1-32.

[POWERS, William T.,] "Causation in behavior," unpublished paper, [1] page.

POWERS, William T., "COMMENT: The trouble with S-delta," CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (14), 5-6.

[POWERS, William T.,] "Cybernetics, utility, control theory, who dealt this mess?" unpublished paper, [2] pages.

[POWERS, William T.,] Draft of untitled book, [August], paginated by chapter.

[POWERS, William T.,] Draft of untitled book, [September], paginated by chapter.

[POWERS, William T.,] "Simulating behavior," unpublished paper, [12] pages.

[POWERS, William T.,] "Simulating organisms," unpublished paper, [9] pages.

[POWERS, William T.,] "Technical comment: An adaptive neural model," unpublished paper, September 6, 3 pages. (With letter from SCIENCE editor to author dated September 28, 1988.)

[POWERS, William T.,] "The good, the true, and the real," unpublished paper, 3 pages.

[POWERS, William T.,] "The ontology of observing: A condensation," unpublished paper, [15] pages. (With nearly identical "The ontology of observing, in language I understand," [14] pages.)

[POWERS, William T.,] "Tracking and principles of human behavior," unpublished paper, [7] pages.

RICHARDSON, George P., FEEDBACK THOUGHT IN SOCIAL SCIENCE, unpublished manuscript, October, paginated by chapter.

ROBERTSON, Richard J., David M. GOLDSTEIN, Michael MERMEL, and Melanie MUSGRAVE, "Testing the self as a control system," unpublished paper presented by first two authors at CSG meeting, Haimowoods, Wisconsin, September-October, 18 + [3].

RUNKEL, Philip J., CASTING NETS AND TESTING SPECIMENS: METHOD IN SOCIAL SCIENCE, unpublished manuscript, June, 616 pages.

SMITH, Thomas J., and Karl U. SMITH, "The cybernetic basis of human behavior and performance," CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (15), Winter, 1-28.

STEIER, Frederick, “On cybernetics as reflex understanding,” CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (12), Spring, 7-8.

TUCKER, Charles W., “Demonstrating control theory,” CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (12), Spring, 16-19.

1989

AITKEN, Michael Luke, “Dick and Jane at Haimowoods,” [4] pages. (Included with CONTROL SYSTEMS GROUP NEWSLETTER, November 3.)

AITKEN, Michael Luke, “Interplay: On the need to adapt educational method to organismic process,” CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (17), Summer, 10-13.

BANDURA, Albert, “Human agency in social cognitive theory,” AMERICAN PSYCHOLOGIST 44(9), September 1989, 1175-1184.

BOURBON, Tom, “Comments for the Smiths,” CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (16), Spring, 4-5.

BOURBON, W. T., K. E. COPELAND, V. R. DYER, W. K. HARMAN, and B. L. MOSELY, “Data from, ‘Reliability and stability of predictions by control system theory,’” unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, September, [3] pages.

BOURBON, W. T., and W. T. POWERS, “Data from, ‘Models and their worlds,’” unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, September, [4] pages.

BOURBON, W. T., and W. T. POWERS, “Data from, ‘Using control system theory to simulate other models of behavior,’” unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, September, [1] page.

CONTROL SYSTEMS GROUP NEWSLETTER, November 3, 6 pages.

“CSG [meeting],” 14 audiocassettes. (Recorded by Greg WILLIAMS.)

[CSG NEWSLETTER,] [May], 6 + [1] pages.

CSG NEWSLETTER, [November (addendum)], [1] page.

DELPRATO, Dennis J., “When self-regulation is not regulation by a self (or the tortuous career of naturalistic behavioral control),” CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (16), Spring, 1-2.

FORD, Edward E., “A diagram of control theory,” distributed by author at CSG meeting, Haimowoods, Wisconsin, September, [1] page.

FORD, Edward E., “Counseling and control theory,” unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, September, 12 pages.

FORD, Edward E., FREEDOM FROM STRESS, Brandt Publishing, Scottsdale, Arizona, xvi + 184 pages. (Both first and second, slightly revised, printings.)

FORD, Edward E., “FREEDOM FROM STRESS based on cybernetic control theory,” card printed on both sides.

FORD, Edward E., “LOVE,” card printed on both sides. (Original edition 1986.)

FORD, Edward E., “Problems people have trying to understand control theory,” unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, September, [1] page.

FORD, Edward E., “TEACHING DISCIPLINE based on cybernetic control theory,” card printed on both sides.

FORD, Edward E., “TEACHING RESPONSIBILITY based on cybernetic control theory,” card printed on both sides.

[FORD, Edward E., Diane GOSSEN, and William GLASSER,] “Gossen lecture—1989 RT convention,” unpublished paper, 7 pages. (Sent by Edward E. FORD.)

GIBBONS, Hugh, JUSTIFYING LAW[:] SPRING 1989, Franklin Pierce Law Center, Concord, New Hampshire, [115] pages. (Reprinted from LAW AND PHILOSOPHY 3(number?), 1984, 165-279.)

GOLDSTEIN, David M., “Control theory applied to clinical psychology[:] 1989 update,” unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, September, [6] pages.

GOSSEN, Diane, CONTROL THEORY: IN ACTION, Chelsom Consultants, Saskatoon, Saskatchewan, [vi] + 96 pages.

GOSSEN, Diane, CONTROL THEORY[:] MY CHILD IS A PLEASURE TO LIVE WITH, Chelsom Consultants, Saskatoon, Saskatchewan, [iv] + 67 pages.

HERSHBERGER, Wayne A., editor, VOLITIONAL ACTION[:] CONATION AND CONTROL, ADVANCES IN PSYCHOLOGY 62, Elsevier Science Publishers (North-Holland), Amsterdam, xiii + 572 pages. (With typesetting materials and galley proofs of some chapters.)

