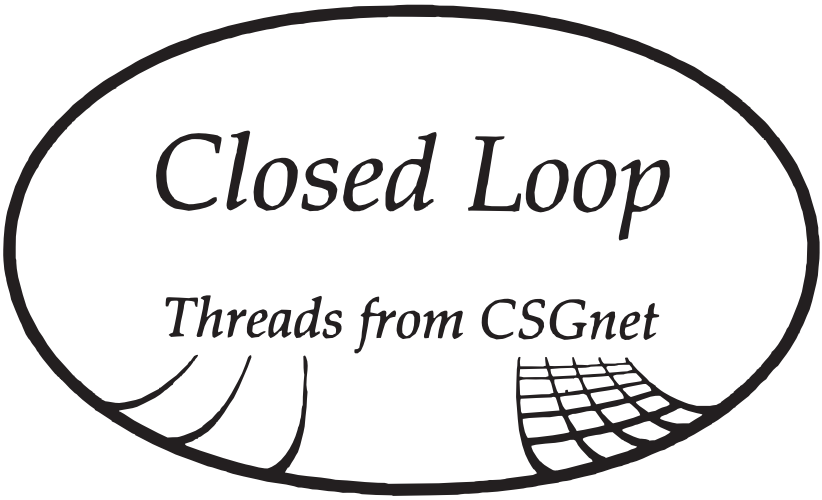

This reproduction of *Closed Loop* was created by Dag Forssell in 2001. Addresses and phone numbers have not been updated. Most are obsolete.

Posted at www.pctresources.com

Proofread as of March 8, 2001

Winter 1993

Volume 3

Number 1

Front cover

Closed Loop

Threads from CSGNet

Winter 1993 Volume 3 Number 1

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

CONTENTS

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

CSGnet, the electronic mail network for individuals interested in control theory as applied to living systems, is a lively forum for sharing ideas, asking questions, and learning more about the theory, its implications, and its problems. The "threads" in each Closed Loop, stitched together from some of the net's many conversations, exemplify the rich interchanges among netters. Some issues of Closed Loop also feature research reports by netters, in hopes of initiating new conversations.

There are no sign-up or connect-time charges for participation on CSGnet. The Internet address is "CSG-L@UIUCVMD" while CSG-L@UIUCVMD is the Bitnet address. Messages sent to CSGnet via these addresses are automatically forwarded to over 120 participants on five continents, as well as to hundreds of NetNews (Usenet) sites where CSGnet can be found as the newsgroup bit.listserv.csg-l. CSGnet also can be accessed via CompuServe, AT&T Mail, MCI Mail, or any other computer communication service with a gateway to Internet or Bitnet. For more information about subscribing to CSGnet, contact Gary Cziko, the network manager, at G-CZIKO@UIUC.EDU, phone him at (217)333-8527, or send a FAX to (217)244-7620.

Inside front cover

Conflict, Belief, Standards: Part II

Martin Taylor: In the posts I've looked at on "standards" and how they should be interpreted in the PCT world, no one has interpreted standards as I would. In interpersonal communication of any kind, including language, you can best achieve control of your percepts if you have some notion of what the other is likely to do that affects your sensory organs. If you don't want to perceive yourself being hit with a 2 x 4, you don't antagonize a Hell's Angel. You model the partner in some way. It seems to me that standards allow you to pre-empt a possibly painful random reorganization by permitting you to set references that are appropriate if the other behaves in a conventionalized way—according to standards. Likewise, if you behave according to standards, your references will be set so that your observable behavior conforms to the expectations of others—they will know what you are controlling for at the relevant level, and they will be able to interpret low-level acts/behaviors as supporting that control.

If there are any "absolute" standards, they will be those that have allowed the social groups using them to survive and prosper. A standard that allows group members to kill one another for fun is not one that is likely to be found in a long-surviving group. Our standards have been evolving since at least the time humanoids diverged from other primates, and there are clearly some sets of standards that work well together but are different from other sets that also work well together. One standard that worked well when relatively isolated tribes wandered around competing for resources involved wariness and intolerance for people not of one's own group. Killing them meant more for one's own group. Racism comes from this. But recently there has come to be only one communicating group in the world, and this long-useful standard seems to be one that will not allow this single group to survive long if it maintains its currency as a model for how to set a reference level.

Standards for grammatical usage seem to have exactly the same theoretical standing as standards for good social behavior. One sets references for using "correct" grammar because it eases the task of communicating partners who use the same standards. If a subgroup uses different standards, there's no problem except that their communication with the main group becomes less effective. If one person decides on a different set of reference levels, they cause communication problems with all of their partners. There's no moral good or bad about it, only a consideration of efficiency.

Bruce Nevin: The previous discussion of "standards" substituted that term for "principles," as in hierarchical perceptual control theory's (HPCT's) level 10. Standards, meaning "norms" or "conventions," can be on any level. Modeling others to facilitate cooperative action with them involves perceptions on many levels. Martin, the convergence of your discussion with the prior one is perhaps this: that people are aware of norms, conventions, and models of others mostly on the principle level, the level at which they attribute motivations and make moral judgments.

Rick Marken: Martin suggests that "standards" should be viewed as conventions that make it easier to cooperate. I agree that there is much to be gained from conventionalized behavior. This is particularly true in the technological world, where it helps enormously to design systems that have a standard response to actions. Thus, we can be pretty confident that a clockwise turn will result in the screw going in or the power going on or increasing. What we tend to conventionalize is the feedback function that relates our outputs to our inputs.

Martin says: "If there are any 'absolute' standards, they will be those that have allowed the social groups using them to survive and prosper." Conventional standards (like the clockwise-turn standards) can be "absolute" to the extent that we can get all objects to abide by this convention. This can be done in principle, though it's difficult (and sometimes not desired) in practice; some people might have a need for a counter-clockwise-in device. But the goal of absolute standards (conventions) is at least feasible for inanimate objects, because these objects have no purposes of their own that might conflict with the convention. Such is not the case with living systems.

The problem is that people are not inanimate objects—and certain individuals in certain circumstances might find that acting according to a particular convention is impossible, not because the person is bad or contrary or immoral, but because he or she is a hierarchical control system that simply cannot act like the knob on a radio. So my argument against "absolute" standards applies as much to standards as social conventions (like grammars) as it does to standards as moral principles. I am all for standards as conventions. The notion of absolute standards—no matter how technologically and socially helpful their existence might be—is inconsistent with human nature (if people are hierarchically organized perceptual control systems). This does not mean that I believe everybody should just go off and do their own thing. I'm just saying that this fact about human nature must be taken into consideration when we think about how people can act cooperatively.

The people who want there to be absolute standards are not "bad" people (from my point of view). The desire for absolutes is quite rea-

sonable. I can understand that desire—especially with respect to people. People should never kill each other or end a sentence with a preposition; people want predictability. All I'm saying is that people are not switches; they cannot abide by such absolute conventions, even if they try. This does not mean that social chaos is inevitable; what I think it means is that we have to find ways to cooperate that take into account the true nature of human control systems. The fact that cooperation is possible in the context of this reality (the inability of control systems to control relative to absolute conventions) is evidenced (I think) by the general spirit of cooperation found despite the diversity (in terms of many conventions) among members of the Control Systems Group itself. It can be done.

Martin Taylor: What I intended was to suggest that a "standard" provides a convenient level at which a reference value can be set, one that has often been found (perhaps by other people over history) to result in a desirable percept. But even with "absolute standards;" there's no compulsion on anyone actually to use them as reference values. As Rick says, such use might conflict with the ability to achieve other reference values. Some day, you might have to try to kill someone if you are to maintain other desired percepts, such as personal survival or freedom.

The existence of absolute standards depends on whether over evolutionary time certain behaviors (in the perceptual-control-theory sense) have benefited the survival and gene-propagation of the people (or others) using those behaviors. If they have, then either by gene transmission or by social transmission, the ordinarily effective behaviors will result in absolute standards. (On social transmission, see F. Boyd and P. J. Richardson, *Culture and the Evolutionary* Process, University of Chicago Press, 1985.)

I find no moral connotation to the idea of "standard," whether absolute or not. The idea of "absolute standard" as "you have to do what I say is right" is, I think, morally and practically repugnant, for many of the reasons adduced by Rick. But "absolute standard" as "that's what people have learned as a usually effective way to behave" is simply a practical concept that improves social interaction.

Rick Marken: I thought Martin was proposing that "standards" be understood as conventions for behavior. For example, there is a convention in the U.S. that we drive on the right. So, when I am on a road, I try to keep my car in a lane to the right of the center line. With regard to perceptual control theory (PCT), this means that I set my reference for the relationship between car and center line at "right of" rather than at some other value, like "left of." I was agreeing that standards of this sort are quite useful for successful social interaction.

Martin says that he "intended... to suggest that a 'standard' provides a convenient level at which a reference value can be set, one that has often been found (perhaps by other people over history) to result in a desirable percept." I agree, except that I think many of these standards (such as which side you drive on) are fairly arbitrary—they work as long as there is agreement among those who need to abide by them in order to avoid interpersonal conflict.

Martin goes on to say that "even with 'absolute standards; there's no compulsion on anyone actually to use them as reference values." Well, there is some social coercion. People can have unpleasant run-ins with the police if they pick the wrong side to drive on. Of course, one is still under no compulsion to set their reference at the conventional level, since he or she is the one setting it.

Then Martin says: "Me existence of absolute standards depends on whether over evolutionary time certain behaviors (in the perceptual-control-theory sense) have benefited the survival and gene-propagation of the people (or others) using those behaviors. If they have, then either by gene transmission or by social transmission, the ordinarily effective behaviors will result in absolute standards." If by "behaviors" you mean "references for certain inputs," then I agree; there might be absolute (fixed, built into the individual, unvarying) references for certain inputs. Such references are almost certainly at the cellular, if not the genetic, level—they are called "intrinsic references" in PCT. If, however, by "behaviors" you mean particular actions, then I don't see how this can be correct; evolution could not possibly select for actions that would have to produce their effects in a disturbance-prone environment. I think a lot of sociobiologists imagine that certain behaviors (in terms of actions) can evolve; for example, they talk about evolution of "aggression." It sounds like they are talking about the evolution of certain visible patterns of outputs. I think the only thing that might be able to evolve is a preference for a certain level of sensory input resulting from these (and/or other) actions.

Finally, Martin says: "But 'absolute standard' as 'that's what people have learned as a usually effective way to behave' is simply a practical concept that improves social interaction." It sounds like you are saying that an "absolute standard" is only *relatively* absolute (it is usually effective at improving social interaction, but not always). If this is what is usually meant by "absolute standard;" then it turns out that I have been advocating a version of this approach to "absolute standards" all along. I've just been saying that some standards are usually effective for lots of people—but not always (they don't work for some of the people some of the time). I just wish some of the others in the discussion of absolute standards would have clarified this point for me. Does this mean that the Ten Commandments are "absolute standards" in your sense of

absolute standards—it is usually effective to not steal, but not always? Is that what Judeo-Christians think God meant? What about that first one: thou shaft have no other god before me... usually? Some people got stewed for not obeying that one. Are some standards more absolute than others?

Bill Powers: Just a few ideas to add to the standards discussion. Any given standard, such as "helping the poor," has at least five aspects:

1. The verbal description or name of the standard ("helping the poor").
2. The perceptual meaning of the description or name of the standard: that is, how you can tell when a poor person is being "helped"?
3. The reference level for the standard: that is, what degree of the helping is the desirable degree?
4. The program of actions used to achieve the standard: that is, what actions will help the poor to the desired degree?
5. The system concepts exemplified by the standards: that is, the concept of human nature and of society that defines the goal achieved by helping the poor.

Most discourses on standards focus on the verbal description or name of the standard, under the (incorrect) assumption that it indicates the same principle to everyone. So when old-style Democrats speak of helping the poor, they mean giving them money, advice, and services that the poor people can't obtain for themselves. When Republicans speak of the same thing, they mean doing something that will eliminate the need for giving things to poor people—enabling them to get what they need for themselves, teaching "self-reliance."

The Republicans quite rightly claim that simply giving things to poor people will keep them dependent and poor (they don't learn how to control their own lives). The Democrats quite rightly point out that simply demanding self-reliance ends up punishing people for being poor and creates callousness toward human suffering. Republicans assume that people work in order to maintain a viable economic system that's essential to everyone, and because of financial rewards and incentives. They assume that the healthy society is one in which the members compete for wealth and predominant positions or power. Democrats assume that people work to improve the quality of their lives outside the economic framework, and that the healthy society is one in which nobody has to labor overly long, under unpleasant or dangerous conditions, or in a state of social inferiority. At least that is my view of the "canonical" positions of the two parties. I speak, of course, as a time traveler from a different era.

It's impossible to agree on standards without agreeing on system concepts: the kind of society we live in and our own human natures. Simply

hurling the names of standards back and forth and claiming that they are good gets us nowhere. Even agreeing on the means of achieving standards requires a shared concept of human nature. Those who enjoy power and wealth quite rightly appreciate the advantages of these things; they advocate principles based on the assumption that everyone would be better off with power and wealth, and principles that will help those who already have power and wealth to keep them. Those who value other goals assume correctly that nobody is permanently better off with power and wealth unless everybody has them, and favor principles that spread the wealth even at the expense of those who lose out by accepting the principles.

When we speak of standards as shared principles, we tend to forget how little of this sharing there really is. The story of standards in human societies is a story of conflict, not sharing. This is true in all sizes of groups from the dating couple, through the family, through a whole country. Even when people say, in words, that they agree on a standard, they perceive it differently; even when they perceive it in more or less the same way, they differ on the reference level. We can agree that many poor people need immediate financial aid. Which people? How much? To be spent how? In whose Congressional district?

The other comment I have is more general. We tend to speak of standards in terms of their effects when they are shared, in terms of their roles as characteristics of a society, or in terms of what they do for social interaction. From the theoretical point of view, however, the questions are not just *what* standards are adopted and why they are adopted, but what a standard is, and *how* it can have any effect.

How does a standard influence the behavior of any individual? How does it get communicated? What has to happen inside an individual before the words describing a standard come to have meaning to that person? And what has to happen inside the person in order for any particular interpretation of such a description to attain the force of a reference condition? Without these processes internal to the individual, no standard can have either meaning or effect. We have to understand standards as they exist in and operate in a single person before we can understand how they work in a world populated by many persons.

Finally, we often speak of the advantages or influences that standards have in a society. I think that, very often, these advantages or influences are hypothetical—they're what *should* occur. But I doubt that such things very often *do* occur.

Rick Markers: Bill Powers says: "Finally, we often speak of the advantages or influences that standards have in a society. I think that very often, these advantages or influences are hypothetical—they're what *should* occur. But I doubt that such things very often *do* occur." This has

been my point all along—at least in terms of personal (and, to some extent, in terms of interpersonal) control. I see no way in which perceiving certain standards at certain reference levels can *necessarily* lead to successful control of any other perceptual variables—or intrinsic variables. Yet this is an article of faith for many people in society. I imagine that if one of us showed, quantitatively and experimentally, that this faith is not correct, he or she would soon be the victim of a holy war. Now that I think of it, if really good science on control of principles and system concepts were done, it is possible that the results would make the religious/political/scientific establishment take steps that would make the Catholic Church's treatment of Galileo look like a picnic.

One of the things that has particularly irritated me about the current political dialog about values (the one going on in the outside world—not on CSGnet) is that the people who are pushing "family values" most ardently are also the people who have most ardently pushed one of the most fundamental (and, I think, destructive) values of U.S. society—the value of *conflict* (also called *competition).* Every red-blooded American knows that competition is what makes for successful economies. The basic idea (as pink-blooded little me understands it) is that consumers are like judges at a beauty contest (a uniquely American event, itself). Producers (or goods and services) compete to win the patronage of the customers. This competition leads to better and better products from producers (in the sense that they are the products that best meet the customer's needs or wants).

This scenario has one little problem that only Americans with pink-tainted blood might ever even deign to point out; in competition like this, there are generally winners and losers. What happens to the losers? America doesn't like losers, so we ignore them or blame the loss on personal failings (not being a *real man* or a real *woman).* Pinko types like me, however, don't think that losers are just valueless trash; they are worthwhile control systems, with intrinsic reference signals of their own. I worry about the losers because societies with lots of them around tend to be very precarious—and have to take strong measures to make sure that the losers don't try to just take stuff from the winners.

I don't like the "value" attached to competition in this society. I like the "value" of cooperation and community. I think society's emphasis on the importance of "being #1" or "fighting to get to the top" is far worse than the lack of emphasis on "family values." But I doubt that Quayle and Bush will come out in favor of the value of "cooperation" and "community." Do I have bad standards? Is it wrong to dislike competition and to like cooperation?

I will admit that competition (conflict) can accelerate the development of technologies that might help the parties to the conflict "win." Thus, two companies making widgets might progress faster toward the goal

of making the "best" widget (the one that satisfies the market best) because they are in conflict (they have to keep improving the widget—the output of each system—or lose the conflict—have their market share of widgets become much lower than their reference).

I think it is this "good" result of competition that has impressed economists. But is this the only way to organize an economy that produces the widgets that we all need to control what we want to control? Must there be winners and losers in order to have an economy that meets the requirements of its members (the winners, anyway)? Can't we organize a society in which everybody is a winner (can control what they need and want to control)—and can't we do it without coercion (the approach that communism used)? It seems to me that the economies of some of the Scandinavian and Western European societies approach a nice compromise between capitalistic individualism and socialistic communalism. Why don't we learn from those economies?

Kent McClelland: Rick, although I agree with and indeed applaud your sentiments favoring cooperation over competition, I wonder whether you're making the choice sound a little too simple. An interesting book by Michael Billig and associates *(Ideological Dilemmas: A Social Psychology of Everyday Thinking,* London: Sage, 1988) has convinced me that such things are not a matter of either/or, at least not in our usual modes of thinking. The book traces the history of Enlightenment thought and shows how contradictory values are built into the public discourse on such issues. Racists, for instance, will typically preface their biased remarks with a disclaimer to the effect that they themselves aren't prejudiced against blacks, but you really can't get away from the fact that, etc., etc. I have no doubt that Bush and Quayle could come up with many heart-warming remarks about the value of community and not see any contradiction at all between that sort of rhetoric and their views on competition.

