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

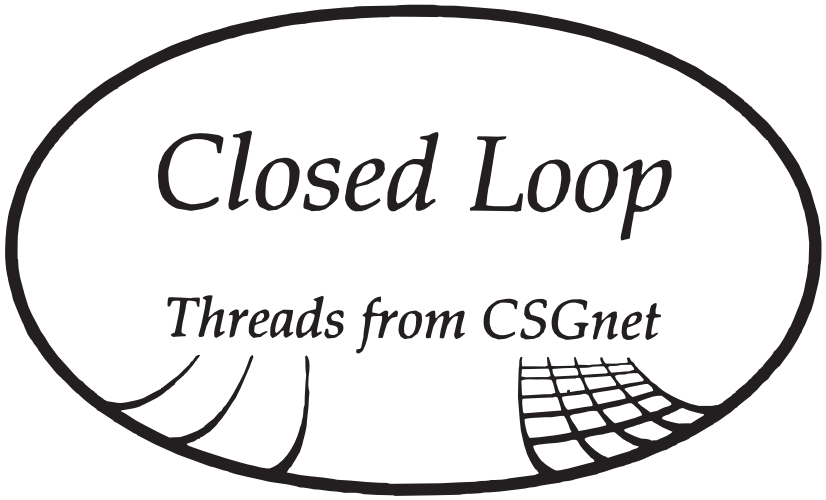

This reproduction of *Closed Loop* was created by Dag Forssell in 2009. Addresses and phone numbers have not been updated. Most are obsolete.

Posted at www.pctresources.com

Spell checked, but not proofread. Please report any errors to webmaster. Email address at website.

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 at-

tention 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 po-

litical 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 con-

trol 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 behavior-

ist. 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 consid-

eration 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 vol-

untary 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 in-

teracting 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.

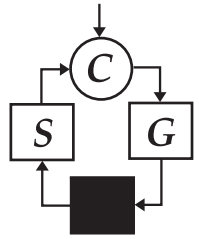

The Control Systems Group is a membership organization which supports the understanding of cybernetic control systems in organisms and their environments: *living control systems.* Academicians, clinicians, and other professionals in several disciplines, including biology, psychology, social work, economics, 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.