MARKEN, Richard S., "Three positive points," CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (16), Spring, 13.

MARKEN, Richard S., and William T. POWERS, "Random-walk chemotaxis: Trial and error as a control process," BEHAVIORAL NEUROSCIENCE 103(6), December, 1348-1355. (Also have manuscript, April 18, 1989.)

MILLER, Michelle A., CONTROL SYSTEM THEORY: A STUDY OF THE THEORY AS A TOOL FOR MANAGEMENT, M.A. Thesis, Stephen F. Austin State University, Nachadoches, Texas, December, iii + 69 pages. (Thesis Director: Tom BOURBON.)

PAVLOSKI, Raymond, "A control system approach to cardiovascular reactivity: Behavioral models that behave," PSYCHOPHYSIOLOGY 26(4), July, 468-481. (Galley proof.)

PAVLOSKI, Raymond, "Furedy's straw man: Wrong purpose," PSYCHOPHYSIOLOGY 26(4), July, 485-487. (Galley proof.)

POWERS, Bill, "Things I'd like to say if they wouldn't think I'm a nut," unpublished paper, [1] page.

POWERS, Mary A., "Footnote to Felton," CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (16), Spring, 9.

POWERS, William T., "A challenge for theorists," unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, September, 10 pages.

POWERS, William T., "A note to my colleagues in the CSG," unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, September, [1] page.

POWERS, William T., "After Galileo," unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, September, [7] pages. (Printout from computer file; paper was not distributed at meeting.)

POWERS, William T., "By any other name," CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (16), Spring, 11-12.

POWERS, William T., "Control theory and statistical generalizations," unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, September, 6 + [2] pages.

POWERS, William T., "Human nature[:] from the standpoint of control theory," unpublished paper, [12] pages.

POWERS, William T., LIVING CONTROL SYSTEMS[:] SELECTED PAPERS OF WILLIAM T. POWERS, The Control Systems Group, Inc., Gravel Switch, Kentucky, xx + 300 pages. (With typesetting materials.)

POWERS, William T., "Modeling and the matching law," unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, September, 5 pages.

POWERS, William T., "On the uses of the term 'purpose,'" unpublished paper, 111] pages.

POWERS, William T., "Problems with the Smiths," CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (16), Spring, 3-4.

POWERS, William T., "Some basic concepts of cybernetic control theory," distributed by author at CSG meeting, [1] page.

[POWERS, William T.,] "The fallacy of behaviorism[:] Version 1," unpublished paper, [6] pages.

[POWERS, William T.,] "The meaning of cybernetic control theory," unpublished paper, [9] pages.

POWERS, William T., "Using control theory in higher realms of psychology," unpublished paper presented at CSG meeting, Haimowoods, Wisconsin, September, 22 pages.

ROBERTSON, Richard J., "At last, serious dialogue on cybernetics," CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (16), Spring, 3.

ROBERTSON, Richard J., and William T. POWERS, editors, INTRODUCTION TO MODERN PSYCHOLOGY: THE CONTROL THEORY VIEW, Richard J. Robertson, Chicago, iii + 316 pages.

RONJAK, Gary, "Book review [of FREEDOM FROM STRESS by Edward E. Ford]," CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (18), Fall, 13-14.

RUNKEL, Philip J., "Please spread the word," CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (16), Spring, 2-3.

[WILLIAMS, Gregory,] "New book by William T. Powers," CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (18), Fall, 14.

[WILLIAMS, Gregory,] "VOLITIONAL ACTION: CONATION AND CONTROL," CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (18), Fall, 14.

WILLIAMS, W. D., "Making it clearer," CONTINUING THE CONVERSATION: A NEWSLETTER OF IDEAS IN CYBERNETICS (16), Spring, 9-10.

1990

ANONYMOUS, "Glossary [of control-theory terms]," AMERICAN BEHAVIORAL SCIENTIST 34(1), September/October, 119-121.

BOURBON, W. Thomas, "Invitation to the dance[:] Explaining the variance when control systems interact," AMERICAN BEHAVIORAL SCIENTIST 34(1), September/October, 95-105.

BOURBON, W. Thomas, Kimberly E. COPELAND, Vick R. DYER, Wade K. HARMAN, and Barbara L. MOSLEY, "On the accuracy and

reliability of predictions by control-system theory," PERCEPTUAL AND MOTOR SKILLS 71(number?), date?, 1331-1338.

["CSG meeting,"] 24 audiocassettes. (Recorded by Greg WILLIAMS and Dag FORSSELL.)

CSG NEWSLETTER, February, 6 pages. [CSG NEWSLETTER], [May], 7 + [2] pages. CSG NEWSLETTER, October, 14 pages.

CARVER, Charles S., and Michael E SCHEIER, "Origins and functions of positive and negative affect: A control-process view," PSYCHOLOGICAL REVIEW 97(1), January, 19-35.

CRAWFORD, C. Milton, III, SELF-DETERMINED FEEDING AND DRINKING BY HOODED RATS, M.A. Thesis, Stephen F. Austin State University, Nachadoches, Texas, December, vii + 80 pages. (Thesis Director: W. T. BOURBON.)

D'AGOSTINO, Brian, "Self perception and national security policy," unpublished paper presented at CSG meeting, Indiana, Pennsylvania, August, 5 pages.

DENNIS, Brent G., "'You are a living control system': An explanation for clients," unpublished paper distributed at CSG meeting, Indiana, Pennsylvania, August, 1 page.

DYER, Victor R., THE CONTROL OF LIMB MOVEMENT: AN EXAMINATION OF NEGATIVE FEEDBACK AND EFFERENT MOTOR COMMANDS, M.A. Thesis, Stephen F. Austin State University, Nachadoches, Texas, June, vii + 74 pages. (Thesis Director: W. Thomas BOURBON.)