I think the question is how stable system concepts can come to be constructed from an amalgamation of values or principles that are often contradictory in practice. But maybe such mental and moral flexibility is necessary for us to maintain the perception that the world we observe is consistent with our preferred system concepts. As I believe Rick pointed out, a control system that was stuck with a single reference signal for a principle like honesty (or cooperation!) would be unable to vary its outputs to maintain control of perceptions of the next higher level, just like an arbitrary restriction to a single setting for arm position would cripple your physical control of bodily movements.

Bill Powers: Kent and Rick, cooperation and conflict are outcomes of a social interaction. If people's goals are aligned, there will be coopera-

tion, or at least non-interference. If they are not, there will be conflict and competition.

Competition arises in our society as a consequence of system concepts and principles. One of these concepts has to do with position in a social hierarchy. The idea of the superior person, with others being inferior, sets the stage in some people for a desire to be, or be acknowledged as being, at the top of this social hierarchy. As achievement of this goal requires a relative ranking of people, it is impossible for everyone in the society to achieve it. If even two people wish to be perceived as number one, a conflict must arise, because by definition only one person can be number one (or number anything). The existence of number one creates number two: number n implies number n -1. If one person wants to become a leader, followers must be found, and others who also want to be leader must be fended off, undermined, or otherwise prevented from succeeding. The striving for social position is a pernicious ill in our society, which accounts for a great many of its problems.

I've heard all of the arguments in favor of competition. I don't believe them. I don't think that people with contradictory goals accomplish anything but building up their muscles and cancelling the effects of someone else's muscles, leaving little effort available for real progress. I don't believe there is a "top" in the social hierarchy—I don't even believe there is a social hierarchy. And as long as I don't believe that, there is no social hierarchy for me. This doesn't endear me to people who want such a hierarchy to exist, but that's their problem. There's nothing I want from anyone that would make it worthwhile to play that game. Not even the privilege of living.

And I know for certain that when, in some microsociety, people manage to do without this concept of Number One, everything magically works better: shared goals are accomplished smoothly, easily, and with great pleasure. People get smarter, because they aren't wasting their time and effort trying to counteract what someone else is doing.

I haven't got this system concept worked out in any detail—talking about it too much tends to reduce it to procedures and slogans, anyway. But what I do understand of it, I want to sell. It defines the kind of world that I find worth living in. All I can do to create that world is to persuade others who will persuade others that it's worth a try.

"Leadership," it seems to me, is a role in a social hierarchy. It requires followers. It opens the door to competition and conflict ("I can lead better than he can, so follow me and not him"). The worst result, from my point of view and in my circumstances, is that followers learn from a leader how to follow, not how to explore, teach, and learn.

The attitude of followers toward leaders, in my experience, often tends to be one of admiration, deference, blind loyalty, and even hero

worship. It's the attitude of a child toward a favored adult. Many leaders like being on the receiving end of this attitude. It confers power, it allows the leader to indulge in egocentric thinking, it protects the leader from criticism and accountability. The leader can arrive too easily, with the connivance of the followers, at the idea that he or she makes fewer mistakes than ordinary people do. The leader can point to the support of the followers as a way of showing others, outside the group, that there must be something superior about the leader (so they would be better off becoming followers, too). Leaders are corrupted by their followers, and willingly.

Dag Forssell: The idea that a leader is defined as someone who has followers is indeed the predominant interpretation in our society. I believe it is an unfortunate one. It is not the only one available.

I can't conceive of a control system wanting to follow. What a control system wants is good system concepts to inspire good principles, so you can select effective programs, ..., so you can maintain your body chemistry. A control system is *designed to lead itself;* to satisfy its own purposes as it perceives them. "Purposeful Leadership," as I define it, is the development and communication of good information that allows every individual to lead himself /herself in full autonomy. It is a non-manipulative, non-coercive, non-violent approach.

With good information shared and internalized voluntarily, people will be aligned and will automatically cooperate on the mutual concerns.

Bruce Nevin: Dag says: "I can't conceive of a control system wanting to follow." Oh, come on, Dag! You can't mean that, can you? Aren't there many occasions when one control system wants to follow the lead of another control system? And is this in itself pernicious? (Though it can be abused—on both sides of the dyad, be it said! Nor does it end with childhood. Nor is it always childlike, though abuse of childrens' dependency does seem to result in many adolescents and adults coming to abhor and scorn it and fear exposure of it in themselves. One of the sure recipes for childishness.)

Have you ever taken a dance class?

Dag Forssell: OK, Bruce. A difficulty on this net is that anything can be and is taken so damn literally. You have a point, of course. I did learn to lead in dance once upon a time. It is important in ballroom dancing to give clear signals to your partner (follower) with a steady hand. The follower chooses to follow and concentrates on that.

Leadership is often understood to mean that you tell someone *what to do,* then they follow by doing what you tell them to do. This emphasis

on doing and instructions fits nicely in a cause-effect world.

I am trying to redefine and sell "Leadership" as the idea that if you want to lead, the *most effective* way is to offer good (a description of that Boss Reality that is as good as you can make it or negotiate it in open discussion) information for your "followers" to evaluate and make part of their own system concepts if they want to (understand, no conflicts with pre-existing concepts, relevant, etc.). Then you step out of the way and let the "followers" control to their hearts' content. You will not need to supervise or "control" their actions, because that is built in.

This form of leadership is inherently non-violent. Teaching it will not work if the top management in a company is coercive, as I perceive most to be. Therefore, the idea *must* be sold at the very top, to the very people who are used to insisting on results or else. I am counting on finding a few who will see it my way, but don't expect many. A few is all I need. Once the process is understood, the leadership /information can come from anyone in the organization.

Bruce Nevin: Following surely cannot mean producing the identical behavioral outputs. We know this because of the variability of behavioral outputs with respect to the reference signal. (Or with respect to the outcome, more or less equivalent depending on success of control.) Nor can it mean assuming the identical reference signals for identical (or equivalent) controlled perceptions. We know this because all the follower has to go by are the behavioral outputs of the leader, among other environmental variables, plus memory and imagination, of course, which are the means for projecting, anthropomorphizing, and so on, which we necessarily do all of the time.

There are two corresponding questions for the other member of the dyad: Can a control system want to lead another control system? Can a control system lead another control system?

From the existence of a large literature and a long history of 'leadership," it seems clear that a control system can want to lead another.

It seems to me clear that A can lead B only to the extent and in the manner that B wants to follow A. This is why virtually all of traditional thinking about leadership boils down to "motivation"—getting others to want to follow you. (Ditto for pedagogy)

Assume that B wants to follow A. The extent and manner depends on B's other goals. B can follow just in terms of proximity. This kind of following ranges from detailed mimicry (mirroring) to very slight correlations, such as B following A with his or her eyes.

Much of what we mean by "follow" is metaphorical, with this literal sense as a basis. We can easily identify the metaphor when we say B is "following A's argument" or "following A's line of thought." The metaphor is not so obvious, perhaps, when we talk of B following A in the

sense of coming to A for directions, going off and executing them, and coming back to A for more.

"Following directions" seems to mean to control one's perceptions so that they mimic ("follow") the perceptions that one imagines on hearing or reading the directions.

"Following A's argument" seems to involve imagining the argument for oneself and finding that the imagined line of argument corresponds with what A has said and is saying.

To paraphrase P. T. Barnum, some of the people want to follow all of the time, all of the people want to follow some of the time, but not all of the people want to follow all of the time.

If B is not confident and purposeful in a given situation, B might seek someone to follow until in a situation where B is more confident and purposeful. (Purposeful: has clear goals, is controlling for them without major conflict) We can discuss why this is so.

If I am B in such a situation, I will follow one who appears confident and purposeful rather than one who appears unconfident and irresolute. We can discuss why this is so.

Some people are unconfident and irresolute and conflicted in much of their waking experience. I suspect that many such people came to be so because of childhood experience with adults who emphasized conformity with external authority and arbitrary standards, enforced in punitive ways.

It can happen that such a person feels confident and purposeful in an institutionalized social context with clearly assigned roles and relationships of relative authority, in accord with standards established for those institutions. Such people can become "leaders" within that framework. They know "the system." They become very anxious outside it, and resist contradiction to it. I think that outside the system they fear unexpected punishment; my experience is that outside the system (that is, in circumstances in which they can no longer interpret their perceptions as within the familiar institutional context) they become unconfident and irresolute. They often despise indecision and lack of confidence. (Such people, by the way, are unlikely to be drawn to HPCT at this stage in its history. And this parallels the familiar left/right ideological dichotomy.)

I suggest that charisma depends in part upon the appearance of confidence and purposefulness. As you have suggested, Dag, this connects with sales and marketing, where the pumped-up appearance often outstrips the basis of confidence, and the real purposes are ulterior. But charisma can be genuine. When you're looking for the exit in a crowded waiting room, a person walking quickly in one direction with a suitcase has some charisma.

The ad hoc situational leadership and functional (not authoritarian)

hierarchies of anarchism depend upon this, especially in cases where the participants lack detailed knowledge of another's capacities. "You seem to know what you're doing. How do you think we can make this go?"

Now: Can a control system manipulate another control system? Can a control system exploit another control system? I believe these are some of the negative senses of "leadership" and "charisma" that you are resisting. Am I right? I think HPCT does not show that these do not exist. 1t only shows that they cannot work as intended. Social institutions can help people persist in being slow learners about this.

Dag Forssell: Bruce, when I said that "I can't conceive of a control system wanting to follow," I did not mean to be so literal. I meant that it is not the nature of a control system to "follow," whatever that is. I appreciate your post. Your restating my points, paraphrasing rather, is a very good thing. It shows me how my careless wording can be (mis-)interpreted. You are doing a good job of sorting out technical alternatives and aspects of "following."

I am resisting what I perceive to be extremely common stereotype interpretations of leadership and sales, where I sense an interpretation that leadership and sales are indeed "manipulation" and "exploitation." This I read into Bill's original refusal to lead and some comments about sales at past CSG conferences. In turn, this leads to an aversion to consider these major applications of HPCT. Still my perceptions, of course.

If you substitute "manipulate" with "inform" or "guide" or "enlighten" or "teach," and "exploit" with "mutual benefit," the substance of the interaction does not change from an HPCT point of view, but the emotional, stereotype flavor changes dramatically. We are still talking about leadership and sales and mutual economic advantage.

Certainly the members of this net want to sell HPCT to the world. Is this "manipulation" and "exploitation"? I would not label it that. But mention leadership and sales. What comes to mind? Some brutal, selfish "leader" on the one hand, and pusher of overpriced junk nobody wants or needs on the other.

These terms are among the unexamined "human pie slices"—system concepts from pre-HPCT days—that can benefit from some HPCT light. By looking closely at this, perhaps a way to sell HPCT can be found, vastly superior to the frustrating sales efforts in the psychological journals that are discussed on the net, but are not labeled as such. (These journals are a minuscule market compared to the rest of the world, and the one market where we know that HPCT is not welcome).

The way there is to forget about "manipulation" and "exploitation" and instead examine the best interest of and control processes in the

other autonomous control system, whether we call it follower or buyer. This done with full visibility to said follower and buyer, of course. There is no need to hide the interest and control processes of the leader or salesperson, either. The exchange of goods or services should benefit both parties. Otherwise we have reverse manipulation and exploitation.

Leadership and sales both can be honorable. HPCT can show how.

Rick Marken: PCT will always have a hard time. People just don't like to believe in autonomy for anyone but themselves. And they will apparently continue to wage war on autonomy even though the consequences of that war are precisely the opposite of what they hope to produce. I am speaking of the "war on drugs"—the greatest and most sustained crime creation program in history. Here is a clear case of trying to do, at a societal level, what we have agreed is useless on an individual level: society is trying to forcibly change the reference level of a controlled variable (drug usage)—trying to force it to zero for everyone. I object to this idiocy, not because I want to take drugs (the usual assumption about those who want to end this drug war idiocy), but because things I care about are endangered.

The only solution is to go up a level ("what do you folks really care about?") or have a police state (a temporary "solution," at best). It looks to me like a solid majority would choose the police state in a second.

If there is a fundamental postulate of PCT, it is that organisms are control systems. A functioning control system is able to make its perceptual experience match its references for that experience; I call this "autonomy"—the normal operation of a control system. Anything that prevents normal operation is the cause of a malfunction. Conflict is an example of a control system malfunction; conflict prevents autonomy—i.e., the ability to control.

The drug war is an example of control systems *in conflict. So* the drug war is an example of control systems that are *malfunctioning.* There is no moral judgment here; that would imply that I *like* the goals of one group (the drug warriors) better than I like those of another (the drug takers). In fact, I personally don't care for either of their goals, but that *is not* why I don't like the drug war. I don't like it because there is *conflict* between control systems; this conflict *might* have unpleasant side effects for me (I might get robbed by a druggie who has to pay high prices for highly abundant substances, or have my house broken into by an overzealous SWAT team that's off by a digit on the address of a crack house). But the chances of those side effects are fairly low. I really object because conflict prevents the functional operation of the control systems involved; neither party (drug warrior, drug taker) is able to function as a full-fledged hierarchical perceptual control system.

The conflict would be solved, of course, if the druggies decided to stop taking drugs *or* the warriors decided to stop fighting drug takers (and suppliers). Either approach would end the conflict and people could start functioning again. I favor a solution to this conflict based on the drug warriors changing, because they are the ones who created the conflict by trying to control other control systems. The other control systems (druggies) maintain the conflict by maintaining their references for the perception that the drug warriors want them to change. But somehow (and it's hard for me to articulate it without becoming moralistic) it seems to me that it's a lot easier for the drug warriors to stop controlling for what the druggies are controlling than it is for the druggies to change their own reference for what they are controlling.

The reason the drug warriors are the problem is because they must push against another control system in order to control the variable they want to control. The victim (the druggie) could (and did until the drug warrior came along) control the variable s/he is controlling without creating conflict in another control system at all. So one set of control systems (the warriors) are *creating* conflict by trying to inhibit the autonomy (not consciously, but that is what they are doing) of others. Since the warriors don't understand PCT, they are creating this malfunction out of ignorance. So I still have no moral complaint here. The drug war is just malfunction—producing idiocy (stupidity) that results from a failure to understand the nature of autonomy. So problems like the drug war can be solved, not by trying to articulate better moral principles, but simply by understanding how control systems work. A person who understands control theory simply shakes his/her head in dismay at drug warriors—just as a person who understands plumbing shakes his/her head in dismay at somebody pouring grease down a drain. Both are just watching people create malfunction.

Ed Ford: It seems to me that when two or more living control systems find themselves in the same environment, in order for them to live in harmony and cooperatively, they have to agree on a way things ought to be, a system of concepts, which are best expressed and set forth by agreeing to a set of standards upon which they base their choices as they attempt to find satisfaction while living together. (I see standards as synonymous with rules, criteria, principles, guidelines, etc.) Thus the needed harmony between levels of the hierarchy in social groups. As they live their lives, trying to satisfy their own individual goals, the choices they make, if based on agreed-to standards, will more than likely make it easier for them to live in harmony with each other.

In the order of nature, we first learn to follow standards as children at home and then, ultimately, to set our own. For us to live in harmony, we must always set rules while respecting the rights of others. Whether

at home, at school, at work, in an institutional setting, or just buying gasoline, we are constantly surrounded by standards and rules. Thus the need to learn to follow standards as well as to set our own standards while resolving our internal conflicts.

Any time I deal with anyone, whether in private practice or elsewhere, standards and rules are a part of life. Whether it is setting standards for the kind of spouse we want (thus to help us make a choice), or wanting to get along with a parent, or getting through school, or interacting properly on the net, etc., standards are a part of life. However, in order to help living control systems resolve their own internal conflicts and to teach them how to deal with their lives, there has to be a basic understanding of standards and rules and of how consequences and choices are integrated into the standards concept already established in the setting where they are being taught.

My experience over the years has taught me that there are tremendous differences in the understanding of the role of standards, the meaning and place of consequences and choices.

Rick Marken: I just don't get it, Ed. What does "teaching standards" have to do with a PCT-based view of human nature? What I get from PCT is the idea that nothing could be less important—the actual substance of a person's references for relationships, programs, categories, principles, "standards," etc. matters only in terms of how these satisfy higher level goals. The system should just be error-free—and this happens by having working (conflict-free) control systems. Of course, such systems will be setting the "right" references for perceptions like your "standards," but they are right from the perspective of the control systems (they combine appropriately with prevailing circumstances to achieve the higher level goals). What is at any time a "right" setting for a particular standard from the point of view of the control system might very well appear to be a wrong setting from the point of view of someone who "knows the right standards." I know that some of the people you are dealing with have interfered seriously with other people in their efforts to achieve their goals. So, obviously, your goal is to teach them to act without hurting others, i.e., "follow the rules." I think this is great, but you should be clear that this focuses your treatment strategy on getting a person to act in ways that are better *for* you—and, incidentally, for the person him/herself.

A person who wants to perceive him/herself as socially cooperative would be creating a big conflict for him/herself if, for some reason, the reference for a perception with a socially accepted reference (like wearing clothes in public) were changed to a different value. But I don't believe that there are *any* "standards" perceptions which, if controlled at a particular reference level, would be intrinsically internal-conflict

producing; conflict depends on what other perceptions a person is controlling and at what level. I think Ed believes that there are certain intrinsically intrapersonal-conflict-producing standards-perception reference settings.

I do believe that there are settings for references for standards perceptions that produce interpersonal conflict—there are *lots* of them. Such conflicts occur because carrying out the purpose tends to produce disturbances to *intrinsic* variables in the other person; there is a biological basis to much (but not all) interpersonal conflict.

Martin Taylor: Rick, when there is conflict, there might be reorganization, and as Bill has often pointed out, that reorganization will tend to drive the conflicting systems into less conflict. If I do not conform to your standards, we both experience conflict if you care enough to try to make my actions conform (you can't see what I am "doing," but you can see my actions), and if your efforts make me unable to satisfy some references. So, point 1, it is not just me who experiences conflict and might reorganize. You might, too.

If a community has developed /evolved a set of standards that results in low levels of conflict when everyone adopts those standards for their actions (again, not for what they are "doing"), the standards will be rather stable. They work, because whatever people are "doing," their actions permit them to control their percepts adequately. That's what is meant by low levels of conflict. If the "standards" don't have this effect, and people find that they experience high levels of error when acting according to the standards, some people will reorganize one way, some another (it's random, after all), and the standards will disintegrate, perhaps to re-form as a new set of standards that provide lower overall error rates. Sets of standards that lead to sustained high error levels in many people are not stable. So point 2 is that if many people adhere to standards, it is because those standards do not conflict with the ideal of low intrinsic error.

I agree that there probably are no standards that we could call "intrinsic," but there are probably some reference levels that cannot be components of stable community standards. These will not be found in the standards of viable communities. But sets of standards probably fit together in clusters that are stable as a group that can be taken into or left out of a total system of standards. Different sets of precepts based on the teachings of long-lasting religions probably form such groups. I would imagine that the number of such sets that could be stable is unlimited, but the societies of the world might have found only a few tens of them.

If an individual lives in a community with stable standards, but does not use them to set the relevant reference levels, that individual will

find conflict in many of his/her interactions with other members of the community, whereas the other members will find conflict only in interactions with the deviant. The deviant is more likely to reorganize than are the other members, and if the standard set is truly stable, this reorganization will continue until the deviant acts according to the standards and, at the same time, finds his/her reference levels generally attainable by non-deviant actions.

I suspect that most sets of social standards are not truly stable, and perhaps there are no possible sets of standards that lead always to zero error as a consequence of interactions. In a non-stable, or conflict-retaining, set of standards, all members of the community are liable to reorganization, and the standards themselves will drift in a way directly comparable to linguistic drift—and for the same reason. The result could be the breakoff of heretical groups, or a more or less unconscious shift of mores, or other shifts.

Serious problems arise when individuals who belong to different communities with incompatible standards have to interact. The incompatibility of the standards sets is defined by the existence of conflict when one individual uses one set and the other the other set. One or both must reorganize. When you have large numbers of individuals from each community meeting, then either one community will lose its standards (its "culture") to the other, or both will have to develop supplementary standards to deal with the interactions. That way lies stereotyping of members of "other" communities, but it might be a necessary way to handle the modern possibilities for world-wide interaction.

It's all based on the iterated interactions of individuals, and one-on-one reorganization based on the conflict that occurs.

Rick Marken: Martin says: "If I do not conform to your standards, we both experience conflict..." Not necessarily true. Here you are talking about interpersonal conflict; we "experience" it only in terms of the success (or lack thereof) of our efforts to control variables (in my case, perceiving you as conforming to my standards; in your case, perceiving no loss of control as a result of my efforts—beating you, starving you, locking you up, etc.—to get you to conform to my standards). If one person is a lot stronger than the other, s/he will "experience" no conflict at all in this conflict; s/he will just get the result they want. If both people are about equally strong, they will experience loss of control, i.e., error with respect to some variables they are trying to control. Of course, being people, each will also be able to perceive the cause of his/ her lack of control: the other person.

My problem with this whole analysis is just the emphasis on "standards"-setting as a basis for harmonious interactions in groups of

control systems. I think this is almost certainly a crock. As humans, we do happen to be able to perceive at the system level, but that doesn't mean that controlling perceptions at this level is any more important than controlling perceptions at other levels. Herds of animals, for example, work together just fine without agreement on (or ability to experience) standards, system concepts, principles, categories, or whatever. Most everyday conflicts between people are usually over control of perceptions that are at lower levels than "standards"—and people work them out just fine.

I think organisms in groups "get along" when there are a sufficient number of perceptual degrees of freedom (df) to be controlled—and sufficient environmental df to allow all members of the group to control their perceptions. This means that the organisms must be able to perceive the environment in a way that allows simultaneous solution of the perceptual df problem in the constraints of the environmental df. Tom Bourbon's studies of two people controlling the relative distance between lines on a screen contributes more to our understanding of what makes it possible for multiple control systems to "get along" than does all our blathering about standards-setting. Standards are just one thing people have to be able to control—no more or less important than controlling sensations, configurations, transitions, etc. When people can control their perceptions—and when each individual in a group can control his or her own perceptions—then there will be no interpersonal conflict. This is an achievable goal, but to get there, we have to look in the right place; *not* at figuring out what standards people should set, but at figuring out how to provide people—*all people*—with the degrees of freedom necessary to control their own perceptions. We already know how to do this, actually. PCT just shows *why* this is important: The ways to do it are (1) *population control* (to preserve the available df); (2) *education* (to learn about the available df for controlling our own perceptions—and how to control those perceptions more effectively).

People have tried to solve their problems by finding the right standards for centuries (from the beginning of recorded history)—it not only doesn't work, it is the cause of most of our intractable problems (nationalism, religious wars, etc.). I suggest that we approach the problem of interpersonal interaction from a PCT perspective; if people really are input control systems, then PCT should have some scientifically and practically useful things to say about how multiple control systems can get along without conflict. I think the answer is "degrees of freedom," not "standards."

Martin Taylor: Rick, I'm a bit confused. I knew this was foolhardy territory to get into, but I can't see what the discussion of interaction

procedures that evolve into conventions known as standards has to do with the system level. You seem to be implying throughout your response that the only place where standards exist is at the system level. I intend the term to apply at all levels relevant to interactions among people, and I think it applies probably more to actions than to behavior (using the PCT distinction that behavior is the control of one's own perceptions, whereas actions are not). Standards include greeting patterns, dress codes, thank-you notes for gifts, and all sorts of things for which the external appearance is what matters.

If I can act according to the standards of my community, and nevertheless control my perceptions with little error, I won't reorganize much, and I will continue to act according to the standards. If I don't act according to the standards and nevertheless am able to control my percepts, I won't reorganize. But in most cases, if the standards matter to many of the people with whom I interact, I will find that not acting according to the standards might impede my control (I might not get the job because I didn't wear a suit to the interview; I might not get a gift from Aunt Mabel because I didn't send a note thanking her for the last one), and I am likely to reorganize. When my reorganization leads me to act in such a way that I maintain control of my percepts, I will no longer reorganize.

Real community standards are those that tend to induce reorganization in people who don't act according to them. As I said before, their stability is determined by the degree of conflict occasioned on average in people who abide by them, because error will lead to reorganization, and if there is a set of standards not very different from the current set but that tend to lead to less error, then the community standards will drift in the direction of that set. The word "community" is diffuse here. It is clearly weighted by the probability of interacting with any particular person, so for most people, I suspect the standards one develops will be closest to those of the parents and older siblings, at least when interacting with them, though other sets of standards might be developed for interaction with others (such as the local gang).

As you can see, I don't think standards are anything that people control (or other herd animals, for that matter). They are the products of reorganization, not percepts. They are the ways that percepts can be controlled when other people are involved in the actions that together form the controlling behavior. All the same, I suppose that people can model desirable organizations, talk about them, and explicitly teach them to the young. But the problem here is how you teach any behavior deliberately. The "standard" you can talk about is a model or a simulation, not the result of a structural reorganization. "Standards" is the result, not the instigator. It is the manifestation of the dynamics of an uncontrolled interaction among control systems, not a prescription for

what should happen. It becomes a prescription by methods fully intelligible within classical PCT.

Rick Marken: Martin, the talk about standards is highly ambiguous—sometimes I think people are talking about reference levels and sometimes about perceptual variables. I thought we had clarified it earlier a bit—my conclusion was that Ed Ford (the main "standards" guy) uses the word "standards" to refer to "higher level perceptual variables"—types of perceptions that might be described by words like system concept, principle, value, belief, etc. He tries to help people set the "right" reference levels for these perceptions. So my reply to your post was really aimed at Ed—I just don't think control of higher order variables is any more important in social interactions than control of other perceptual variables.

You say: "Standards include greeting patterns, dress codes, thank-you notes for gifts, and all sorts of things for which the external appearance is what matters." So what you mean by standards is "perceptual variables that involve another person." Well, now we have another possible meaning for "standard." Why don't we just stick to the PCT-model terminology (and semantics)?

And you say: "As you can see, I don't think standards are anything that people control..." Boy, you've got me. In the quote above, it sounded like standards were social perceptions. Now they are something that can't be controlled. And yet people reorganize when controlling them produces conflict. So it must not be failure to control standards that is leading to reorganization. But the reorganization leads to new, stable standards. So standards are a perceptible (to Martin) side-effect of reorganizing to control perceptions that are not standards? In other words, people control perceptual variables; this can appear to an observer as a process of converging on social standards. Is this it? If so, I completely agree.

Martin Taylor: Rick says: "So what you mean by standards is 'perceptual variables that involve another person.' Well, now we have another possible meaning for 'standard.' Why don't we just stick to the PCT-model terminology (and semantics)?" I've been trying to stick very precisely to the PCT model, but I don't know of any standard terminology to handle what we are talking about.

The problem with any definition of "standard" is that it is something (let's not say what) that one person applies to the observable actions of another. A person might apply standards to himself or herself, but only as an observer, possibly in imagination, of his or her own actions.

Standards have a funny status. I cannot control your behavior, because I have no sensory information that allows me to perceive it. But I

can control my perception of your actions, in the same way I can control my perceptions of the inanimate world. To control perceptions of anything, I perform actions (not controlled; the product of all past reorganization), and if my reorganizations have been effective, my perceptual signals come closer to their reference levels. If I hold reference levels for my perceptions of your actions, the same applies. I act, and if the error signal does not decrease, there is a reasonable probability that I will reorganize. You, too. We both reorganize if controlling each other's actions is not proceeding successfully (and if it matters—i.e., if we are working at a reasonably high gain). Our mutual reorganization will probably wind up eventually in a situation where our perceptual errors are not too large. Then, each of us is acting according to the other's standards. This cannot happen if it causes a more-than-compensating increase in errors related to control of percepts outside the interaction. The most likely end-result is that most people in a community use much the same set of standards.

Naturally, the end-result of reorganizing through social interaction and the control of the actions (not the behaviors) of each other will be the existence of perceptual functions in each person that relate to patterns of actions in other people (and perhaps in themselves). Specific reference levels for these perceptions will be associated with the probability of low errors in other perceptual signals, and those reference levels might become the kind of "standards" that Rick was originally talking about.

Ed Ford: Rick says: "I just don't get it, Ed. What does 'teaching standards' have to do with the PCT-based view of human nature?" I am not teaching standards, but the intelligent evaluation and use of the ones people create for themselves. Or, I am trying to help people deal with the standards in the environment in which they find themselves to satisfy their own goals. An example would be helping a person to think through the best way to satisfy the goal of getting released from a lockup facility within the reality of his/her present environment.

When you're down in the trenches, you have to be very practical. People can only achieve their goals by establishing in their own mind criteria (standards, rules, guidelines) upon which they are going to base their decisions. You say the system should be "error-free, " and I agree, and then you go on to say that "this happens by having working (conflict-free) control systems," to which I agree. My question is: how do you help another system get to that point? When you set a reference for driving on the freeway, for establishing a closer relationship with a member of your family, for satisfying an employer or improving your job performance in a working environment, for employing a worker, or just for buying food at a grocery store, you surely do have standards or

criteria based on your references for the choices you're going to make to achieve your goal. And in order for you to function in the environment in which you find yourself in some of the above situations, you are going to have to become aware of what the agreed-to standards or rules are that others in that environment have agreed upon to live by so that you and they can function cooperatively. You can't see a reference for safe driving, but you can see stop signs and speed limits, and you are made aware when you get a driver's license of the various rules or standards for driving.

Obviously, my friend, you have never read my book *Freedom from Stress,* which goes into great detail explaining the relationship of standards, principles, or whatever you want to call them, to the other levels of the hierarchy of control, and how all of that understanding helps people control much more effectively and efficiently for references or goals. You just don't deal exclusively with the highest goal. And, more importantly, the various people with whom I work evidence a need for help in learning how to use their system more efficiently so that they can function more effectively and get what they want.

People are able not only to articulate, prioritize, and evaluate references, but also to set appropriate standards or rules or criteria that will help them reach their goals. Also, these rules or standards will then act as guides for the various choices they have to make *if* they have learned to use their systems properly. What I am trying to say is that you teach people how to use their own systems, to set their own goals, their own standards upon which they can make choices, because PCT teaches me that that is how the system is designed. This hierarchical system is highly interconnected, cross-connected, and interdependent; being able to satisfy goals often demands the awareness and evaluation of all of these various levels. And you know what? It all works.

I never, ever push people to act in a way that would be better for me. That is absolutely wrong. Please explain to me how this focuses my treatment strategy on my goals (except that of helping them to function more effectively and responsibly on their own). Have you ever seen me work with anyone or explain what I do through a role-play demonstration? I suggest that you read the role plays in chapters nine and ten in *Freedom from Stress.* If these people are a part of my life, a necessary part of the environment in which I attempt to live and work cooperatively with others (for example, at work or at home), I have to find out what their goals are, what they are planning to do, how they perceive things, so that I can deal with my life within the reality of the choices these others are making. In my counseling, it is the clients who are asking for help in learning how to deal with their world in such a way that they can satisfy their internal reference signals, including getting along with the people in their lives. It is these living control

systems who are asking for help. They are asking to be taught the skills of functioning more efficiently and to learn how to reach their goals with the least hassles. My goal is to help them with what they want. The last thing I want to do is impose my values or beliefs on them. They are going to have to deal with the consequences that are a result of the goals, standards, and choices they make. I teach them how to manipulate themselves, to ride their own bike, to make their own choices, to satisfy their own goals. To manipulate people in such a way as to get them to do what I want is totally against good, sound counseling and teaching, and totally against the PCT design. It is totally repugnant to everything I believe.

Rick Marken: Ed, you say: "People are able not only to articulate, prioritize, and evaluate references, but also to set appropriate standards or rules or criteria that will help them reach their goals." Well, you might be working down there where the rubber meets the road, but you are dealing with some enormously prescient people; apparently, they are able to know what the state of the world (disturbances) will be when they set their standards, rules, and criteria so that these will be appropriate and allow them to reach their goals. How can they do this when the disturbances they will actually encounter are unpredictable and, often, undetectable? I thought that PCT made it clear that the only appropriate settings for *any* references are those that, when the outputs resulting from these reference settings are combined with prevailing disturbances, produce the intended perceptual results. Thus, you *might* be able to direct a person's attention to the perceptual variables that *might* improve his/her ability to *control* other perceptual variables (the ones that he/she came in complaining that he/she could not control), but you cannot *possibly* know in advance the appropriate *settings* for the references for these variables.

And you say: "The last thing I want to do is impose my values or beliefs on them." I never meant to suggest that you did; I know you don't. I am just questioning the idea (at least as you describe it, and as I understand it) that one can help another person control better (which is what I imagine to be the goal of PCT therapy) by suggesting that there might be appropriate settings for one's references for any perceptual variables—rules, standards, principles, whatever. The "appropriate" setting of a reference must vary with circumstance if the intended result is to be produced. So it's not that I think you are trying to impose *your* values—it's that you are suggesting that there are values that are *right* for the client. This is correct, as far as it goes, but the *rightness of* that value is relative; it depends on what they are trying to achieve at a higher level (which I think you clearly understand), *and* it depends on prevailing (and unpredictably changing) circumstances—so that the

setting for the value that achieved the higher order goal at one time almost certainly won't do at another time. It is this latter aspect of "setting standards" that I don't hear reflected in your ideas.

Ed Ford: Rick, it is easy to say things in words in the theoretical realm. I wish you would use several examples. It would be much easier for me to understand and to deal with precisely what you are saying. In any event, I will try to respond to what you've said.

Any time we have a goal (reference), and we attempt to achieve this goal, the standards or criteria we set can be set for many reasons, many having to do with other references that interconnect or interrelate to the main reference we have. Whether I am trying to decide on which university to attend, or a young woman to marry, or to drive on a freeway, or to exercise, or to eat "healthier" foods, or where I want to live, or just to call a friend, all of these references are going to involve my making choices which are going to involve other important references. I might set some standards for the kind of woman I want to marry, but in my attempt to satisfy this goal, I might have to adjust my standards if my choices reject me. I might have certain standards for the way a happily married couple should live; obviously, those of us who are (or were) married have found a constant adjusting of standards very necessary to meet the "happy and warm, loving relationship" goal. As a vegetarian, I have very strict standards for what I eat, but I don't try to impose these standards on those with whom I live. Often, when asked to dine at the home of a friend, I am willing adjust some, but not all of my standards (I'll eat some cheese, but never meat or fish). I have been successful at maintaining a no-smoking policy in my house by asking visitors who must smoke to please do it outside.

All of the perceptual variables with which I am trying to deal can be controlled only by satisfying all of the other interconnected references, as well as the one I'm trying to satisfy. Standards can describe in specific terms the kinds of variables you are controlling for; they can also describe the outer limits you are willing to go to to reach or achieve your references, including how much disturbance you are willing to tolerate. Standards can also be tied to other references that are definitely interconnected or interrelated to the present references which you are trying to satisfy.

I am certainly not getting people to articulate "appropriate" settings for references. Rather, they articulate and then evaluate their present settings for their references and see if these particular settings are the most efficient or best settings and the best standards for helping them to reach their goals.

It is the person who has to discover the specifics of his/her conflict and the essential elements within the conflicting area that need to be

evaluated, including the references and priorities they've set, the standards and criteria they've established, and the choices they've made—and whether anything in this conflicted area needs to be changed or altered to reduce the conflict. You see, Rick, this is what I've been trying to do. It isn't the counselor who has to discover all of this, it is the person who is having the conflict who has to discover it. All the counselor is doing is helping or teaching the person to better use their system more efficiently and effectively.