ELGIN, Suzette Haden, "The GENTLE ART syntonics system of verbal self-defense," unpublished paper distributed by Luke AITKEN at CSG meeting, Indiana, Pennsylvania, August, 4 pages.

FORD, Edward E., FREEDOM FROM STRESS, revised edition, Brandt Publishing, Scottsdale, Arizona, xvi + 184 pages.

FORD, Edward E., "On understanding control theory[:] Learning some difficult concepts," AMERICAN BEHAVIORAL SCIENTIST 34(1), September/October, 117-118.

FORSSELL, Dag C., "A systems framework for understanding the Deming management system," unpublished paper, February 4, 15 pages.

FORSSELL, Dag C., "One person's route to enthusiasm about Deming's 14 point management system. Or one way to teach the Deming management philosophy," unpublished paper, February 11, 14 pages.

GIBBONS, Hugh, THE DEATH OF JEFFREY STAPLETON[:] EXPLORING THE WAY LAWYER'S [sic] THINK, Franklin Pierce Law Center, Concord, New Hampshire, [i] + 197 pages.

GOLDSTEIN, David M., "Clinical applications of control theory," AMERICAN BEHAVIORAL SCIENTIST 34(1), September/October, 110116.

GOLDSTEIN, David M., "Control theory approach to studying the self-image in a clinical setting," unpublished paper presented at CSG meeting, Indiana, Pennsylvania, August, 9 pages.

GOLDSTEIN, David M., "Psychotherapy based on cybernetic control theory: What is it and how does it compare to other therapy approaches?" unpublished paper presented at CSG meeting, Indiana, Pennsylvania, August, 5 pages.

HANCOCK, Thomas Emerson, "A test of aspects of a control model of instructional feedback," unpublished paper presented at CSG meeting, Indiana, Pennsylvania, August, [17] pages.

HERSHBERGER, Wayne A., "Control theory and learning theory," AMERICAN BEHAVIORAL SCIENTIST 34(1), September/October, 5566.

HORVATH, Stephen, Letter to CSG members about special issue of AMERICAN BEHAVIORAL SCIENTIST, October 16, [2] pages. (With enclosures.)

KEMP, Steven M., "Homeostasis vs. stochastic processes," unpublished paper, [1] + 14 pages. (With related materials, including rejoinder by William T. POWERS.)

"Life Problems Control Institute," flyer distributed by David M. GOLDSTEIN at CSG meeting, Indiana, Pennsylvania, August, [2] pages.

MARKEN, Richard S., "A science of purpose," AMERICAN BEHAVIORAL SCIENTIST 34(1), September/October, 6-13.

MARKEN, Richard S., "Spreadsheet analysis of a hierarchical control system model of behavior," BEHAVIOR RESEARCH METHODS, INSTRUMENTS, & COMPUTERS 22(4), date?, 349-359. (Filed with MIND READINGS by Richard S. MARKEN, 1992.)

McPHAIL, Clark, and Charles W. TUCKER, "Purposive collective action," AMERICAN BEHAVIORAL SCIENTIST 34(1), September/October, 81-94.

MEYER, Jean-Arcady, and Agnes GUILLOT, FROM ANIMALS TO ANIMATS: EVERYTHING YOU WANTED TO KNOW ABOUT THE SIMULATION OF ADAPTIVE BEHAVIOR, Technical Report BioInfo-90-1, Ecole Normale Superieure, Paris, September, 40 pages.

PAVLOSKI, Raymond P., Gerard T. BARRON, and Mark A. HOGUE, "Reorganization[:] Learning and attention in a hierarchy of control systems," AMERICAN BEHAVIORAL SCIENTIST 34(1), September/October, 32-54.

PIRSIG, Robert M., "Quality," reprinted from page 215 of ZEN AND THE ART OF MOTORCYCLE MAINTENANCE (1974), distributed anonymously at CSG meeting, Indiana, Pennsylvania, August, [1] page.

PLOOIJ, Frans X., and Hedwig H. C. VAN DE RIJT-PLOOIJ, "Developmental transitions as successive reorganizations of a control hierarchy," AMERICAN BEHAVIORAL SCIENTIST 34(1), September/October, 67-80.

POWERS, William T., "Control theory: A model of organisms," SYSTEM DYNAMICS REVIEW 6(1), Winter, 1-20. (With typesetting materials.)

POWERS, William T., "Control theory and statistical generalizations," AMERICAN BEHAVIORAL SCIENTIST 34(1), September/October, 2431.

POWERS, William T., "Standing at the crossroads," unpublished paper presented at CSG meeting, Indiana, Pennsylvania, August, 9 pages.

POWERS, William T., "The x phenomenon," unpublished paper distributed by author at CSG meeting, Indiana, Pennsylvania, August, 4 pages.

ROBERTSON, Richard J., and William T. POWERS, editors, INTRODUCTION TO MODERN PSYCHOLOGY: THE CONTROL-THEORY VIEW, The Control Systems Group, Inc., Gravel Switch, Kentucky, x + 238 pages. (With typesetting materials and correspondence among authors and Greg WILLIAMS.)

RUNKEL, Philip J., CASTING NETS AND TESTING SPECIMENS: TWO GRAND METHODS OF PSYCHOLOGY, Praeger, New York, Westport, Connecticut, and London, xiv + 216 pages. (With "Book News" publisher's announcement, [2] pages.)

RUNKEL, Philip J., "Research method for control theory," AMERICAN BEHAVIORAL SCIENTIST 34(1), September/October, 14-23.

WILLIAMS, Greg, collector, "Materials related to living control systems theory," Version 1, self-published, July, [22] pages.

WILLIAMS, Greg, collector, "Materials related to living control systems theory," Version 2, self-published, August, [22] pages.

WILLIAMS, Greg, "What is a dream, that a control system should have it?" unpublished paper presented at CSG meeting, Indiana, Pennsylvania, August, [2] pages. (With related materials.)

WILLIAMS, William D., "The Giffen effect[:] A note on economic purposes," AMERICAN BEHAVIORAL SCIENTIST 34(1), September/October, 106-109.

1991

BANDURA, Albert, "Human agency: The rhetoric and the reality," AMERICAN PSYCHOLOGIST 46(2), February, 157-162. (Rejoinder to "Commentary on Bandura's 'human agency" by William T. POWERS.)

BREWER, Charles L., Review of INTRODUCTION TO MODERN PSYCHOLOGY: THE CONTROL-THEORY VIEW by Richard J. ROBERTSON and William T. POWERS, CONTEMPORARY PSYCHOLOGY 36(5), May, 450.

CLOSED LOOP (2), May, 24 pages. (Included with [CSG NEWSLETTER,] [February]. With typesetting materials.)

CLOSED LOOP (3), July, 36 pages. (With typesetting materials.)

CLOSED LOOP (4), October, 52 pages. (Included with [CSG NEWSLETTER] November. With typesetting materials.)

[CSG meeting,] 28 audiocassettes. (Recorded by Greg WILLIAMS and Dag FORSSELL.)

[CSG NEWSLETTER,] [February,] 19 pages. (CLOSED LOOP [(1)] is part of this issue.)

[CSG NEWSLETTER,] May 28, 6 pages. [CSG NEWSLETTER,] November, 11 pages.

COPELAND, Kimberly E., JUVENILE DELINQUENTS' PERCEPTIONS AND CONTROL, M.A. Thesis, Stephen F. Austin State University, Nachadoches, Texas, August, vii + 108 pages.

CZIKO, Gary A., "The past, present and future of CSGnet," unpublished paper presented at CSG meeting, Durango, Colorado, August, 4 pages.

D'AGOSTINO, Brian, "Self perception and national security policy," unpublished paper distributed by author at CSG meeting, Durango, Colorado, August, paginated by section.

DEWEY, John, "A short catechism concerning truth," unpublished paper distributed by Chuck TUCKER at CSG meeting, Durango, Colorado, August, 6 pages.

DOOLEY, Jeffrey E., GENETIC EPISTEMOLOGY, CONSTRUCTIVISM, AND AREALISM: KNOWLEDGE FOR SELF-ORGANIZATION OF THE PHYSICO-COGNITIVE WHOLE, M.S. Thesis, San Jose State University, California, December, vii + 176 pages. (Approved by William POWERS and others. With letter from Jeff DOOLEY to Greg WILLIAMS dated April 13, 1992.)

FORSSELL, Dag C., THE DEMING MANAGEMENT PHILOSOPHY[:] OBSERVATIONS, INTERPRETATION, COMMENTARY, Purposeful Leadership, Valencia, California, 26 pages.

FORSSELL, Dag C., "Purposeful Leadership," flyer distributed by author at CSG meeting, Durango, Colorado, August, [2] pages.

GERONILLA, Linda, "Control theory/reality therapy," unpublished paper distributed by author at CSG meeting, Durango, Colorado, August, [2] pages.

GERONILLA, Linda, "Evaluating my roles using reality therapy," unpublished paper distributed by author at CSG meeting, Durango, Colorado, August, [2] pages.

GERONILLA, Linda, "Hierarchy of control—eating behaviors," unpublished paper distributed by author at CSG meeting, Durango, Colorado, August, [1] page. ("Systems theory" on verso.)

GERONILLA, Linda, "Systems theory," unpublished paper distributed by author at CSG meeting, Durango, Colorado, August, [1] page. ("Hierarchy of control—eating behaviors" on verso.)

GERONILLA, Linda, "What do you want?" unpublished paper distributed by author at CSG meeting, Durango, Colorado, August, [2] pages.

GOLDSTEIN, David M., "Towards a perceptual control theory psychotherapy," unpublished paper, 23 pages.

HARMAN, Wade Kenan, USING PERCEPTUAL CONTROL THEORY TO SIMULATE THE CONTROL OF FOUR DEGREES OF FREEDOM FOR AN ERGONOMIC DESIGN, M.A. Thesis, Stephen F. Austin State University, Nachadoches, Texas, December, vi + 89 pages. (Thesis Director: W. Thomas Bourbon.)

HERSHBERGER, Wayne A., "Volitional action: The control of consequences," unpublished paper distributed by author at CSG meeting, Durango, Colorado, August, 10 pages.

HERSHBERGER, Wayne A., and J. Scott JORDAN, "Visual direction constancy: Perceiving the visual direction of perisaccadic flashes," unpublished paper distributed by principal author at CSG meeting, Durango, Colorado, August, 38 + [4] pages.

KLEIN, Howard J., "Control theory and understanding motivated behavior: A different conclusion," MOTIVATION AND EMOTION 15(1), March, 29-44.

LaMANTIA, Frank, TRACKING PERFORMANCE AS A PREDICTOR FOR IMPAIRED DRIVING, M.A. Thesis, Stephen F. Austin State University, Nachadoches, Texas, July, 52 pages.

LAZARE, Mark A., "Thesis proposal: A control theory analysis of adaptation and social interaction," unpublished paper presented at CSG meeting, Durango, Colorado, August, 24 pages. (Committee Chairman: W. Thomas BOURBON.)

LOCKE, Edwin A., "Goal theory vs. control theory: Contrasting approaches to understanding work motivation," MOTIVATION AND EMOTION 15(1), March, 9-28.

MARKEN, Richard S., "Degrees of freedom in behavior," PSYCHOLOGICAL SCIENCE 2(2), March, 92-100. (Filed with MIND READINGS by Richard S. Marken, 1992.)

McCLELLAND, Kent, "Perceptual control and sociological theory," unpublished paper ("DRAFT 8/1/91") distributed by author at CSG meeting, Durango, Colorado, August, 31 + [3] pages.