You say: "So it's not that I think you are trying to impose *your* values—it's that you are suggesting that there are values that are *right* for the client." No, I am not. I am suggesting that the client find the standards or criteria that work best for him/her in the situation in which he/she finds himself/herself. My job is to teach them how to use their hierarchical systems, as suggested by PCT. When I ask them about their various levels, I am actually teaching them to think "level-wise" and to think about the interconnectedness and the interrelationships involved. From that, they are better able to articulate to themselves (and to me) the specifics of what is going on in their worlds. The more they understand how their living control systems work, the more they are able to use it to their own advantage. Therein lies the beauty of PCT, and especially the levels. When the levels are understood in light of how we function, they become much more useful to us, and our ability to manipulate our own system to our own advantage is enhanced, so that we can satisfy our own internal goals and thus eliminate or reduce conflict to a point where we can live with it.

You say "it depends on what they are trying to achieve at a higher level (which I think you clearly understand), *and* it depends on prevailing (and unpredictably changing) circumstances—so that the setting for the value that achieved the higher order goal at one time almost certainly won't do at another time. It is this latter aspect of 'setting standards' that I don't hear reflected in your ideas." I have nowhere suggested that once someone articulates their individual standards to me, they are locked into those standards. It is the ability of people to recognize and utilize these levels to their advantage to deal with their conflicts that is important. We all change standards all of the time. It is important that they first recognize the existence of the standards, the part they play in how we think, their usefulness in setting and achieving references through the choices they make. I don't care whether they change their standards or not. We all change standards all of the time. It's being able to change within the context of avoiding or reducing conflict that is critical.

When counseling (read teaching) others, it's not what I think, it's what *they* think—my job is to teach them to think by helping them to build confidence in their thinking ability. When they learn PCT and what

goes into making up a living control system, they have the road map. My job is to teach them how to use it. They have to learn to use it when I'm not around.

Rick Marken: Ed, your last post on standards cleared up a lot. I know that what you do is teach people to control their own lives more effectively; sometimes, I take issue with the way you describe some of your therapeutic goals. But your last description was excellent and quite compatible with my own sentiments about therapy (and they are just sentiments, since I would never be able to actually *do* therapy as skillfully as you do it).

Ed Ford: I've always believed that there should be no conflict between science and religion. I've recently found evidence of this. My granddaughter, Ruth, age five, from California, was visiting Hester and me. Hester had taken Ruth and her first cousin, Sally Ann, age four, who lives here in Phoenix, to a Christmas tree display, and on the return trip, the two children were in the back of the car, talking. The conversation went as follows:

Sally Ann: My immune system takes care of me.

Ruth: Well, my guardian angel takes care of me all the time.

Sally Ann: All the time?

Ruth: Yes, all the time. She's always with me, everywhere.

Sally Ann: Well, if you just leave the body alone it will take care of itself.

Ruth: Well, my guardian angel takes care of me all the time.

Sally Ann: Well, my immune system takes care of me.

They then went on to another subject.

Greg Williams: Ed, regardless of the potential and (I believe) actual conflicts between science and certain religious ideas, it appears that the major problem is religion vs. religion.

Rick Marken: I think Ed's young relatives were having a religious dispute—no science involved at all. Using scientific terms (like immune system) to describe the cause of perceptions (health) doesn't make it science. "Science" and "religion" are words that refer to lots of different perceptual variables. For me, the best definition of science was given by Bill Powers: "disciplined imagination"; we invent models (imagination) and then test to see if we observe in perception what the model does when "switched on" (discipline). This is a nice definition because it makes it easy to juxtapose it to what I think of as the essence of religion: "faithful imagination." The crux of the difference is the way you ultimately test whether your imaginings are "right." In science, the final arbiter is

God—i.e., the cause of one's perceptual experience (we call her Boss Reality). In religion, the final arbiter is People—perceptions are made to fit the faith (too often, violently). (I should note that, by this definition, much that is called "science" is not. Lysenkoism in the USSR is an example of religion—faith in inheritance of acquired characteristics—posing as science.) To my knowledge, there is no religion that would qualify, by this definition, as a science.

Everybody seems to be making up different stories about god(s) and what they say about the meaning of life and how we should behave in it. Seems like what we've got here are variable means to achieve a higher order result—varying across people, anyway. Wouldn't it be marvelous if we could learn to vary these means within one person—ourselves? Then a "Serb" could see that s/he is "Bosnian," too, and vice versa; an Israeli could see that s/he is Palestinian, a Catholic could see that s/he is Lutheran, an Atheist could see that s/he is Muslim, etc. The solution to the problem of religion (like the solution to any conflict resulting from inflexible goals) is not to eliminate the goal but to *rise* above it; PCT can help people get their consciousness to the level that is served by controlling religious perceptions. Once you get up there, you will see that religious goals are arbitrary—but useful for satisfying the needs of that higher level. When you get up there, you see that choosing a religion, ethnicity, nationality, etc., is just as useful (and arbitrary) as choosing a nice book to settle in with on a rainy day; sometimes you want a romance and sometimes only a thriller will do.

Dag Forssell: Rick says: "The crux of the difference is the way you ultimately test whether your imaginings are 'right.' In science, the final arbiter is God—i.e., the cause of one's perceptual experience (we call her Boss Reality). In religion, the final arbiter is People..." What a marvelous, lucid insight. And people can create and defend any system-concept religion they want, teach it, fight for it, and die for it. Witness the sorry spectacle in India. No Boss Reality arbiter there.

Rick also says: "To my knowledge, there is no religion that would qualify, by this definition, as a science." Some years ago, I attended Religious Science, Science of Mind (several times). They would take a text from the Bible, another from the Koran, a third from some Buddhist book. They suggested that there have been many good teachers, but that none is a God any more or less than you and 1. In every affirmation, song, and message, I was able to substitute the word God with "laws of nature." The one thing that was supernatural was "treatment." So I guess they fall down like all of the others.

Religion is more than a system concept, though. It is also a social club. There is where much of the strength and value comes from. And the coercion. If you don't say you believe in what we say we believe in, you

can't play in our sandbox. You might get ostracized from your family, friends, and community. Better go to church on Sunday.

Ed Ford: The purpose of my post about Ruth and Sally Ann was to share with friends a delightful and amusing interchange between two innocent children. What I enjoyed most about their conversation was how different their perceptions were and how they tolerated that difference. It was meant to be light and amusing and not a serious comment on or about religion.

Years ago, the foreword to a movie with a religious theme read as follows: "To those who believe, no explanation is necessary; to those who don't, no explanation is possible."

Rick says: "The solution to the problem of religion (like the solution to any conflict resulting from inflexible goals) is not to eliminate the goal but to *rise above* it; PCT can help people get their consciousness to the level that is served by controlling religious perceptions." It depends on whether, for a particular individual, his/her religion presents a conflict. For me, it doesn't. It is highly compatible with everything else that goes on in my head, including and especially PCT. Secondly, I see it at my highest level. It does give me satisfaction at the very highest (in terms of priorities) system concept I have, which, for me, is to be one with my Maker. Obviously, there are standards that flow from that system, and choices I make based on those standards.

And Rick says: "Once you get up there, you will see that religious goals are arbitrary—but useful for satisfying the needs of that higher level." I think it would depend on an individual's own perceptions of when he or she was up there, and how he or she perceived his or her "religious goals" and the other needs at the higher level. It is interesting when someone comes on the net and claims to understand PCT, and then says things that are obviously different from what we all have experienced through our own individual work. Even among those in the CSG, we all understand PCT according to how we have created it in our perceptual systems, from what we have done, read, observed around us, perceived as useful, experienced in creating ideas from it, perceived from building models based on it, etc. Many of us have understandings that others will never have. My wife's understanding of what it is like to have a child is quite different from mine, and, obviously, I'll never understand her experiential knowledge.

Rick also says: "When you get up there, you see that choosing a religion, ethnicity, nationality, etc., is just as useful (and arbitrary) as choosing a nice book to settle in with on a rainy day; sometimes you want a romance and sometimes only a thriller will do." Again, "you see" refers to what *you* perceive, not what *everyone will* perceive "when you get up there." I think that you presume a lot when you state that you under-

stand everyone's knowledge of how they will perceive an experience you've only had in terms of your own individual perception of your own created experience. Your knowledge of religion is limited to what you presently perceive, just like my knowledge of PCT (or anything else) is limited to only what I have built into my own perceptual system.

Dag says: "Some years ago, I attended Religious Science... They suggested that there have been many good teachers, but that none is a God any more or less than you and I." That was their perception, and you accepted that as yours. And I respect that. However, I don't happen to agree with that statement. That's my perception.

And Dag says: 'Religion is more than a system concept, though. It is also a social club." Again, that is your perception. And again, I don't happen to agree with that statement. My own particular religion is based on fact, not fiction. It is also based on 50 years of thought, study, research, and lots of reading.

Rick Marken: Ed quotes: "To those who believe, no explanation is necessary; to those who don't, no explanation is possible." Apparently, that's true. What I want to understand is *why* it is true. I want an explanation of believing, itself, whatever the beliefs themselves might be.

Ed says: "It depends on whether, for a particular individual, his/her religion presents a conflict." I didn't mean that religion is a problem because it creates intrapersonal conflict. I'm sure most devout people are quite unconflicted about their religious beliefs. The problem with religion (and other high-level goals of the same sort that become fixed—ethnicities, nationalities, etc.) is *interpersonal conflict. I* don't know if you've looked at your local newspaper lately, but mine is *filled* with violent, interpersonal conflicts over religions, nationality, ethnicity, etc. People are fighting their brains out to defend perfectly arbitrary goals; I consider this a problem-one that is so unnecessary that it is unbelievable. And the solution, of course, is for each person to be able to see that their own ethnic, religious, or national goals, though important to themselves, are perfectly arbitrary; that it's like arguing over whether cars should be driven on the left or right.

To me, religion is (as I have said before) just something that people do—like being a control theorist. PCT is trying to understand *all* of human behavior, and religion is certainly one of the most important (and troublesome) things that people do. It should be something we in PCT try desperately to understand.

Bill Powers: It will not be possible for science and religion to get together until both realize that neither is Revealed Truth, and that both are human ideas. Of course, that is precisely what both sides have been

rejecting since the start of science. One side points the finger at Nature, the other at God. Neither side, apparently, notices whose finger is doing the pointing.

Greg Williams: I'd like to add one additional observation to Bill's post regarding the possibility of the religious and scientific "sides" getting together. In several forms of religion, and some (at least historical) forms of science, accepting authority and having faith have been/are now valued more (sometimes *much* more) than adjusting beliefs in the light of new evidence. Modern science at least gives lip service to the idea that one's *own* finger should be doing the pointing, unencumbered by pleas or threats from others. But that is anathema to some modern religions. One reason that a discussion of "science vs. religion" is appropriate on CSGnet, in my opinion, is that the issue of self- vs. otherdetermination is right at the heart of what control theory has to say about the chances of an individual successfully coping in a disturbance-filled world. On the other hand, high-level reference signals (within a broad spectrum) appear to be very loosely coupled to day-to-day survival (assuming you aren't in a holy war, of course), so I don't feel much missionary zeal for going around begging folks to recant what they accept on authority. And if I did, I wouldn't rail against the beliefs themselves so much as against why they are held. As a general principle (based on PCT ideas), it would appear that breaking correcting loops (e.g., accepting dogma uncritically) is dysfunctional. Yet people do it all the time and seem none the worse for it. Of course, their neighbors might be *much* worse for it!

May your neighbors not be *extremely* dogmatic.

Rick Marken: I partly disagree with Bill's post on religion and science getting together. I think there are many scientists (the good ones) who understand that their models are human ideas, and that "nature"—the cauldron in which these ideas are tested—is just their own perceptions. I think there are also religionists who understand that religious models (myths) are human ideas, and that "spiritual experience"—the cauldron in which these ideas are tested—is just their own perceptions (human experience). I would venture to guess that there are far more scientists like the above than there are religionists. The reason for this is that implicit (or explicit) in most religions is the idea that you *must* believe that these ideas are Revealed Truth or else you, your people, or the human race are in deep trouble. I don't think this latter assumption is part of science—although I agree that many scientists act as *though* such an idea were part of the game; that is where science and religion become one—as Bill says, when their ideas (models, myths) are treated as revealed truth rather than human invention—invented for a *purpose*.

Bill Powers: There are two sides to religion. One of them, the good side, consists of the attempt to adopt and live out principles that make civilization possible. As most people never think about such things except in the context of a religion, one wonders what the world would be like without such formalized social systems of belief.

The bad side shows up because people have different religions. If those living under principles of love and tolerance could actually live up to those principles, all would be well. But aside from the fact that not all religions preach universal brotherhood, it doesn't seem possible for people to live up to their religious principles when those principles disagree with someone else's.

The basic reason, I think, is the assumption of supernatural origin of the religious principles. When you believe that you are in receipt of the word of God, directly or through an authorized dealer, there can be no tolerance for deviations. The word of God is absolute. This means that if a different group claims to have heard a different word, or a different interpretation of words, the other group must simply be wrong. Every religious group must feel this way about every other group, no matter what they say. Very quickly, this comes down to the choice of converting the other group to the true belief ("saving" them), isolating the other group, or eliminating the other group.

Each group, of course, must resist all attempts by the other groups to evangelize, because succumbing would be going against the word of God. The loop gain, with respect to adhering to the word of the Infinite, must be infinite. This means that even minor differences of doctrine can lead to maximum conflict.

All that saves us from continuous violent confrontation between religions is that very few people are actually as religious as they think they are, or claim to be.

The greatest mystery of the human mind, in my view, is the phenomenon of Belief. Nazis are easy to deal with, because their beliefs are threatening to our physical safety, and we can flatly reject them. But what about other belief systems, invented and accepted apparently at random? Is the human mind just naturally susceptible to any belief that comes around, no matter how childish and full of holes? Is there something about our highest levels of organization that demands some belief, any belief, to fill the vacuum?

It seems to me that before we can have anything approaching sanity on our planet, we must begin to understand how belief systems get formed and how to keep them from overpowering people—how to leave a little freedom of belief, so that knowledge about the *whole* world of experience can play a part in forming belief systems. I haven't the slightest idea of how to do that, except by continuing to point out that different people believe different things, a fact that ought to give

anyone pause who is convinced that his/her own belief system is the only right one.

Or is this a level at which we are all helpless, including me?

Rick Marken: Bill says: "The greatest mystery of the human mind, in my view, is the phenomenon of Belief." I agree. We should explore this from a PCT perspective. The problem, of course, is that, when it comes to many of one's own beliefs, they are not treated as beliefs, but as knowledge. I think many of our most tenacious intrapersonal and *interpersonal conflicts* are the result of controlling perceptions based more on beliefs (replayed reference signals) than on Boss Reality.

I think it would be worthwhile to say what beliefs are in the context of the PCT model—to describe examples of the everyday beliefs that people are walking around with (from the divine, like religious beliefs, to the profane, like beliefs about the "right" foods to eat); also, it would be nice to discuss the difference (from a PCT perspective) between belief and knowledge. I know this is a difficult discussion to have, precisely because beliefs are so important to people. With Bill, I ask, "*Why* is this so?" Why do people fight to prove that what they do not know is so? There must be a reason that this species has been willing to persecute itself for millennia over fantasies. It must be an aspect of our nature as control systems. What is it? I think that this could be a very satisfying (and even therapeutic) investigation.

Bill asks: "Or is this a level at which we are all helpless, including me?" No. I think people, like you (and me?), who are willing to consider the possibility that *anything* we think could be just a belief and, more importantly, are willing to wonder what a belief is, are not helpless victims of our beliefs (at least, when we are able to keep our awareness "above" the levels that create those beliefs—something that I don't do nearly as often as I would like). I think it requires some effort to defeat some of the insidious consequences of belief, but it can be done, I think.

Ed Ford: It seems to me that belief systems are formed by living control systems as they try to establish harmony within themselves as a result of their attempts to find satisfying experiences from the environment in which they find themselves. The choices we make and the standards we've set ultimately evolve into systems of ideas, or the way we think things ought to be. I think this harmony, this internal peace or internal integrity, is what the living control system is continually striving toward. Obviously, our knowledge of what's available is limited by our perception of the environment in which we find ourselves, plus what becomes available to us through reorganization. What we create out of what we perceive is what ultimately becomes what we are.

I think humans tend to accept the system concepts of those who they perceive love them and whom they love or admire. Obviously, if there are internal peace and harmony where we live, then the prevailing system concepts of our parents/friends are most likely to be perceived as acceptable. Those systems are ultimately tested when children (and later adults) are faced with choices which are in conflict with the prevailing or accepted system concepts. But to me, the ultimate test of a system concept is that first it brings internal harmony or peace to the person.

I don't believe a belief or value system (system-concept level) overpowers a person. I believe many people choose systems and elements of those systems and create their own standards from how they perceive those systems to justify the choices they're making, in their attempts to find that elusive peace and harmony that all living control systems are trying to establish.

When a person harms another living control system, his/her system concept is brought into disrepute. And this shouldn't be. I don't think it's right to blame Christianity for the acts of those who, claiming to be Christians, do harm to others, any more than it's fair to blame any system of ideas on those who claim to be adherents, but who go about harming others.

The second important test of any system concept is the respect shown to those "who don't belong, who don't believe." Therein lies the critical test of any system of beliefs, namely, that everyone is shown respect, as having value as a person. That, to me, is the real test of a valid system of beliefs. If from a system concept I am able to establish standards and make choices that bring me the internal peace and harmony within my system *and* at the same time that system concept leads me to see value in others and respect their right to make choices, then the system has value. In short, when we harm others, we harm ourselves, and in the process the very harmony and peace we are seeking are lost.

When a person is in conflict and uses a system concept to justify actions which bring harm to others, I don't think the system concept is wrong, I think the person is wrong. And I don't think the belief system overpowered them, they merely used the system "to justify their own means." I think people tend to overpower themselves by setting impossible standards or goals, by trying to change things over which they have no control, or by making ineffective choices in a desperate attempt to bring harmony or peace to their system.

Because I'm a living control system by design, my system concepts are very unique to me. No one quite perceives things the way I do. And I think the test for whether our systems of beliefs are valid are our own internal harmony and peace, and the respect and value we assign to others.