McPHAIL, Clark, THE MYTH OF THE MADDING CROWD, Aldine de Gruyter, Hawthorne, New York, xxx + 265 pages.

McPHAIL, Clark, William T. POWERS, and Charles W. TUCKER, "Simulating individual and collective action in temporary gatherings," unpublished paper presented at CSG meeting, Durango, Colorado, August, 14 + [12] pages.

[POWERS, William T.,] "Beckman Inst. WTP," March 15, 1 audiocassette. (Taped at University of Illinois, Urbana-Champaign.)

POWERS, William T., "Commentary on Bandura's 'human agency,'" AMERICAN PSYCHOLOGIST 46(2), February, 151-153.

[POWERS, William T.,] "WTP@UIUC," March 15, 1 audiocassette. (Taped at University of Illinois, Urbana-Champaign.)

RICHARDSON, George P, FEEDBACK THOUGHT IN SOCIAL SCIENCE AND SYSTEMS THEORY, University of Pennsylvania Press, Philadelphia, x + 374 pages.

ROBERTSON, Richard J., "Instinctual issues in the CT approach to psychotherapy," unpublished paper distributed by author at CSG meeting, Durango, Colorado, August, 11 pages.

ROBERTSON, Richard J., “Students’ work for a course grade as an instance of higher-level system function,” unpublished paper distributed by author at CSG meeting, Durango, Colorado, August, 5 pages.

TUCKER, Charles W., “Sociology 525[:] Selves and social transactions,” unpublished paper distributed by author at CSG meeting, Durango, Colorado, August, [1] page.

WILLIAMS, Greg, collector, “Materials related to living control systems theory,” Version 4, self-published, July, 25 pages.

WILLIAMS, Greg, “The fable of the radio,” unpublished paper presented at CSG meeting, Durango, Colorado, August, [3] pages.

WILLIAMS, Greg, and Pat WILLIAMS, “An alternative model of control-hierarchy organization?” unpublished paper distributed by principal author at CSG meeting, Durango, Colorado, August, [1] page.

1992

ACKOFF, Russell L., Seminar on interactive planning, Purposeful Leadership, Valencia, California, paginated by section. (Lectures originally given in 1978, transcribed and annotated by Dag C. FORSSELL.)

AMUNDSON, Roy, Ronald C. SERLIN, and Richard LEHRER, “On the threats that do NOT face educational research,” EDUCATIONAL RESEARCHER 21(9), December, 19-24. (Reply to “Perceptual control theory” by Gary A. CZIKO.)

ANONYMOUS, “Checklist for ordering recent materials in psychology,” PERCEPTUAL AND MOTOR SKILLS 75(number?), August, 336-338. (Mentions LIVING CONTROL SYSTEMS II by W. T. POWERS.)

ANONYMOUS, “Checklist for ordering recent materials in psychology,” PERCEPTUAL AND MOTOR SKILLS 75(number?), October, 672-674. (Mentions MIND READINGS by R. S. MARKEN.)

ANONYMOUS, “Checklist for ordering recent materials in psychology,” PSYCHOLOGICAL REPORTS 75(number?), August, 336-338. (Mentions LIVING CONTROL SYSTEMS II by W. T. POWERS.)

ANONYMOUS, “Checklist for ordering recent materials in psychology,” PSYCHOLOGICAL REPORTS 71(number?), October, 672-674. (Mentions MIND READINGS by R. S. MARKEN.)

CAMPBELL, Donald T., “Distinguishing between pattern in perception due to the knowing mechanisms and pattern plausibly attributable to the referent,” unpublished paper, 13 pages.

CLOSED LOOP[:] THREADS FROM CSGNET 2(1), Winter, 72 pages. (With CSG NEWSLETTER, January, [3] pages, and typesetting materials.)

CLOSED LOOP[:] THREADS FROM CSGNET 2(2), Spring, 73 pages. (With CSG NEWSLETTER, [April], [3] pages, and typesetting materials.)

CLOSED LOOP[:] THREADS FROM CSGNET 2(3), Summer, 72 pages. (With typesetting materials.)

CLOSED LOOP[:] THREADS FROM CSGNET 2(4), Fall, 72 pages. (With typesetting materials.)

[CSG meeting,] 16 audiocassettes. (Recorded by Greg WILLIAMS and Dag FORSSELL.)

CSG NEWSLETTER, January, 3 pages.

CZIKO, Gary A., “From rubber bands to robots: Portable demonstrations of perceptual control,” unpublished paper distributed at CSG meeting, July-August, 4 pages.

CZIKO, Gary A., “Perceptual control theory: One threat to educational research NOT (yet?) faced by Amundson, Serlin, and Lehrer,” EDUCATIONAL RESEARCHER 21(9), December, 25-27.

CZIKO, Gary A., “Purposeful behavior as the control of perception: Implications for educational research,” unpublished paper distributed at CSG meeting, Durango, Colorado, July-August, 13 pages. (“Revised April 1992.”)

CZIKO, Gary A., “Purposeful behavior as the control of perception: Implications for educational research,” EDUCATIONAL RESEARCHER 21(9), December, 10-18, 27. (With letter from Gary CZIKO to Greg WILLIAMS dated December 7, 1992, and letter from Greg WILLIAMS to Gary CZIKO dated December 12, 1992.)

D’AGOSTINO, Brian, “Self images of hawks and doves,” unpublished paper distributed at CSG meeting, July-August, [2] + 18 + [8] pages.

FORD, Edward E., "LOVE/QUALITY TIME" and (on verso) "TEACHING RESPONSIBLE THINKING," card printed on both sides.

FORSSELL, Dag C., "Cassette tape script: Introduction, explanation[;] Demonstration[:] 'Game board' and illustrations[;] Program agenda," Purposeful Leadership, Valencia, California, July 20. (Distributed by author at CSG meeting, Durango, Colorado, July-August 1993.)