Rick Marken: Ed says: "I don't believe a belief or value system (system-concept level) overpowers a person." A belief and a system concept are not the same thing. A belief in PCT (I think) is an imagined perception: this means that beliefs can occur at any level of the hierarchy (except for the lowest); we can believe that the sky is blue (sensation), that it will rain (fluid transitions?), that we're loved (relationship), etc. We can also have beliefs that are system-concept-level perceptions —I can believe that I am a control theorist.

Beliefs (by my definition) can also differ in terms of one's ability to produce or experience them as perceptions (rather than just as imaginations). I believe my car is in the lot, and I can produce that perception; I believe that Mozart was the means by which God spoke to humanity, but I can't produce that perception (I can certainly produce the imagination).

Our ability to "believe" is, I think, one of the things that makes life fun; it makes it possible to be entertained by stories, plays, and such. I think it also makes life a bit more tolerable (as Ed said, it helps us "find that elusive peace and harmony that all living control systems are trying to establish"). It does this by "filling in" the unachieved aspects of the perceptions we are controlling; we believe that we are "loved," for example—and we create a perception that is based mostly on Boss Reality but that is "filled in" a bit by belief (imagination) so that our control seems a bit better than it might actually be.

But you can see that what is good about belief is what could also make it a problem; belief makes stories fun because we treat the imaginations as though they were "real" perceptions; but what happens when we forget that they *are not and never were real perceptions?* We get what we see—people willing to die or kill to control for imagined perceptions.

I think it is interesting that when the "filling in" done by belief gets to be a bigger part of perception than the part constrained by Boss Reality, we call that "insanity." But when the "filling in" is *total*—so that there is no constraint of Boss Reality—just belief based on made-up stories (the Bible, the Koran, etc.), we (some of us) call that "wisdom" I suggest that we call it what it is: "total insanity."

Page 36 intentionally left blank

The Blind Men and the Elephant: Three Perspectives on the Phenomenon of Control

Richard S. Marken
10459 Holman Avenue, Los Angeles, CA 90024

Abstract

Behavior has been described as a response to stimulation, an output controlled by reinforcement contingencies, and an observable result of cognitive processes. It seems as if these are descriptions of three different phenomena, but they are actually descriptions of three different aspects of the same phenomenon: control. Control is like the proverbial elephant studied by the three blind men; what one concludes about it, and how one tries to explain it, depends on where one stands. I suggest that the best place to stand is where one has a view of the *whole* phenomenon, be it elephant or control.

Introduction

The behavior of living organisms (and some artifacts) is characterized by the production of consistent results in an unpredictably changing environment, a phenomenon known as control (Marken, 1988). Control can be as simple as maintaining one's balance on uneven terrain, or as complex as maintaining one's self-esteem in a dysfunctional family. Control is a pervasive aspect of all behavior, yet it has gone virtually unnoticed in psychology. What *has* been noticed is that behavior appears to be a response to stimulation, an output controlled by reinforcement contingencies, or an observable result of cognitive processes. Each of these appearances is what would be expected if people were looking at control from different perspectives. The situation is similar to that of the three blind men who were asked to describe an elephant. The one near the tail described the elephant as a snake, the one near the leg described it as a tree trunk, and the one near the side described it as a wall. Each description gives an accurate picture of some aspects of the elephant, but a false picture of the elephant as a whole. If behavior involves control, then psychology has given an accurate picture of some aspects of behavior, but a false picture of behavior as a whole.

The basic requirement for control is that an organism be in a negative-feedback situation with respect to its environment. A negative-feedback situation exists when an organism's response to sensory input reduces the tendency of that input to elicit further responding. Negative feedback implies a closed-loop relationship between organism and environment; sensory input causes responding that influences the sensory cause of that responding, as shown in Fig. 1. It is hard to imagine an organism that does not exist in such a closed-loop situation, because all organisms are built in such a way that what they do affects what they sense. Eyes, for example, are located on a head that moves, so that what the eyes see depends on what the head does. To the extent that what the head does depends on what the eyes see (such as when the head turns in response to an attractive passerby), there is a closed loop; sensory input causes responding (head movement), which affects the cause of responding (sensory input). The feedback in this loop must be negative, because behavior is typically stable (organisms do not normally exhibit the "runaway" behavior that characterizes positive-feedback loops, such as the feedback from a microphone that amplifies its own output).

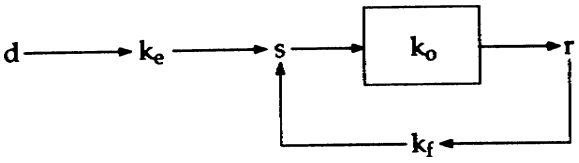

Figure 1. Closed-loop feedback relationship between an organism, represented by the rectangle, and its environment, represented by the arrows outside of the rectangle. A sensory variable, s, influences responding, r, via the organism function, k_o. Responding influences the sensory variable via the feedback function, k_f. The sensory variable is also influenced by an environmental variable, d, via the environmental function, k_e.

The fact that organisms exist in closed negative-feedback loops means that two simultaneous equations are needed to describe their relationship to the environment. These are given as equations (1) and (2) below. The terms in these equations are summarized here for reference in the discussion that follows:

s = sensory variable;
r = response variable;
s* = reference value for sensory variable, such that r = 0 when s = s*
d = environmental variable;
k_o = organism function relating sensory variable to response variable;
k_e = environmental function relating environmental variable to sensory variable; and
k_f = feedback function relating response variable to sensory variable.

For simplicity, I assume that all functions are linear, and that all variables are measured in the same units.

Equation (1) describes the effect of sensory input on responding, so that

(1) $r = k_o(s^* - s)$.

This equation says that responding, r, is a linear function of sensory input, s. The sensory input is expressed as a deviation from the value of input, s*, that produces no responding; s* defines the zero point of the sensory input. Equation (2) describes the effect of responding on sensory input. For simplicity, I assume that responding, r, adds to the effect of the environment, d, so that:

(2) $s = k_f r + k_e d$.

The variables r and d have independent (additive) effects on the sensory input, s. The nature of the environmental effect on sensory input is determined by the environmental function, k_e. The feedback effect of responding on the sensory cause of that responding is determined by the feedback function, k_f.

Equations (1) and (2) must be solved as a simultaneous pair in order to determine the relationship between stimulus and response variables in the closed loop (see Appendix, below). The result is

(3) $r = (1 / ((1 / k_o) + k_f))s^* - (k_e / ((1 / k_o) + k_f))d$.

Equation (3) can be simplified by noting that the organism function, k_o, transforms a small amount of sensory energy into a large amount of response energy (such as when a pattern of light on the retina is transformed into the forces that move the head). In control engineering, k_o is called "system amplification" or "gain," which can be quite a large number.

With sufficient amplification (such that k_o approaches infinity), the $1/k_o$ terms in equation (3) approach zero, so equation (3) reduces to

(4) $r = s^* / k_f - (k_e / k_f)d$.

Equation (4) is an input-output equation that describes the relationship between environmental (stimulus) and response variables when an organism is in a closed-loop, negative-feedback situation with respect to its environment. The result of being in such a situation is that the organism acts to keep its sensory input equal to s*, which is called the reference value of the input. Equation (4) shows that the organism does this by varying responses, r, to compensate for variations in the environment, d, that would tend to move sensory input away from the reference value; this process is called control.

Three Views of Control

All variables in equation (4), with the possible exception of s*, are readily observable when an organism is engaged in the process of control. The environmental variable, d, is seen as a stimulus, such as a light or sound. The response variable, r, is some measurable result of an organism's actions, such as bar pressing or speaking. The reference value for sensory input, s*, is difficult to detect because an observer cannot see what an organism is sensing. But s* is the central feature of control, since everything an organism does is aimed at keeping its sensory inputs at reference values. The value of s* can be constant or variable, its value at any instant being determined by properties of the organism itself.

Because reference values are difficult to detect, it will not be obvious to an observer that an organism is engaged in the process of control. What will be obvious is that certain variables, particularly the environmental and response variables and the relationship between them, will behave as described by equation (4). Thus, equation (4) can be used to show how control appears to someone who does not know that it is occurring. It turns out that there are three dearly different ways of looking at control, depending on which aspect of the behavior described by equation (4) one attends to.

The stimulus-response view. This view of control sees behavior as a direct or indirect result of input stimulation. An example of stimulus-response behavior is the so-called "pupillary reflex," where changes in a stimulus variable (illumination level) lead to changes in a response variable (pupil size). The stimulus-response view is the basis of several current approaches to understanding behavior, such as the "synergistic" or "coordinative structure" theory of motor coordination. Warren, Young, and Lee (1986), for example, describe a synergistic model of running in which "vertical impulse is directly modulated by the optical variable..." (p. 264). The behavior of running is seen in stimulus-response terms; an optical stimulus

variable determines ("modulates") the value of a response variable (vertical impulse). The stimulus-response view is also the basis of a recent theory of attention (Cohen, Dunbar, and McClelland, 1991) in which connections between printed-word stimuli and verbal responses in the Stroop effect are modulated by connections in a neural network.

Equation (4) shows that behavior will look like a stimulus-response process when the reference value for sensory input, s*, is a constant. If s* is zero, then responding is related to environmental stimuli as follows:

(5) $r = -(k_e / k_f)d$.

Equation (5) shows that, when there is a fixed reference level for sensory input, it will look to an observer of behavior as though variations in an environmental stimulus, d, cause variations in a response, r. This is what one sees in the pupillary reflex, where pupil size, r, is proportional to illumination level, d. Of course, this relationship between pupil size and illumination level is precisely what is required to keep a sensory variable (sensed illumination) at a fixed reference value (s* = constant).

When looking at an apparent relationship between stimulus and response, one's inclination is to assume that the nature of that relationship depends on characteristics of the organism. Equation (5) shows, however, that when an organism is engaged in control, this relationship depends only on characteristics of the environment (the functions k_e and k_f); the organism function relating sensory input to response output, k_o, is rendered completely invisible by the negative-feedback loop. This characteristic of the process of control has been called the "behavioral illusion" (Powers, 1978).

The reinforcement view. This view of control sees behavior as an output that is shaped by contingencies of reinforcement. A reinforcement contingency is a rule that relates outputs (like bar presses) to inputs (reinforcements); in equation (4), this contingency is represented by the feedback function that relates responses to sensory inputs, k_f. The reinforcement view is the basis of at least one influential theory of generalization and discrimination (Shepard, 1987). In a connectionist implementation of the theory, a reinforcement contingency is used to shape the formation of generalization gradients (Shepard, 1990). The reinforcement view is also the basis of modern theories of operant behavior. According to Domjan (1987), the contemporary perspective on operant behavior focuses on how contingencies "restrict freedom of action and... create redistributions of various types of activities" (p. 562). In other words, contingencies shape (redistribute) responses (activities).

Equation (4) shows that it will look as though contingencies (the feedback function) control responses when s*, d, and k_e are constants, as they are in typical operant conditioning experiments. In these experiments, s* is the organism's reference value for the sensory effects of the reinforcement. The environmental variable, d, is the reinforcement, which, if it is food, is typically a constant size and weight. The sensory effect of a reinforcement can be assumed to be directly proportional to its size and weight, making k_e -1. So, equation (4) can be re-written as

(6) $r = S^* / k_f - D / k_f$,

where S* is the constant reference value for sensed reinforcement, and D is the constant value of the reinforcement itself.

The only variable in equation (6) is the feedback function, k_f, which defines the contingencies of reinforcement. One simple contingency is called the "ratio schedule," in which the organism receives a reinforcement only after a certain number of responses. The ratio corresponds to the function k_f in equation (6). When the ratio is not too demanding, it is found that increases in the ratio lead to increased responding. More demanding ratios produce the opposite result; increases in the ratio lead to decreased responding (Staddon, 1979). Either of these results can be produced by manipulating the relative values of S* and D in equation (6). The important point, however, is that the apparent dependence of responding on the feedback function, k_f, is predicted by equation (6). To an observer, it will look like behavior (responding) is controlled by contingencies of reinforcement. In fact, the relationship between behavior and reinforcement contingencies exists because the organism is controlling sensed reinforcement; responding varies appropriately to compensate for changes in the reinforcement contingency, so that sensed reinforcement is kept at a constant reference value, S*.

The cognitive view. This view of control sees behavior as a reflection or result of mental plans or programs. This kind of behavior is seen when people produce complex responses (such as spoken sentences, clever chess moves, or canny investment decisions) apparently spontaneously; there is often no visible stimulus or reinforcement contingency that can be seen as the cause of this behavior. The cognitive view is the basis of numerous psychological theories that propose mental algorithms to explain the appearance of cognitive behavior. Examples of such theories are the ACT (Anderson, 1983) and SOAR (Newell, 1990) models of cognition, and hierarchical models of the generation of movement sequences (Rosenbaum, Kerry, & Derr, 1983).

Cognitive behavior is most obvious when environmental factors (such as stimulus variables and environmental and feedback functions)

are held constant. When this is the case, equation (4) becomes

(7) $r = s^*/F + K$,

where F is the constant feedback function, and $K = (k_e / k_f)(d)$, a constant. Since s* is typically invisible, equation (7) shows that there will appear to be no obvious environmental correlate of cognitive behavior. An observer is likely to conclude that variations in r are the result of mental processes-and, indeed, they are. But it is actually variations in s*, not r, that are caused by these processes; variations in r are the means used to get sensory inputs equal to s*. Thus, chess moves are made to keep some sensed aspect of the game at its reference value. When the environment is constant, r (the moves) might be a fair reflection of changes in the reference value for sensory input. However, under normal circumstances, r is only indirectly related to s*, variations in r being mainly used to compensate for variations in the environment that would tend to move sensory input from the reference value, s*.

Looking at the Whole Elephant

The blind men never got a chance to see the whole elephant, but if they had, they would have instantly understood why it seemed like a snake to one, a tree trunk to the second, and a wall to the third. Psychologists, however, can take a look at control and see why the appearance of behavior differs, depending on one's perspective. What is common to the three views of behavior discussed in this paper is the reference for the value of sensory input, s*. Organisms behave in order to keep sensory inputs at these reference values (Powers, 1989). They respond to stimulation in order to keep the sensory consequences of this stimulation from moving away from the reference value, so it appears that stimuli cause responses. They adjust to changes in reinforcement contingencies by responding as needed in order to keep the sensory consequences of reinforcement at the reference value, so it appears that contingencies control responding. And they change their responding in order to make sensory input track a changing reference value for that input, so it appears that responding is spontaneous.

What appear to be three very different ways of describing behavior can now be seen as legitimate ways of describing different aspects of one phenomenon—control. Each is just a different way of describing what an organism must do to keep its sensory inputs at their reference values. Indeed, once one understands that the appearances called "behavior" are the visible consequences of an organism's efforts to control its sensory inputs, the problem of explaining behavior changes completely, from an attempt to build models that simulate the appearance of

behavior (stimulus-response, reinforcement, or cognitive) to an attempt to build models that control the same sensory inputs as those controlled by real organisms. In order to build the latter type of model, it is necessary to learn what sensory variables are actually being controlled by organisms. This type of investigation cannot be done by simply looking at the appearance of behavior. Methods based on control theory can be used to test which sensory variables an organism might be controlling at any time (Marken, 1992). These methods make it possible to take off the blindfolds and see the whole elephant—the phenomenon of control.

Appendix

Given the two system equations

(1) $r = k_o(s^* - s)$ and (2) $s = k_f r + k_e d$,

we want to solve for r as a function of s. First, substitute equation (2) for s in equation (1) to obtain

(A.1) $r = k_o(s^* - (k_f r + k_e d))$,

which expands to

(A.2) $r = k_o s^* - k_o k_f r - k_o k_e d$.

Move all terms with r to the left side of equation (A.2) to obtain

(A.3) $r + k_o k_f r = k_o s^* - k_o k_e d$.

Factor out r on the left side of equation (A.3) to obtain

(A.4) $r(1 + k_o k_f) = k_o s^* - k_o k_e d$.

Divide both sides of equation (A.4) by $(1 + k_o k_f)$ to obtain

(A.5) $r = (k_o / (1 + k_o k_f))s^* - (k_o k_e / (1 + k_o k_f))d$.

Finally, divide k_o out of the numerators on the right side of (A.5) to get equation (3):

(3) $r = (1 / ((1 / k_o) + k_f))s^* - (k / ((l / k_o) + k_f))d$.

Acknowledgments

I would like to acknowledge the earlier ideas of Wayne A. Hershberger and W. Thomas Bourbon on this topic, which provided inspiration for the integrative analysis presented in this paper.

References

Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.

Cohen, J. D., Dunbar, K., & McClelland, J. L. (1991). On the control of automatic processes: A parallel distributed processing account of the Stroop effect. *Psychological Review,* 97, 332-361.

Domjan, M. (1987). Animal learning comes of age. *American Psychologist*, 42, 556-564.

Marken, R. S. (1988). The nature of behavior: Control as fact and theory. *Behavioral Science*, 33,196-206.

Marken, R. S. (1992). *Mind readings: Experimental studies of purpose*. Gravel Switch, KY: Control Systems Group.

Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.

Powers, W. T. (1978). Quantitative analysis of purposive systems: Some spadework at the foundations of scientific psychology. *Psychological Review*, 85, 417-435.

Powers, W. T. (1989). *Living control systems: Selected papers of William T. Powers*. Gravel Switch, KY: Control Systems Group.

Rosenbaum, D. A., Kerry, S., & Derr, M. A. (1983). Hierarchical control of rapid movement sequences. *Journal of Experimental Psychology: Human Perception and Performance*, 9, 86-102.

Shepard, R N. (1987). Toward a universal law of generalization for psychological science. *Science*, 237, 1317-1324.

Shepard, R. N. (1990). Neural nets for generalization and classification: Comment on Staddon and Reid (1990). *Psychological Review*, 97, 579-580.

Staddon, J. E. R. (1979). Operant behavior as adaptation to constraint. *Journal of Experimental Psychology: General,* 108, 48-67.