FORSSELL, Dag C., "Control: What it is; where it applies," July 16, unpublished paper distributed by author at CSG meeting, Durango, Colorado, July-August, 2 pages.

FORSSELL, Dag C., PURPOSEFUL LEADERSHIP: MANAGEMENT PROGRAMS FOR PEOPLE WHO THINK: 1 BASIC APPLICATIONS, Purposeful Leadership, Valencia, California, paginated by chapter. (With envelope containing "Introduction" audiocassette and related materials.)

GANN, Clifford P., III, NONVERBAL COMMUNICATION IN A COOPERATIVE TASK MODELED BY PERCEPTUAL CONTROL THEORY, M.A. Thesis, Stephen F. Austin State University, Nachadoches, Texas, May 16, v + 62 pages. (Thesis Director: Tom BOURBON.)

JUDD, Joel Bennett, SECOND LANGUAGE ACQUISITION AS THE CONTROL OF NON-PRIMARY LINGUISTIC PERCEPTION: A CRITIQUE OF RESEARCH AND THEORY, PhD. dissertation, University of Illinois at Urbana-Champaign, v + 142 pages. (Committee Chairperson: Gary CZIKO.)

KIRAN, Carol, "Reorganization in learning," unpublished paper distributed by Richard ROBERTSON at CSG meeting, 3 pages.

KURTZER, Marc Louis, "Rice as continual and quantal disturbances; A followup on Powers's bean demonstration," unpublished paper presented at CSG meeting, 10 + [2] pages.

LAZARE, Mark A., A PERCEPTUAL CONTROL THEORY ANALYSIS OF COOPERATION AND SOCIAL INTERACTION, M.A. Thesis, Stephen F. Austin State University, Nachadoches, Texas, 43 + [8] pages. (Draft distributed by author at CSG meeting, Durango, Colorado, July-August.)

LAZARE, Mark A., A PERCEPTUAL CONTROL THEORY ANALYSIS OF COOPERATION AND SOCIAL INTERACTION, M.A. Thesis,

Stephen F. Austin State University, Nachadoches, Texas, iv + 126 pages. (Thesis Director: W. Thomas BOURBON.)

MALER, Oded, "Hybrid systems and real-world computations," paper presented at the 1st Workshop on Hybrid Systems, Lyngby, United Kingdom, [10] pages. (With letter to Greg WILLIAMS dated March 5, 1993.)

MALER, Oded, "Why should we build artificial worms and how?" unpublished paper, January 2, 12 pages.

MALER, Oded, Bruno ARNALDI, Georges DUMONT, and Cristophe LECERF, "Simulation experiments with a single-segment artificial worm on a two-dimensional surface without gravity," unpublished paper, January 2, 12 pages.

MARKEN, Richard S., MIND READINGS[:] EXPERIMENTAL STUDIES OF PURPOSE, The Control Systems Group, Gravel Switch, Kentucky, x + 212 pages. (With typesetting materials and correspondence among author, Greg WILLIAMS, and Bill POWERS.)

McCLELLAND, Kent, "Perceptual control and sociological theory," unpublished paper distributed by author at CSG meeting, Durango, Colorado, July-August, 32 + [2] pages. ("DRAFT 7/21/92.")

McPHAIL, Clark, William T. POWERS, and Charles W. TUCKER, "Simulating individual and collective action in temporary gatherings," SOCIAL SCIENCE COMPUTER REVIEW 10(1), Spring, 1-28.

OLSON, Mark W., "Volition and causation: Theoretical and methodological considerations," unpublished paper distributed by author at CSG meeting, Durango, Colorado, July-August, 45 pages.

POWERS, Bill, "Questionnaire-based testing — A stab at a new way of doing it," unpublished paper distributed by author at CSG meeting, Durango, Colorado, July-August, 2 pages.

POWERS, William T., CT PSYCHOLOGY AND SOCIAL ORGANIZATIONS, [Dag Forssell, Valencia, California,] 37 pages. (Reprinted from LIVING CONTROL SYSTEMS II.)

POWERS, William T., "Information: A matter of perception," unpublished paper, [5] pages.

POWERS, William T., LIVING CONTROL SYSTEMS II[:] SELECTED PAPERS OF WILLIAM T. POWERS, The Control Systems Group, Gravel Switch, Kentucky, xviii + 277 pages. (With typesetting materials and correspondence among author, Richard S. MARKEN, and Greg WILLIAMS.)

POWERS, William T., "The origins of purpose: The first metasystem transitions," unpublished paper distributed by author at CSG meeting, Durango, Colorado, July-August, 12 + [1] pages. ("Preprint draft: not for citation.")

[RANDALL, Allan E, compiler,] "Bibliography [on neural control networks]," unpublished paper, 64 pages. (With letter from Allan RANDALL to Greg Williams dated April 2, 1993.)

RANDALL, Allan F., "Neural control networks: A literature survey," unpublished paper, February, NTT Systems, Inc., Toronto, 33 pages.

[RANDALL, Allan F., compiler], "[William T.] Powers citation index, 1955 - August 1992," unpublished paper, [8] pages.

ROBERTSON, Richard, "First awkward steps toward modeling higher-level control functions: Problems with modeling students control of a course grade," unpublished paper distributed by author at CSG meeting, 6 pages.

RUNKEL, Philip J., "Making 'things' and explaining things," unpublished paper distributed by author at CSG meeting, Durango, Colorado, July-August, 24 pages.

TAYLOR, M. M., and D. A. WAUGH, "Principles for integrating voice I/O in a complex interface," paper presented at 63rd AGARD Avionics Panel Symposium: Advanced Aircraft Interfaces: The Machine Side of the Man-Machine Interface, location?, 12 pages.