Warren, W. H. Jr., Young, D. S., & Lee, D. N. (1986). Visual control of step length during running over irregular terrain. *Journal of Experimental Psychology: Human Perception and Performance, 12,* 259-266.

Models and Their Worlds

W. Thomas Bourbon
(Department of Neurosurgery, The University of Texas Medical School-Houston, 6431 Fannin, Suite 148, Houston, TX 77030)

William T. Powers
(73 Ridge Place, CR 510, Durango, CO 81301)

Abstract

Many seemingly plausible models of behavior demand implausible models of the physical world in which behavior occurs. We used quantitative simulations of a person's performance on a simple task to compare the models of causality and of how the world works in three theories of behavior: stimulus-response, cognitive, and control-theoretic. Our results demonstrate that if organisms in fact functioned like the first two models, they could survive only in implausibly stable worlds; if like the third, they could survive in a changeable world. Organisms inhabit a changeable world that does not satisfy the demands of popular behavioral theories. For the sciences of behavior, the implications are clear: either cling to theories that do not mesh with knowledge of how the world works, or abandon many cherished notions about how and why behavior happens in favor of models that deal adequately with change.

Models and Their Worlds

The question usually addressed by behavioral theorists is "Why do organisms behave the way they do?" One group answers "Because the world outside them is the way it is"; another group answers "Because the minds or brains inside them are the way they are." In either case, behavior is at the end of a linear sequence of cause and effect, a consequence of antecedent stimuli from the environment or antecedent commands from the mind or brain. As an alternative, one can propose that organisms behave to control what happens to them. In the process, their actions affect the world outside of them. 'Why is the world the way it is? Partly because organisms behave the way they do."

"The world" is the part of the surroundings on which an organism can act, and which, in turn, affects the organism. Every statement about

the antecedents or consequences of behavior either includes or implies notions about how the world operates. Every theory of behavior is, in part, a theory about the world in which behavior occurs.

In this paper, we reduce three models of behavior to elemental form to identify and test their ideas about causality. Two models represent core assumptions in most popular theories; the third is the model from perceptual control theory (PCT). We require each model to simulate and predict the same behavioral events that occur when a person performs a simple task, but we go a step further. For each model, we determine whether its implications about how the world and behavior affect one another are reasonable and true to what is known about the physical world.

Three Models

For convenience, we call the two popular models the "stimulus-response" (S-R) model and the "cognitive" model. Our simple versions of these models are not intended to represent, in detail, any specific variations on those two themes, but we believe they faithfully represent core assumptions about causality embraced in those themes. Our method of testing requires that each model predict moment-by-moment values of several continuous environmental variables, a challenge to which behavioristic and cognitive models are rarely subjected; hence, simple computational versions of those models are not readily available, and we constructed our own. Anyone who rejects our versions of those theories should identify acceptable versions and then require their models to duplicate the quantitative results we report here.

The stimulus-response model. Our S-R model represents all theories that say external influences determine behavior. Such models sometimes (but by no means always) recognize that motor actions produce environmental consequences, but all insist that action is a dependent variable. A behavioral episode begins with an independent antecedent (stimulus, context, event, occasion, relationship, or treatment), followed (in some theories) by an effect on the organism, then (in all theories) a behavior as a dependent variable, and finally the consequences of that behavior. Environmental consequences of action simply follow from what the environment did to the organism; if any consequences of action modify subsequent influences on the organism, that is merely another change in the independent variable, followed in a lineal causal chain by another action and another consequence.

We expect most behaviorists to say that our S-R model is "reflexological"—a version of behavioristic theory many behaviorists disavowed years ago—and to echo the comment: "There may not be a reflexologist alive" (Shimp, 1989, p. 163). Protests aside, at the

core of every behavioristic theory is a claim that the environment controls behavior. From the beginning, behaviorists have asserted, like Donahoe and Palmer, "Although the organism is the locus of environmental action, it is the environment, and not the organism, that is the initiator and shaper of behavior" (1989, p. 410). When Hayes and Brownstein (1986) discussed prediction and control as criteria for evaluating behavioristic analyses of behavior, they said, "One could ask, for example, how do we know that *this is* the relevant stimulus for *this* behavior? The answer is of the general form that when we change *this* stimulus (and not *that* stimulus), we get a change in *this* behavior (and not *that* behavior)" (p. 178, emphases in the original). And Skinner claimed, "The ways in which behavior is brought under control of stimuli can be analyzed without too much trouble..." (1989, p.14).

Here, we merely test results that would ensue were it in fact true that independent environmental stimuli specify instantaneous details of behavior and its consequences.

The "cognitive" model. Our cognitive model stands for all theories that say actions originate not from current external events, but from internal causes, inner traits, tendencies, propensities, sets, plans, attitudes, aspirations, symbol-generating processes, programs, computations, coordinative structures, or some kind of systematic endogenous brain activity. No major theory of this sort proposes that behavior is entirely spontaneous; in one way or another they say the internal causes of present behavior formed and changed slowly, during past experience with the outside world—the recent past in some theories, the geologically distant past in genetic theories of behavior. In cognitive theories, the link between present behavior and influences in the present external world ranges from weak to almost nonexistent. In many texts on cognitive theory, there is no mention of overt action, much less an attempt to explain such actions. When there are explanations, the causal chain runs from input to cognition to command to action to consequence.

Kihlstrom (1987) succinctly identified the linear causal model in cognitive theory: "Cognitive psychology comes in various forms, but all share an abiding interest in describing the mental structures and processes that link environmental stimuli to organismic responses..." (p. 1445). Each step of the assumed chain from stimulus (input) to response (output) is described in detail by various cognitive theorists. For example, Real (1991) describes how inputs from a variable world would be transformed, in three sequential stages, into cognitive "representations":

> ... three stages may be viewed... as three components of a single dynamical system mechanistically tied to the organism's nervous

> system. The encoding of information would... correspond to initial inputs, computational rules correspond to transient dynamics, and representations would correspond to the equilibrium configurations resulting from the transient dynamics. The animal reaches a representation of the environment through the operation of specific computational rules applied to a particular pattern of incoming sensory information (p. 980).

In a discussion of computations which they assume cause movement, Bizzi, Mussa-Ivaldi, and Giszter (1991) complete the chain between representations and actions: *"...* the central nervous system must transform the neural representation of the direction, amplitude, and velocity of the limb, represented by the activity of cortical and subcortical neurons, into signals that activate the muscles that move the limb" (p. 287).

Some theories combine cognitive and S-R models. In their simplest forms, hybrid models say that the mind-brain receives "inputs," then produces direct transformations of coordinates from "perceptual spare" to "action space" that are required to initiate commands to move the body or part of the body to a point specified in the input (as examples, see P. M. Churchman, 1986; P. S. Churchman, 1986). Such models reduce cognition and neurology to a simple table-look-up.

A more complex hybrid S-R/cognitive model was endorsed by the cognitive theorist Allen Newell (1990) in the 1987 William James Lectures. Newell spoke of how "It is possible to step back and treat the mind as one big monster response function from the total environment over the total past of the organism to future actions..." (p. 44). On a more immediate scale, he said, "The world is divided up into microepics which are sufficiently distinct and independent so that the control system (that is, the mind) produces different response functions, one after the other" (p. 44). For strategic purposes, Newell places his theory in the category of cognitive theories that he says do not effectively explain how perception and motor behavior are linked to central cognitive processes. Then he says that such theories *"... will* never cover the complete arc from stimulus to response, which is to say, never to tell the full story about any particular behavior" (p. 160). In his allusion to the reflex arc, Newell remarkably implies the equivalence of the causal models in his cognitive theory and in reflexological theory.

In either their simple or complex forms, hybrid S-R/cognitive models produce results identical to those of S-R models, so we will not discuss them further.

The perceptual control theory model. The PCT model, which we discuss later at some length, is the least familiar of the three models. In brief, it proposes that there is a simultaneous two-way interaction between organism and environment (see Hershberger, 1989; Marken, 1990; and

Powers, 1973, 1989, 1992). In PCT, the basic unit of behavior is not the linear input-output chain, but the negative-feedback loop, which has properties different from the units of the other two models and implies interesting consequences about the way an organism's actions alter the outside world.

"Models"

We use the term "model" in the very narrow sense in which an engineer would use it: a precise quantitative proposal about the way some system operates in relation to its environment. Most behavioral scientists use *descriptive* models, which merely rephrase (usually in words; sometimes in mathematical form) previously observed relationships between organism and environment. There are unlimited ways to restate behavioral data. If each of them passes as a *model* of behavior, then the list of seemingly plausible models is also limitless. The availability of many equally plausible descriptive models is behind the mistaken assumption, common in behavioral science, that models are poor substitutes for real understanding—that if one understood the phenomenon at hand, one would state the facts, not a "mere" theory or model.

But "model" also means, in the present context, a *generative* model, in which the proposed organization is stated in a way that can be used to calculate behavior as a function of moment-by-moment variations in the independent variable. By that usage, a model does not substitute for knowledge. To the contrary, simulation of a well-posed model rigorously tests one's presumed knowledge of the causal principles at work in behavior.

S-R theory as a model. Calculations of the correlation between a dependent and independent variable produce a correlation coefficient, a regression coefficient, and an intercept. In most behavioral research, little attention is paid to the regression coefficient and intercept, one reason being that the typical scatter of the data is large enough to make a linear regression line almost useless for predicting behavior. But, by the logic of the S-R approach, the regression equation constitutes both a generative model and a description. It is a first approximation to a proposed law of behavior: at every moment, the behavioral measure is proportional to the magnitude of the independent variable. If that law is true, one can vary the independent variable and calculate (predict) the dependent one strictly from the previously determined regression equation.

It can be argued that this strict interpretation of a regression equation is inconsistent with the state of the art in behavioral science—all we can hope for now, in most cases, is to establish the presence or absence of a statistically significant relationship. Our reply gives the benefit of the

doubt to the theory underlying the S-R concept. If, given as many years as necessary, methodologies improve, sources of variance are eliminated, and better data are obtained, then regression equations will become meaningful. When they do, there will be an obvious test for whether a proposed regression coefficient is a law of behavior. In the regression equation, one can impose a new pattern of the independent variable and calculate the resulting pattern of behavior, the dependent variable. The modeled result can be compared against what happens when the organism encounters the altered independent variable. In more elaborate form, this process of testing a model against actual events is the basic methodology of the physical sciences. Used in this way, the regression equation is a generative model.

We use an alternative to waiting for years for data to improve: we apply this method in an experiment so simple that the regression line is highly meaningful, and random variation is a minor factor. We subject the S-R model to a test under conditions that should make it work as well as it ever will.

Cognitive theory as a model. We give the cognitive model a similar treatment. Cognitive models are more difficult to test and defend than S-R models; there is no simple way to determine whether a given cognitive model is correct, as well as plausible. No matter how well a model proposing a specific organization of the mind-brain predicts behavior, one cannot test the model objectively by, for example, deriving a regression line based entirely on observable variables. There is no way to know whether some other cognitive model would not work as well or better. There is only one regression line that best fits the behavioral data, but there are many seemingly plausible cognitive models.

Kugler and Turvey (1987) aptly described the problem of non-unique computational models for behavioral output:

> Whereas physical events are said to follow uniquely from their causes, internally consistent, logical descriptions of the causal process are multiple .... How does one get from the existence of multiple (logical) descriptions to a unique (causal) description? Dressing up logical formulae in instantiable programs does not resolve the uniqueness problem. Many programs can give rise to the same sequence of machine outputs (p. 28).

To avoid problems of this sort, we give cognitive models the same benefit of the doubt that we give S-R models. Given proper knowledge of the history and properties of the environment, and the correct internal computations, the ideal cognitive model should calculate exactly the motor outputs required to produce a preselected result. Of course, even a perfect cognitive model would require experience with an en-

vironment to build up knowledge of its properties: if the environment changed, the model would need new interactions with the altered form before it could again compute the correct action.

We test the cognitive model by assuming that it is perfect: it makes optimal use of information and computes the same required action on successive trials, and the motor systems perfectly obey its commands.

The reasoning behind our approach to the models is simple: in a well-defined experiment, if quantitative predictions by both the S-R and cognitive models, given the benefit of every doubt, are incorrect, and the PCT model predicts correctly in the same experiment, there will be excellent reason to say that the control-theoretic model is right and the other two are wrong, for that experiment. How far one generalizes the result depends on how dear are the parallels with other experiments and the simple one we use: we leave such judgments to the reader.

Perceptual Control Theory as a Model

Perceptual control theory always considers two simultaneous relationships: (a) the observed dependence of stimulus inputs on behavioral outputs and independent events, and (b) a conjectured dependence of behavioral outputs on stimulus inputs.

The environment equation. The first relationship the PCT model describes is how the input to an organism depends on the organism's actions and on disturbances arising simultaneously with behavior but independently of it in the external world. To simplify this part of the model, we restrict all variables in the experiment to change in a single dimension, described later. Consequently, the variable at the organism's input is simply the sum of a physical effect from the organism's output and another physical effect from an independent disturbance. The apparatus (a computer system) records exactly what these relationships are and exactly what disturbance is acting at any moment. This part of the model is completely determined by the experimental setup; it is a statement of fact, not a conjecture, and it is illustrated in detail by Bourbon, Copeland, Dyer, Harman, and Mosely (1990).

The organism equation. Perceptual control theorists assume an organism can be modeled as a system that senses some aspect of the environment that is then represented internally as a one-dimensional perceptual variable. The magnitude of this variable is compared continuously against a reference signal (or reference magnitude) inside the organism or the model of the organism. Any difference between the reference signal and the perception is a non-zero "error signal" which drives action, again in a single dimension of variation.

This part of the model can be treated exactly as a regression equation. The slope of the regression line represents the incremental ratio of output to input, and the intercept represents the setting of the internal reference signal. The slope reflects measured output as a function of measured input; the intercept is the magnitude of input for which the output does not change. Control theorists assume that the value of the input for which the organism produces no change in output is the input that the organism specified in advance.

The system equations. The organism and environment equations form a system of equations; for examples, see Pavloski, Barron, and Hogue (1990, pp. 33-37); Powers (1973, pp. 273-282; 1978, pp. 422-428); and Runkel (1990, pp. 93-99). There are two system variables (the input and output variables) and two equations. The input and output variables appear in both equations, and each must have only one value at a time. Consequently, the system can be solved for each variable as a joint function of any system constants and the values of the two independent variables (the external disturbance and the internal reference signal).

Our experiments use random disturbances that cannot be represented by any reasonable analytic equation. Consequently, in the PCT model, we calculate numerical solutions of the system equations. Numerical solution of system equations, with time as a parameter, is called simulation.

Simulation. Simulation recreates, through computation, a continuous relationship among system variables and independent variables. The experimenter causes a pattern of changes in the independent variables, while the equations for the model continuously compute the states of dependent behavioral variables at the input and output. For a good model, the results of a simulation look very much like a recording of an organism's actions in an experiment where the independent variables change in exactly the same way as during the simulation; for a bad model, the results of the simulation do not resemble those produced by the organism.

Simulation involves at least two stages. The first matches simulated behavior to real behavior, after the fact, by adjusting the parameters in the model. The second stage uses a new pattern of variation in the independent variable, with the model's parameters set as previously determined, and records the behavior of the model. Then the new pattern of variation is applied to the person, whose behavior is recorded and compared with the model's behavior. In the sciences and in engineering, models are often tested in a third stage (as we do here), with both a new pattern of variation for the independent variable and a new kind of environmental disturbance, not used in the original parameter determinations. In this third stage, the model predicts, in simulation, relationships not previously observed.

Reduced to its essentials, the logic of simulation resembles more familiar ways of studying relationships and testing to see if they generalize. It is, however, much more exacting: it compares modeled and actual behaviors instant-by-instant, rather than in terms of static data sets. For the present experiments, the models predict thousands of values for several variables, all of which are compared with the values produced by a participant. The success or failure of a prediction is immediately obvious.

Some people argue that models which work properly in very simple situations might not work when complexities occur. The converse of that hypothesis, also sometimes offered, is that failure of a behavioral theory in a very simple experiment doesn't necessarily mean that it will fail in more realistically complex studies. But engineers, who deal with both simple and complex systems, would not agree. Certainly, a model that works in a simple situation might need considerable revision to work in a more complex situation. But if a model fails to work in the simplest possible circumstances, there is no chance that it will successfully predict more complex phenomena. Complexity can be an excuse for failures of a model in a complex situation, but not in a simple one. If the core assumptions of a model fail in simple experiments like ours, there is no chance the model will work in more complex circumstances.

The Experiment

The Task

Participants in this three-phase experiment move a control handle in one dimension, forward and backward. On a computer screen in front of them is a short horizontal bar, the "cursor," distinct from the background, which moves up as the handle moves forward and down when it moves back. Flanking the path of the cursor are two more bars, the "target," that remain even with one another and move slowly up and down the screen, following a path generated by the computer. The person's task in all phases of the experiment is to keep the cursor exactly between the target lines. (There is nothing special about that relationship between cursor and target; the person could easily select any other.) This task is known as "tracking." When the target is stationary, it is called compensatory tracking; when the target moves, as it does here, it is called pursuit tracking.

We can easily modify the experiment to include perceptual variables other than spatial position. For example, the handle can be set to alter the size or shape of a geometric figure, change the magnitude of a number displayed on the screen, or alter the pitch of a sound. And tracking can occur across stimulus attributes and sensory modalities, as when a

person uses the handle to make the pitch of a sound match the magnitude of a number or the vertical position of a target. All relationships observed during a simple tracking experiment are found in these other tasks; any of them can be used to make the points we make here.

The Conditions: Three Phases

Phase 1. In Phase 1, the target moves up at constant speed to a preset limit, then down at a constant speed to another preset limit, and so on, in a triangular wave. Each excursion up or down takes 2.8 seconds. The person practices as long as necessary to keep the cursor between the targets with an error of no more than three per cent of the total movement averaged over one minute. Data from the final minute of practice when this criterion is reached are saved as the data for the experimental run.