1993

BOURBON, W. Thomas, "An elemental perceptual control system emulates S-R systems and plan-driven systems," unpublished paper presented at CSG meeting, Durango, Colorado, July-August, 11 pages.

BOURBON, W. Thomas, "Concerning error, disturbances, PCT, and neuroscience," unpublished paper presented at CSG meeting, Durango, Colorado, July-August, 7 pages.

CLOSED LOOP[:] THREADS FROM CSGNET 3(1), Winter, 72 pages. (With "Conference announcement" by Mary POWERS, [1] page, and typesetting materials.)

CLOSED LOOP[:] THREADS FROM CSGNET 3(2), Spring, 72 pages. (With "Conference announcement" [by Mary A. POWERS], [3] pages, and typesetting materials.)

CLOSED LOOP[:] THREADS FROM CSGNET 3(3), Summer, 72 pages. (With "The Control Systems Group[,] August 1993" by Mary A. POWERS, [1] page. and typesetting materials.)

CLOSED LOOP[:] THREADS FROM CSGNET 3(4), Fall, 72 pages. (With CSG Book Publishing advertising sheet and typesetting materials.)

"CSG '93 [meeting]," 14 audiocassettes. (Recorded by Greg WILLIAMS.)

CZIKO, Gary A., and Paul C. THURSTON, "Post-cognitive psychology and educational leadership: A perceptual-control-theory perspective," unpublished paper presented by principal author at CSG meeting, Durango, Colorado, July-August, 20 pages.

D'AGOSTINO, Brian, Memorandum to Gary CZIKO, Dag FORSSELL, Wayne HERSHBERGER, Kent McCLELLAND, Clark McPHAIL, Rick MARKEN, Bill POWERS, Bill WILLIAMS, and Greg WILLIAMS regarding "possible collaboration on a PCT-based, U.S. foreign policy decision-making modeling project," March 30, 3 pages. (With "POLI: An expert system model of U.S. foreign policy belief systems" by Charles S. TABER and two figures from another paper by TABER.)

DUGGINS-SCHWARTZ, Michelle, COMPUTER HELPER INTERACTING WITH A HUMAN: WHEN IS HELPING HELPING? M.A. Thesis, Stephen F. Austin State University, Nachadoches, Texas, August, 98 pages. (Draft distributed by author at CSG meeting, Durango, Colorado, July-August.)

DUGGINS-SCHWARTZ, Michelle, COMPUTER HELPER INTERACTING WITH A HUMAN: WHEN IS HELPING HELPING? M.A. Thesis, Stephen F. Austin State University, Nachadoches, Texas, August, viii + 103 pages.

FORD, Edward E., “Applying perceptual control theory to education and parenting,” Ed Ford & Associates, Scottsdale, Arizona, 21 pages. (Reprinted from FREEDOM FROM STRESS, Revised Edition.)

FORD, Edward E., “Building an ODDM family,” QUALITY OUTCOMES-DRIVEN EDUCATION 2(3), February, 21-24.

FORD, Ed, “CONFERENCE/TIME-OUT procedure—form five,” Ed Ford & Associates, [Scottsdale, Arizona,] [2] pages.

FORD, Edward E., FREEDOM FROM STRESS, Revised Edition, Brandt Publishing, Scottsdale, Arizona, xvi + 205 pages. (With typesetting materials and correspondence between author and Greg WILLIAMS.)

FORD, Edward E., “LOVE/QUALITY TIME” and (on verso) “TEACHING RESPONSIBLE THINKING,” card printed on both sides.

[FORSSELL, Dag, editor,] CONTROL SYSTEMS GROUP ANNUAL CONFERENCES,] FORT LEWIS COLLEGE[,] DURANGO, COLORADO[,] JULY 28-31, 1993, 3 VHS videotapes.

FORSSELL, Dag C., INTRODUCTION, VHS Videotape, Purposeful Leadership, Valencia, California, February 4, 119 minutes.

FORSSELL, Dag C., INTRODUCTION, VHS Videotape, Purposeful Leadership, Valencia, California, March 11, 120 minutes.

FORSSELL, Dag C., INTRODUCTION, VHS Videotape, Purposeful Leadership, Valencia, California, March 11, 117 minutes.

FORSSELL, Dag C., INTRODUCTION: PERCEPTUAL CONTROL THEORY “PCT,” VHS videotape, Purposeful Leadership, Valencia, California, July 17, 63 minutes.

FORSSELL, Dag C., “A person—an autonomous living control system,” Purposeful Leadership, [Valencia, California,] April 30, 1 page.

FORSSELL, Dag C., “Control: A key management insight,” unpublished paper, 8 pages. (“Draft. Printed October 16, 1993.”)

FORSSELL, Dag C., “Control structure,” “Control structure,” “Levels and areas of control,” “Levels and areas of control,” “Levels and areas of control,” sheets stapled together distributed by author at CSG meeting, Durango, Colorado, July-August, [5] pages.

FORSSELL, Dag C., “Introducing: Perceptual Control Theory ‘PCT,’” Purposeful Leadership, Valencia, California, July 17, [1] + 18 pages. (Distributed by author at CSG meeting, Durango, Colorado, July-August.)

FORSSELL, Dag C., “Introduction,” Purposeful Leadership, Valencia, California, July 3, paginated by section. (Distributed by author at CSG meeting, Durango, Colorado, July-August.)

FORSSELL, Dag C., “The next behavioral science illustrated in a thousand words,” unpublished paper presented at CSG meeting, Durango, Colorado, July-August, 3 pages.

FORSSELL, Dag C., “Why a theory-based leadership program?” Purposeful Leadership, Valencia, California, July 4, 5 pages. (“Prepublication draft.” Distributed by author at CSG meeting, Durango, Colorado, July-August.)

KAROLY, Paul, “Mechanisms of self-regulation: A systems view,” ANNUAL REVIEW OF PSYCHOLOGY 44, 23-52.