The relevant parameters are estimated for each model, and then the models reproduce the person's behavior. In the next two phases, we use the parameters thus determined to create a simulated run before the person runs a single one-minute trial. No model is altered, in any way whatsoever, from this point on.

Phase 2. Conditions in Phase 2 are the same as in Phase 1, except that there is a probability of 2/3 that the target speed will differ from the last speed on any given up or down excursion. The speed of each excursion is selected randomly from 1.4, 2.8, or 5.6 seconds per excursion, with a mean of 2.8 seconds over the one-minute experimental run (the same mean excursion time as in Phase 1). The person must still move the handle to keep the cursor between the target marks. A few minutes prior to the person's run, each model is run with the same randomly generated pattern of variations in target speed that the person will experience. The person gets no practice: the first run under these new conditions is the only run for Phase 2.

Phase 3. Conditions are the same as in Phase 2, except that now a smoothed random disturbance also acts on the cursor. The disturbance is created at the start of the entire experiment by smoothing the output of a random-number computer algorithm and storing the resulting waveform. The same disturbance is used in runs by the models and the person. Cursor position is determined by the sum of handle displacement from renter and the momentary magnitude of the disturbance. Again, the person does a single one-minute run with no practice. A few minutes before the person's run, each model predicts the results, with a new pattern of target excursions and with the disturbance acting on the cursor.

The experimental variables. During each 60-second experiment, each variable is sampled every 1 /30 second, for a total of 1800 values per

variable. In the figures illustrating the results, every third value is plotted. There are three measured variables: the positions of the target (T), handle (H), and cursor (C).

Phase 1

The person's data. The person kept the cursor even with the target, as shown in Fig. 1A. The perfectly regular triangular wave in the upper part of the figure is the vertical target position across time. The slightly less-regular wave that closely follows it is the cursor position created by the person. In the lower part is the handle-position record, identical to the cursor-position record because handle position directly determined cursor position. (The handle-position plot is scaled to be the same amplitude as the cursor-position plot; we use this scaling in all figures).

The mean vertical distance between the cursor and target was -0.8 units of screen resolution (S.D. -1.8; total vertical range on the screen = 200 units). The following Pearson correlation coefficients describe the

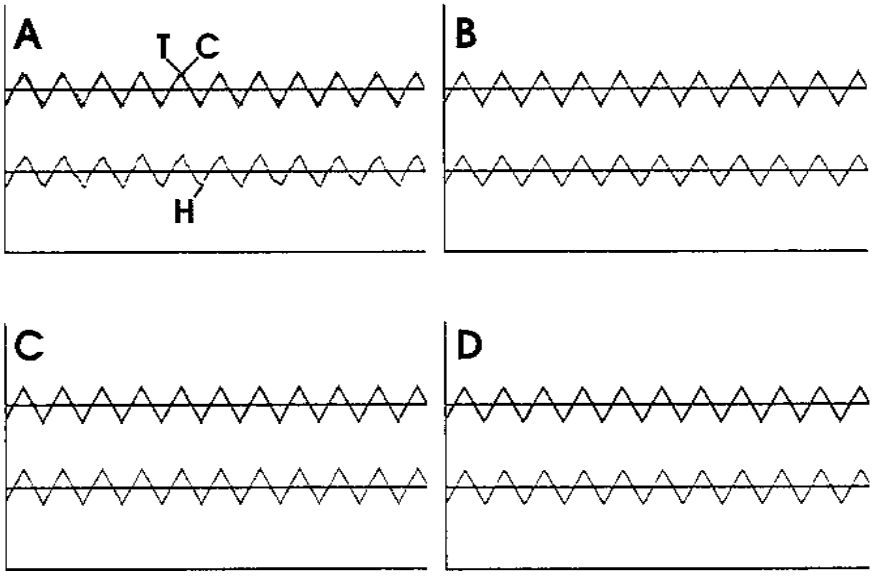

Figure 1. Results of pursuit tracking, Phase 1: data from the person (A); reconstructions of the data by the stimulus-response model (B); by the cognitive model (C); and by the control-system model (D. In A, H = handle, T = target, and C = cursor. For target and cursor, "up" in the figure is toward the top of the computer monitor; for handle, "up" is away from the person. The duration of each experiment is 60 seconds.

relationships among variables in Fig. 1A: between positions of the cursor and target, .977; handle and target, .977; and handle and cursor, 1.0. In the regression of handle on target, the slope was 0.89 (the person moved the handle the equivalent of 0.89 screen units for every movement of one unit by the target), and the intercept was -0.8, identical to the average difference between positions of the cursor and target.

Testing the models: The rationale. In simulations of the models, computations begin with all variables set to the same initial values from the first moment of the run by the person and are repeated 1799 times, once for every 1/30 second in the run by the person. Each model produces handle positions in its unique way, but a common procedure determines cursor positions.

Establishing the S-R model. We remind readers that we do not compare the relative merits of the many varieties of behavioristic theory, nor do we examine or challenge behaviorists' descriptions of conditions in which learning occurs. We merely examine consequences that would ensue were behavior controlled by an independent antecedent variable —were behavior literally "under environmental stimulus control."

Our simple S-R model is rigorously true to the requirements laid down for laws of behavior by B. F. Skinner (1953):

> The external variables of which behavior is a function provide for what may be called a causal or functional analysis. We undertake to predict and control the behavior of the individual organism. This is our "dependent variable"—the effect for which we are to find the cause. Our "independent variables"—the causes of behavior-are the external conditions of which behavior is a function. Relations between the two—the "cause-and-effect relationships" in behavior—are the laws of a science (p. 35).

In our simple experiment, the only independent variable is the position of the target, determined solely by the computer program. The position of the handle depends on the actions of the person, so it is a pure dependent variable, which we model as a response to target position. In Phase 1, the handle determines the position of the cursor, which is a remote (from the person) consequence of behavior, not a cause.

Cursor movement is also a "stimulus," by any traditional definition, but it is not independent of behavior; it lies at the *conclusion* of the assumed causal chain. At best, it might be a "reinforcing" stimulus. Behavioral theorists claim that reinforcement produces long-term changes in the probability of a general class of actions (an "operant"). For example, some might say that, at an earlier time, cursor movement reinforced handle movement, which explains why the person uses the handle now. But reinforcement theory does not explain or predict how

a person produces moment-by-moment changes in behavior and in its consequences.

We use a regression equation as our S-R model. For the handle and target positions in the person's data, shown in Fig. 1A, the slope (m) of the regression of handle on target is 0.89, and the offset (intercept, b) is -0.8. We represent target position as t, handle position as h, and cursor position as c. Therefore, the S-R model for handle position is of the form

$$h = mt + b,$$

and the position of the cursor is modeled as

$$c = h.$$

Results of running the S-R model. To "run" the S-R model, we start with all variables at their values during the first instant of the run by the person, then we multiply the remaining 1799 target-position values, in sequence, by the slope m and add the intercept b, and obtain the successive predicted positions of the handle and cursor, shown in Fig. 1B.

The positions of handle and cursor created by the model resemble those from the person: the correlation between modeled and actual handle positions is .977; between modeled and actual cursor positions, also .977. Our simple reflexological model accounts for % per cent of the variance (r-squared) in the behavioral data from Fig. 1A; the regression equation is highly meaningful.

Establishing the cognitive model. Our goal with the cognitive model is not to compare the many diverse computational algorithms studied by cognitive and brain scientists. We merely examine the consequences that would ensue, were it possible for a system to reliably compute the same output, no matter how it does the computation. Our cognitive model assumes that, during the practice period, some central process learns and models the amplitude and frequency of target movements and computes commands that cause the muscles to move the handle, and thus the cursor, in a pattern as close as possible to that of the target.

A detailed version of this model would use a program loop simulating a "higher cognitive process" to compute handle positions independently of target movements. It would generate commands for the amplitude, frequency, and shape of the movements. But severe phase errors (mismatches in timing between the positions of the target and the model's handle) would develop unless we gave the model exact information about the frequency of the target and started it at exactly the right moment with exactly the right initial conditions. To assure that there were no errors, we would tell the model exactly how

to move the handle to re-create the results of Phase 1. To achieve the same result, without the complex computations, we simply assume that, however the cognitive model works, it works perfectly: it computes handle movements to match the average pattern of previous target movements. For the last minute of practice, it uses information accumulated earlier to command movements that reproduce the movements of the target (of course the model we use here does not actually need any practice).

This makes the cognitive model exceedingly simple: it is of the form

h = t.

Handle movements perfectly reproduce movements of the target that occurred during the practice run, and the resulting cursor movements also perfectly reproduce the movements of the target.

Results of running the cognitive model. A run of the cognitive model is extremely simple: since h = t and c = h, we simply plot the successive target position values as c and as h. The upper trace in Fig. 1C shows target and cursor positions perfectly superimposed; the lower trace of handle position is identical to the upper traces. The positions *of* handle and cursor created by the model are like those from the person: the correlation between modeled and actual handle positions is .977; between modeled and actual cursor positions, also .977.

Establishing the control-theory model. The environment part of the PCT model is just a description of the external situation: cursor position depends on handle position plus the magnitude of any possible disturbance. The environment equation is

c = h + d.

In Phase 1, the disturbance magnitude is zero.

The fact that the cursor is also a dependent variable wholly or partly determined by handle position is not a problem, because both the organism equation and the environment equation form a single system of equations. We symbolize the perceived separation of cursor and target, c - t, as p, which we take as the real input variable. This variable p is compared against a reference level p*, which specifies the state of p at which there will be no change in output; it is the value of p that the person intends to experience. Any difference between p and p* is called "error." The output, which is the handle position h, is the time-integral of error and takes the form

h = k[int(p* - p)].

The constant k is the "integration factor." It represents how rapidly the person moved the handle for a given difference between the perceived separation p and the reference separation p*; k is expressed in units of screen resolution the cursor would move per second for a given amount of perceived error.

To fit the model to the subject's behavior, we estimate p* and k, the only adjustable parameters of the model. We set p* equal to the average value of cursor-minus-target during the person's run in Phase 1. (By estimating p* from the data, we avoid claiming that we know the person is trying to keep the separation of target and cursor at zero. The person can maintain any reasonable separation-there is nothing special about p* = 0.) To estimate k, we insert the estimated value of p* into the model, then we insert an arbitrary value of k and "run" the model, a procedure we explain below. During each of several successive runs of the model, we insert a new arbitrary value of k and calculate the root-mean-square (RMS) difference between all of the cursor positions from both the model and the person. The best estimate of k is the one from the run with the smallest RMS difference.

To "run" the model, we start the handle position at the subject's initial handle position during Phase 1, and then do the following computer program steps over and over, changing the value of t on each step to re-create the target movements:

l: c: = h + d
2: p: = c-t
3: error: = p* - p
4: h: = h + k • error • dt

where dt is the physical duration represented by one iteration of the program steps. In all of the experiments reported here, each iteration represents 1 / 30 second, so dt = 1 / 30 sec. For the various terms in the program steps, k and p* are the system constants: k is the tentative value of the integration factor and p* is the estimated reference signal; t is the momentary target position, c is the cursor position, h is the handle position, and d is the disturbance magnitude-here, 0.

The fourth program step is a crude form of numerical integration; the notation means that the new value of h is computed by adding an amount (k • error • dt) to the old value of h. These are program steps, not algebra: do not cancel the h's! The "colon-equal" sign is the replacement operation, which replaces the previous value of the variable on the left with the new computed value of the argument on the right.

Results of running *the PCT model.* In the person's run during Phase 1, p* was estimated as -1 unit on the screen (-0.8 rounded), which means that, on average, the person kept the cursor slightly below the target.

Following the procedure described above, the estimated best value of the integration constant k was 8.64 in units of resolution per second.

The results of a run of the model with those estimated values of p'" and k are shown in Fig. 1D. The positions of handle and cursor created by the model resemble those from the person: the correlation between modeled and actual handle positions is .989; modeled and actual cursor positions, also .989.

Summary of Phase 1. The person performed the tracking task reasonably well, and simulations of all three models produced results like those from the person. After this round of simulations, all three models remain defensible as explanations of the person's performance.

Phase 2

Next, we use the three models to predict behavior when one condition changes, then the person does a run under exactly the same conditions as those encountered by the models. The changed condition is that the target now moves up and down at randomly varying speeds. The mean speed is still 2.8 seconds per excursion, but on every successive excursion, there is a 2/3 probability of a change of speed that lasts until the end of the excursion, and then the next speed is selected randomly. The random changes are generated beforehand and recorded, so the same changes are presented to all three models and to the person. We have already established the three models, so our descriptions of the results are brief.

The person's data. Fig. 2A shows data from the person's run, after the models made their predictions. The person made the cursor follow the target about as well as in Phase 1. The mean vertical distance between cursor and target was -1.4 units of vertical screen resolution (S.D. = 2.2). The following Pearson correlation coefficients describe relationships among variables in Fig. 2A: between positions of the cursor and target, .966; handle and target, .966; and handle and cursor, 1.0.

Prediction of the S-R model. The linear regression equation developed after Phase 1 accurately predicts the positions of the cursor and handle despite the changes in target speed, as is shown in Fig. 2B. This is possible because, just as in Phase 1, the required handle movement is simply proportional to target movement at every instant. The positions of handle and cursor created by the model are like those from the person: the correlation between modeled and actual handle positions is .989; between modeled and actual cursor positions, also .989.

Prediction of the cognitive model. The results for the cognitive model, shown in Fig. 2C, reveal the first obvious failure of a model. The positions of handle and cursor created by the model are not like those from the person: the correlation between modeled and actual handle

positions is .230; between modeled and actual cursor positions, also .230.

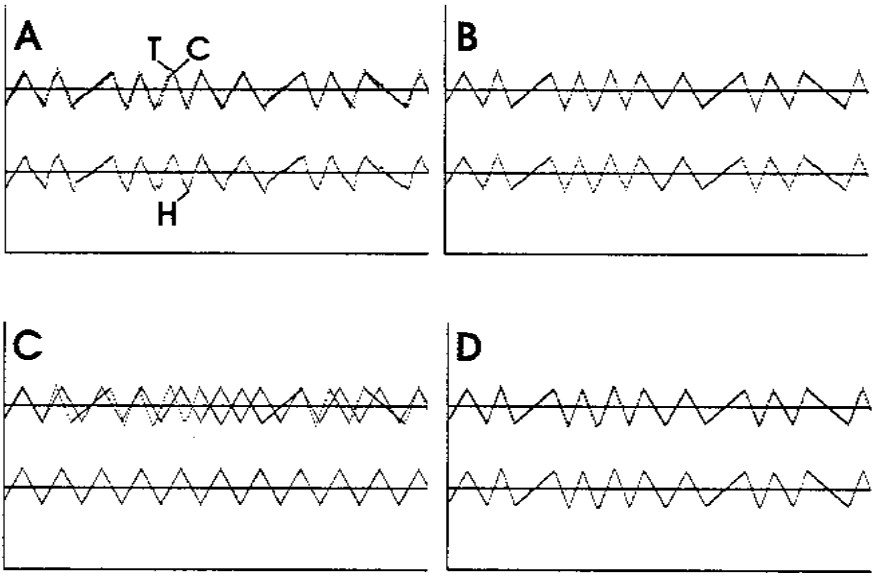

Figure 2. Results of pursuit tracking in Phase 2: data from the person (A); predictions of the data by the stimulus-response model (B); by the cognitive model (C); and by the control-system model (D. In A, H = handle, T = target, and C = cursor. For target and cursor, "up" in the figure is toward the top of the computer monitor; for handle, "up" is away from the person. The duration of each experiment is 60 seconds.

The reason for this failure is obvious. The cognitive model assesses properties of the environment and computes an action that will have a required result. But now the environment, in the form of target movements, is subject to unpredictable variation. The cognitive model gets no information about the next target speed before it is experienced. Thus, the best that a cognitive "central-process" model can do is command its output to match the best estimate of average target speed; in the present case, that average is the speed that occurred throughout Phase 1, when the motor plan was established. The cognitive model continued to produce a triangular wave of handle and cursor movement that conformed to the average waveform of target movement—a form not like the waveform of the target in Phase 2.

One might think of modifying the cognitive model so that the central processor re-assesses the environment's properties on an instant-by-instant

basis. That would solve the problem, but only at the expense of converting the cognitive model into a control-system model intent on making its output match its input: the new model would be a control-system model acting like a stimulus-response model. The core concept of a cognitive motor plan would be abandoned.

Prediction of the control-system model. Fig. 2D shows the results for the control-system model. The program steps from Phase 1, using the same values for the parameters k and p*, successfully predict the person's handle and cursor positions. The correlation between modeled and actual handle positions is .981; between modeled and actual cursor positions, also .981.

Summary of Phase 2. The person performed the tracking task with reasonable accuracy, and simulations of the S-R and PCT models produced results like those for the person. However, the cognitive model continued to make its output follow the path 'learned" during Phase 1; consequently, its cursor did not follow the now-erratic waveform of the target. After this round of simulations, only the S-R and PCT models remain reasonable as explanations of the person's performance.

Phase 3

Now the three models predict behavior under a radical change of conditions. The target still moves up and down at randomly varying speeds, as in Phase 2, but for every time-interval, a new value of a random disturbance is added to the position of the cursor. Now, with the handle held still, the cursor wanders randomly up and down. When the handle moves, the net movements of the cursor are determined by the sum of handle movements and disturbance changes.

In both previous phases, the "d" in the cursor equation, c = h + d, was zero. Now it varies unpredictably, although not rapidly (the bandwidth of variations is about 0.2 Hz). This new disturbance enters after the motor outputs of the person and the accompanying handle movements, "downstream" in the causal chain. The cause of the disturbance is hidden; the only evidence the person has about the disturbance is the deviation of cursor position from the momentary equivalent of the handle position. At any moment, there is no practical way for the person to know the degree to which either the position of the handle or the value of the disturbance affects the position of the cursor.

The person's data. As we show in Fig. 3A, the person still made the cursor track the target (mean distance between cursor and target = -1.0 screen units, S.D. = 3.0), despite the unpredictable variations in target speed and the unpredictable interference of a disturbance. Had the person not moved the handle, the correlation between positions of the cursor and momentary values of the disturbance would have been + 1.0;

that between positions of cursor and target, near 0.0. Instead, the correlation between the disturbance and cursor was only .101, while that between cursor and target was .940.