KELLY-WILSON, Lisa, “Deindividuation theory and research: A review and critique,” unpublished paper distributed by author at CSG meeting, Durango, Colorado, July-August, 14 pages.

McCLELLAND, Kent, “An introduction to Perceptual Control Theory,” unpublished paper presented at CSG meeting, Durango, Colorado, July-August, 17 + [3] pages. (“DRAFT July, 1993.”)

McCLELLAND, Kent, “Conflictive cooperation? Spreadsheet simulations of the social behavior of perceptual control systems,” unpublished paper presented at CSG meeting, Durango, Colorado, July-August, 10 + [12] pages. (“DRAFT 7/23/93.”)

McCLELLAND, Kent, “Perceptual control and social power,” unpublished paper, October, 22 + [2] pages.

McCLELLAND, Kent, “Perceptual control, interpersonal control, and social power,” unpublished paper distributed by author at CSG meeting, Durango, Colorado, July-August, 14 pages. (“DRAFT 7/21/93.”)

McPHAIL, Clark, “The dark side of purpose: Individual and collective violence in riots,” unpublished paper presented at CSG meeting, Durango, Colorado, July-August, 27 pages.

NELSON, Todd D., "The Hierarchical organization of behavior: A useful feedback model of self-regulation," CURRENT DIRECTIONS IN PSYCHOLOGICAL SCIENCE 2(4), August, 121-126.

POWERS, William T., and the Editorial Board of the Control Systems Group, "The dispute over control theory," unpublished paper distributed by principal author at CSG meeting, Durango, Colorado, July-August, 17 pages.

SCHWEINGRUBER, David, "Simulating a collective locomotion experiment with CROWD," unpublished paper presented at CSG meeting, Durango, Colorado, July-August, 17 pages.

SHEEHY, Noel, "Control theory," in Jennifer BOTHAMLEY, editor, DICTIONARY OF THEORIES, Gale Research International Ltd., London, Detroit, and Washington, D.C., 119.

TAYLOR, M. M., "Optimal decision support[:] Part 2: Perceptual Control Theory and human data fusion," unpublished paper distributed by author at CSG meeting, Durango, Colorado, July-August, 42 pages.

TUCKER, Charles W., and Clark McPHAIL, "Selves and Perceptual Control Theory," unpublished paper distributed by authors at CSG meeting, Durango, Colorado, July-August, 6 pages.

WILLIAMS, Gregory, collector, "Materials related to living control systems theory," Version 6, July, self-published, [34] pages.

1994

BENNETT, Milton, summarized by Ruth Lindsey MILLER, "Curiously controversial cybernetic conversations on cybernetics," CSS NEWS 4(2), Summer, 1-4.

CLOSED LOOP[:] JOURNAL OF LIVING CONTROL SYSTEMS, 4(1), Winter, 76 pages. (With typesetting materials.)

CLOSED LOOP[:] JOURNAL OF LIVING CONTROL SYSTEMS, 4(2), Spring, 72 pages. (With "From Mary Powers," [1] page, and typesetting materials.)

FORD, Edward E., DISCIPLINE FOR HOME AND SCHOOL, Brandt Publishing, Scottsdale, Arizona, xvi + 110 pages. (With advertising postal card, typesetting materials, and correspondence among author, Dag FORSSELL, and Greg WILLIAMS.)

[FORSSELL, Dag, editor,] CONTROL SYSTEMS GROUP ANNUAL CONFERENCE[,] FORT LEWIS COLLEGE[,] DURANGO, COLORADO[,] JULY 27-30 1994, 3 VHS videotapes.

FORSSELL, Dag, "Perceptual control: Leading the uncontrollable," unpublished paper, October 30, 8 pages. (With letter from author to Greg WILLIAMS dated November 1, 1994.)

MATTERN, Hal, "Discipline transforms school[:] New expectations improve kids' conduct," THE ARIZONA REPUBLIC, November 20, B1, B10. (Sent by Ed FORD.)

RODRIGUES, Marcos A., and Mark H. LEE, editors, PERCEPTUAL CONTROL THEORY, University of Wales, Aberystwyth, 94 pages.

RUNKEL, Philip J., "Is behavior probabilistic?" unpublished paper, April, 5 pages. (With letter from author to Greg WILLIAMS dated April 18, 1994, and "Replication of experiments" by Philip J. RUNKEL.)

RUNKEL, Philip J., "Replication of experiments," unpublished paper, April, 9 pages. (With letter from author to Greg WILLIAMS dated April 18, 1994, and "Is behavior probabilistic?" by Philip J. RUNKEL.)

MULTI-YEAR COLLECTIONS

Diskettes for IBM-PC compatible computers (360 KB and 1.2 MB) containing CSG-L log files beginning in September 1990.

Envelope containing materials and correspondence regarding ASC meetings, 1983 and 1984. (Sent by Tom BOURBON to Greg WILLIAMS on October 5, 1994.)

[WILLIAMS, Greg,] Schedules for speakers and notes on CSG meetings, 1986-1993.

The Control Systems Group is a membership organization which supports the understanding of cybernetic control systems in organisms and their environments: *living control systems.* Academicians, clinicians, and other professionals in several disciplines, including biology, psychology, social work, economics, education, engineering, and philosophy, are members of the Group. Annual meetings have been held since 1985. The CSG Business Office is located at 73 Ridge Pl., CR 510, Durango, CO 81301; phone (303) 247-7986.

The CSG logo shows the generic structure of cybernetic control systems. A Comparator (C) computes the difference between a reference signal (represented by the arrow coming from above) and the output signal from Sensory (S) computation. The resulting difference signal is the input to the Gain generator (G). Disturbances (represented by the black box) alter the Gain generator output on the way to Sensory computation, where the negative-feedback km/ is closed.