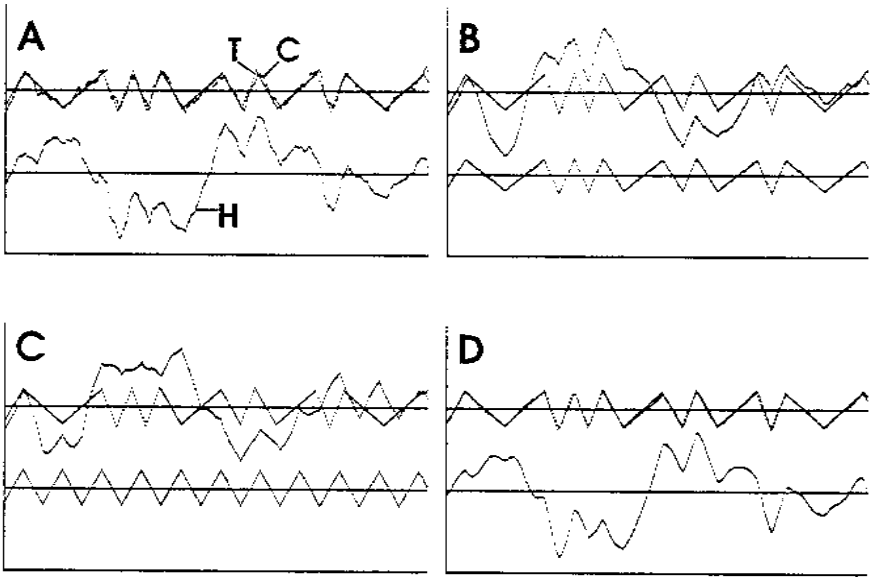

Figure 3. Results of pursuit tracking in Phase 3: data from the person (A); predictions of the data by the stimulus-response model (B); by the cognitive model (C); and by the control-system model (D). In A, H = handle, T = target, and C = cursor. For target and cursor, "up" in the figure is toward the top of the computer monitor; for the handle, "up" is away from the person. The duration of each experiment is 60 seconds.

In Phases 1 and 2, the handle alone determined the position of the cursor: the correlation between handle and cursor was + 1.0. But in Phase 3, the person moved the handle any way necessary to cancel the effects of the random disturbance on the cursor: the correlation between positions of handle and cursor is only .294, that between positions of the handle and the disturbance that moved the cursor away from the target is -.992.

Prediction of the S-R model. As we show in Fig. 3B, the S-R model failed: the correlation between modeled and actual handle positions is .296; between modeled and actual cursor positions, .385.

Successful simulation can no longer be attained by moving the handle in synchrony with target movements. That is why the person moved the handle in a pattern that deviated radically from the pattern of target

movements; the deviations were exactly the ones needed to counteract the effects of the new disturbance. But the S-R model responded to the target stimulus just as before, and moved the handle proportionately to any movement of the target. The simulated cursor, now subject to an independent disturbance, did not follow the target.

To salvage the S-R model, one might propose that the cursor, too, be included in the definition of the stimulus. However, the person's data in Fig. 3A show that the cursor moved in nearly the same pattern as the target, but neither pattern resembled what the handle did. To include the cursor in the definition of the stimulus, we might conclude that the difference between the target and cursor positions is the stimulus. On further examination, we would find that this difference does not match the handle movements, either, but its time-integral does: perhaps the time-integral is the stimulus. That change is acceptable, but if we adopt it, we are left with the fact that cursor position depends, simultaneously, on handle position and the independent random disturbance: now there is no true independent variable in the causal chain, and the core premise of any model of stimulus control over behavior is abandoned. Neither the cursor nor any relationship between the cursor and any other variable can be described as a pure independent variable, because it is also, at every moment, a dependent variable.

Prediction of the cognitive model. Fig. 3C shows that the prediction by the cognitive model failed. The model followed its plan learned in Phase 1 and moved the handle to conform to the average behavior of the target. It should have moved the handle in the erratic pattern produced by the person, shown in Fig. 3A. The correlation between predicted and actual handle positions is .119; between predicted and actual cursor positions, .151.

Even if we gave the cognitive model more practice in the new situation (and the ability to learn), it would revert to essentially the same actions. The average deviation of cursor speed from 2.8 seconds per excursion is zero. The average amount of disturbance applied to the cursor closely approximates zero. Neither the next speed of the target nor the next variation in the disturbance is predictable. No matter how smart one wants to make the central processor when it comes to predictions, we can always make the disturbances still more random. Any cognitive model must compute output that is calculated to have a desired effect. It can base its computations only on experience with properties of the external world. When those properties contain significant instant-by-instant irregularities, as they do in our simple experiment, the core concept of the cognitive model cannot work. Unless, of course, it is modified to compare its plan of the world against its momentary perceptions of the world and to act so as to eliminate any discrepancy, but those modifications would make the model a control-system model.

Prediction of the control-system model. As we show in Fig. 3D, the control-system model produced precisely the outputs required to maintain a pre-selected target-cursor separation, despite two kinds of random variation that called for pronounced changes in the output pattern. The PCT model faithfully predicted the person's behavior. The correlation between actual and predicted handle positions is .996; between actual and predicted cursor positions, .969. Correlations as high as those here, between tracking behavior and predictions by PCT, are commonplace, even when the interval between predictions and behavior is as long as one year as is reported by Bourbon, Copeland, Dyer, Harman, and Mosley (1990).

To avoid drawing this paper out any longer, we omit analyses of other variations that the person and the PCT model can handle, with no change in the model's parameters. Both the person and the control-theory model continue to track accurately if we alter the scaling factor that converts handle movement into cursor movement; if we add a third or a fourth or a fifth independent source of disturbance to target speed or cursor position; if we put nonlinearity into the connection between handle and cursor (the person and the model still move the handle in an inverse nonlinear relationship to target and disturbance); or if we make the ratio of handle movement to cursor movement time-dependent (at a reasonable speed). None of these variations can be handled by the core concepts of the S-R or cognitive models. Yet all of these variations, as well as those shown in the three phases of our experiment, are commonplace in the real environments where real behavior must work.

Discussion

We attempted to determine if core assumptions about the immediate causes of behavior in three different models of behavior are consistent with what is known about the world in which behavior occurs. We compared specific predictions made during simulations of the three models with the performance of a person for three phases of a simple task. We concluded that the causal assumptions in a control-theoretic model are consistent with what is known about the world, while those in any pure stimulus-response (stimulus-control) model, or any pure cognitive-control (neurological-control) model, are not. The control theory model assumes that, when organisms act, they produce correspondences between their immediate perceptions of selected variables in the world and internal (to the organisms) reference states (reference signals) for those perceptions.

We did not ask whether reference signals exist in any particular physical form, or, if they do, whether they are "gained" through interaction

with the world, whether animate, inanimate, or social, or are inherited as part of a "genetic code." Robinson (1976) wrote of this issue in a discussion of Aristotle's concept of "final cause," which refers in part to a person's goals or intentions: "The issue is not *how* a given goal or intention was established. Rather, the issue or proposition is that outcomes are never completely understood until the final cause is apprehended, no matter what 'caused' the final cause" (p. 91, emphasis in the original). In our simulations, by hypothesizing and estimating the magnitudes of "reference signals," whatever their origins, that function in the manner of "final causes" within a control-system model of a person, we can understand and predict the outcomes when the person controls selected perceptions of parts of the unpredictably variable environment.

Modeling as a proper test of theory. The success or failure of our simulations immediately revealed the robustness, or lack of robustness, of alternative models of behavior. Other behavioral scientists recognize the importance of comparing the simulated behavior of models against the actual behavior of organisms. In a critique of conventional statistical methods in psychology, Meehl (1978) remarked:

> In my modern physics text, I am unable to find a single test of statistical significance. What happens instead is that the physicist has a sufficiently powerful invisible hand theory that enables him to generate an expected curve for his experimental results. He plots the observed points, looks at the agreement, and comments that "the results are in reasonably good agreement with the theory." Moral: *It is always more valuable to show approximate agreement of observations with a theoretically predicted numerical point value, rank order, or function form, than it is to compute a "precise probability" that something merely differs from something else" (p.* 825, emphasis in the original).

Similarly, Dar (1987) wrote:

> In physics... theories are tighter and lead to precise predictions. As a consequence, (a) if the numerical result is as predicted (that is, close enough to the predicted point value or curve), it will be very difficult, in contrast to the situation in psychology, to offer a reasonable alternative theory for that. This is because it is difficult to imagine alternative states of nature that will lead to the exact same curve or numerical result. (b) If the experimental result is not as predicted, some serious revision of the theory would be required. This is because a tight theory simply does not allow for significant (I do not mean "statistically significant") discrep-

ancies from predicted outcome (p.148).

And in his review of a book on cognitive theory, the behaviorist Shimp (1989) declared:

> A theory that behaves, that produces a stream of behavior, would seem in an intriguing way to fit better with Skinner's chief criterion for a good theory than do many more common sorts of behavioral theory. Skinner has argued that a good behavioral theory is a theory on the same level as the behavior itself. What is closer to the level of a behavior stream of an organism than a behavior stream of a theory? (p. 170).

We could not say it better. On any given experimental run, our simulations produced multiple simultaneous streams of behavior, altogether comprising thousands of predicted data points. The levels of agreement between the simulations and the behavior of a person allowed us to immediately assess the adequacy of the three models of behavior and of their implied models of the world.

The worlds implied by the models. For all three models, the results reported here would be general. Within its physical limits, any S-R system could make its movements match any target input, no matter how unpredictable. But, as happened with the cursor in Phase 3, if the consequences of those movements were disturbed, they would always deviate from the target by an amount equal to the variations in the disturbance.

Upon its first encounter with a new pattern of input, no cognitive system could compute commands to immediately make its behavior match the input. After some time, of course, an appropriately endowed cognitive system could search for a new pattern of commands. But if the input followed an unpredictable path or were presented only once or too few times for the system to "compute" an appropriate plan, learning would be impossible. Furthermore, if the consequences of its actions were continuously and randomly disturbed, no command-driven cognitive system could compute behavior to keep the consequences in any consistent relationship with the input. To do that, the behavior must deviate from its original pattern by precisely the amount needed to cancel the effect of the disturbance, but the source of the disturbance cannot be sensed in advance to allow anticipatory compensations in the commands for behavior.

The only ways to salvage the traditional models, short of turning them into control systems, rely on whimsical assumptions about the world. For example, the S-R model might still work if it were only necessary that changes in stimulation result in corresponding changes in behav-

ior, with no regard for the consequences of behavior; and the cognitive model might still work, if it were only necessary that movements repeat, while their consequences were allowed to change at random. But those assumptions contradict any reasonable understanding of behavior and its role in survival: behavior is functional, and its consequences matter. An alternative defense is to assume that the antecedents of behavior never change, or that they conveniently change across a small enough set of discrete options so that we can always recognize which one is present and perfectly match it with computed outputs-either that, or we must anticipate the changes by "precognition." And nothing must ever disturb the consequences of behavior. The world demanded by those assumptions is not the one we know.

In contrast, within broad limits, any perceptual control system would vary its behavior to keep its perceptions of a controlled variable at the value specified by a reference signal, even if both the target event and the consequences of the system's actions were subject to unpredictable variations.

We live in a changeable world, in which organisms with behavior determined solely by environmental stimuli or solely by internal commands could not survive; but theories of behavior that postulate control by stimuli or by commands have survived for centuries largely because they are not systematically exposed to the test of modeling. To modify cognitive or S-R models so that, like living systems, they might thrive amidst change, we must abandon the core concept that behavior is at the end of a causal chain, wherever the chain allegedly begins. We must give each model an internal standard and a process for comparing present perceptions against that standard. But then the models would all be control systems, each controlling its input.

Conclusions. The sciences of life reflect a three-century commitment to linear models of cause and effect, with behavior as the final step in a causal sequence. If we are to advance our understanding of life, we must question those venerable models, however plausible they seem. We can no longer embrace them, knowing that they presuppose nonexistent worlds. To question our traditional models raises the specter of difficult change; but if we retain them, with their fanciful worlds, we risk the trivializing and decline of our science.

The search for alternative models of behavior can begin with a simple change in the question we ask, from "Why is behavior the way it is?" to "Why is the world the way it is?" The answer to the new question includes a long-elusive answer to the old one: the behavior of organisms controls many variables in the world.

Acknowledgments

We thank Andrew C. Papanicolaou, Philip J. Runkel, Gregory Williams, and William D. Williams for their critical reviews of several early versions of the manuscript, and Glenn Millard for valuable assistance in preparing the figures.

References

Bizzi, E., Mussa-Ivaldi, F. A., & Giszter, S. (1991). Computations underlying the execution of movement: A biological perspective. *Science,* 253, 287-291.

Bourbon, W. T., Copeland, K. C., Dyer, V. R., Harman, W. K., & Mosley, B. L. (1990). On the accuracy and reliability of predictions by control-system theory. *Perceptual and Motor Skills,* 71,1331-1338.

Churchland, P. M. (1986). Some reductive strategies in cognitive neurobiology. *Mind,* 95, 279-309.

Churchland, P. S. (1986). *Neurophilosophy: Toward a unified science of the mind-brain.* Cambridge, MA: MIT Press.

Dar, R. (1987). Another look at Meehl, Lakatos, and the scientific research practices of psychologists. *American Psychologist,* 42,145-151.

Donahoe, J. W., & Palmer, D. C. (1989). The interpretation of complex human behavior: Some reactions to *Parallel distributed processing,* edited by J. L. McClelland, D. E. Rumelhart, and the PDP research group. [Book review]. *Journal of the Experimental Analysis of Behavior,* 51, 399-416.

Hayes, S. C., & Brownstein, A. J. (1986). Mentalism, behavior-behavior relations, and the behavior-analytic view of the purposes of science. *The Behavior Analyst,* 9, 175-190.

Hershberger, W. A. (Ed.). (1989). *Volitional action: Conation and control.* Amsterdam: North-Holland.

Kihlstrom, J. F. (1987). The cognitive unconscious. *Science,* 237, 1445-1452.

Kugler, P. N., & Turvey, M. T. (1987). *Information, natural law, and the self-assembly of rhythmic movement.* Hillsdale, NJ: Erlbaum.

Marken, R. S. (Ed.). (1990). Purposeful behavior: The control theory approach. [Special issue]. *American Behavioral Scientist, 34(1).*

Meehl, P. E. (1978). Theoretical risks and tabular asterisks: Sir Karl, Sir Ronald, and the slow progress of soft psychology. *Journal of Consulting and Clinical Psychology, 46,* 806-834.

Newell, A. (1990). *Unified theories of cognition.* Cambridge, MA: Harvard University Press.

Pavloski, R. P., Barron, G. T., & Hogue, M. A. (1990). Reorganization: Learning and attention in a hierarchy of control systems. *American Behavioral Scientist, 34,* 32-54.

Powers, W. T. (1973). *Behavior: The control of perception.* Chicago: Aldine.

Powers, W. T. (1978). Quantitative analysis of purposive behavior: Some spadework at the foundations of experimental psychology. *Psychological Review, 85,* 417-438.

Powers, W. T. (1989). *Living control systems: Selected papers of William T. Powers.* Gravel Switch, KY: Control Systems Group.

Powers, W. T. (1992). *Living control systems IL Selected papers of William T. Powers.* Gravel Switch, KY: Control Systems Group.

Real, L. A. (1991). Animal choice behavior and the evolution of cognitive architecture. *Science, 253,* 980-986.

Robinson, D. N. (1976). *An intellectual history of psychology.* New York: Macmillan.

Runkel, P. (1990). *Casting nets and testing specimens: Two grand methods of psychology.* New York: Praeger.

Skinner, B. F. (1953). *Science and human behavior.* New York: Free Press.

Skinner, B. F. (1989). The origins of cognitive thought. *American Psychologist,* 44,13-18.

Shimp, C. P. (1989). Contemporary behaviorism versus the old behavioral straw man in Gardner's *The mind's new science: A history of the cognitive revolution.* [Book review] *Journal of the Experimental Analysis of Behavior,* 51,163-171.

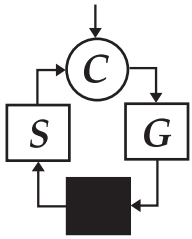

The Control Systems Group is a membership organization which supports the understanding of cybernetic control systems in organisms and their environments: *living control systems.* Academicians, clinicians, and other professionals in several disciplines, including biology, psychology, social work, economics, education, engineering, and philosophy, are members of the Group. Annual meetings have been held since 1985. CSG publications include a newsletter and a series of books, as well as this journal. The CSG Business Office is located at 73 Ridge Pl., CR 510, Durango, CO 81301; the phone number is (303)247-7986.

The CSG logo shows the generic structure of cybernetic control systems. A Comparator (C) computes the difference between a reference signal (represented by the arrow coming from above) and the output signal from Sensory (S) computation. The resulting difference signal is the input to the Gain generator (G). Disturbances (represented by the black box) alter the Gain generator output on the way to Sensory computation, where the negative-feedback loop is closed.

Inside back cover

Currently available:

LIVING CONTROL SYSTEMS

Selected Papers of William T. Powers

ISBN 0-9624154-0-5, 1989, 300 pages, $16.50

Introduction to Modern Psychology

The Control-Theory View

Edited by Richard J. Robertson and William T. Powers

ISBN 0-9624154-1-3, 1990, 238 pages, $25.00

LIVING CONTROL SYSTEMS II

Selected Papers of William T. Powers

ISBN 0-9624154-2-1, 1992, 277 pages, $22.00

MIND READINGS

Experimental Studies of Purpose

Richard S. Marken

ISBN 0-9624154-3-X, 1992, 212 pages, $18.00

Prices are postpaid. KY residents please add 6% sales tax. Order from: CSG Book Publishing, 460 Black Lick Rd., Gravel Switch, KY 40328.

Back cover