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

Posted at www.pctresources.com

Proofread as of March 6, 2001

# *Closed Loop # 2*
## *Threads from CSGNet*

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

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



*Greg Williams*
*606-332-7606*
*May 1991*

## *The Method of Levels and Internal Conflict*

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

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

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

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

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

Not every human problem, given this understanding, is a therapy problem. Therapy will not provide missing higher-level systems. It will not cure goals that are set at the highest levels in ways that guarantee conflict with everyone else. It will not provide the things that education provides: understanding of the world, of other people; acquisition of skills. It will not provide what spiritual searches provide: the sense of harmony and beauty that makes a person feel whole, that makes life

worth bothering with. Those things concern us all, and no one of us is more than a learner in those regards. The end of the line in therapy is not becoming a superbeing, but becoming an ordinary person capable of entering the struggle along with the rest of us—getting up to speed, as it were, for continuing a journey in a direction that is not clear to anyone.

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

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

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

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

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

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

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

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

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

it. The speaker is the ultimate authority.

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

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

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

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

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

only within and by the person in question.

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

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

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

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

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

Even finding that a drug treatment affects the symptoms does not prove a physiological origin, because generally the "psychoactive" drugs used affect functions of the brain that are also affected by normal brain activity. If, for example, there is a dearth of dopamine in some part of the brain, this is because the normal sources of dopamine, a neurotransmitter, are not active. You can either supply dopamine artificially, treating the symptom, or find out why the normal sources have dried up, treating the cause. The cause could be a physical mal-

function in the neurons themselves, or it could be lack of activation from other systems that normally send signals to those neurons. In the latter case, I see no reason why the explanation could not turn out to be conflict that is canceling the normal output of a control system somewhere else.

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

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

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

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

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

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

So here's the "but." My way of applying this method essentially ignores the content of the ongoing conversation, treating every statement, every description, as nothing more than possible evidence about a higher-level point of view that's in operation, in the background. One thing that feels very strange to a victim of this process is that the questioner never really seems to make contact; it's like talking to someone who seems overly fascinated with the way your mouth moves and isn't showing normal reactions to what you're saying. In learning to do this, I have learned how not to let remarks pass as if they were just a sort of innocent accompaniment to the main theme. For the speaker, the background attitudes and thoughts are silent and hardly noticed at all, the way you don't realize that you keep looking at your watch because you're in a conversation that has to end before your plane leaves. These background processes are there in consciousness, but only a very little bit, not enough to warrant full attention. The whole point of the Method

of Levels is that the *listener* can't do the same thing—realize vaguely that the background thingie is there, but let it pass without giving it full attention. The listener has to notice those glances at the watch.

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

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

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

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

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

This is what it's like to be so focused on an external problem that you're totally unaware of where you're coming from. Everything tells you that you're completely justified in needing to solve the problem, that something very bad is happening out there, and that you need to do something about it. That might be the exact truth, in that any normal person would feel the need to do something. In the background, however, there are all sorts of conflicts that keep you from thinking of an effective action, so all you know is that you feel helpless and overwhelmed by the

problem and need help with *it*, not with yourself. This sort of problem is a real attention-grabber. The only thing getting reorganized is what you imagine to be going on behind the scenes, what actions you imagine taking (and immediately give up on because every one of them arouses some sort of conflict).

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

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

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

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

I see symptoms as just that, symptoms. They're not problems, only evidence of problems. Unfortunately, the two are easily confused; that's because we see the symptoms. The conflict is all internal. All symptoms do is give evidence of conflict. The real problem is that somewhere within a client's system there is conflict, or a lack of harmony. When clients come to me, they are obviously reorganizing (who would pay a private counselor when their world is in harmony, when their goals—read reference signals—are being satisfied?). When

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

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

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

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

finding and maintaining a job).

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

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

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

In summary, Bill, don't revise the theory, it's working quite well, thank you (and I've spent 10 years learning it). Also, I have found that level raising does work. Finally, people shouldn't be listening to the therapist. That's because the therapist shouldn't be doing the

talking. The job of the therapist is to question the client and listen, listen and watch the way the client is dealing with his/her world. That's the way you teach people to think (a rare experience in school these days). I question people about their world (reference levels and perceptions to you scientists) and ask them if their worlds compare favorably (perceptual error or no perceptual error). If they don't, I ask them if they want to set a reference condition for working at another way of getting what they want, and then I teach them (because now I'm perceived as a teacher) to get what they want, making sure in the process they establish measurable goals that can be easily compared with internal reference signals.

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

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

## *"Conditioning"*

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

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

By the way, has anyone done an experiment something like the following? Take a "conditioned" Pavlovian dog and fill its mouth with a

working load of saliva before presentation of the conditioned stimulus. Does it then salivate at the bell? CT should say it doesn't.

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

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

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

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

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

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

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

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

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

But could you try this out for something like the startle reaction to a sudden loud sound? What good does jumping out of your chair do

when someone pops a balloon right behind you? In fact, the startle reaction also includes an eyeblink. Is this just a useless side effect of some behavior which is in some way more functional? Perhaps just "priming the pump" to get the systems going for flight or fight?

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

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

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

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

*Wayne Hershberger.* I am disappointed, Gary, that you found my control-theoretic account of respondent conditioning difficult to understand ("Control theory and learning theory," in the special issue of *ABS* edited

by Rick Marken: 1990, 34, pp. 55-66). The audience I had in mind while writing that paper was the psychologist who is familiar with learning theory and conditioning phenomena, but I had supposed that what I was saying would also be clear and convincing to readers familiar with control theory. I have also been cheered by the reprint requests I continue to receive for that paper, believing that my readers understood my message. Perhaps neither assumption is warranted—what a discouraging thought.

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

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

Classical Conditioning

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

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

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

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

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

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

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

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

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

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

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

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

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

Perceiving Environmental Disturbances

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Nice work.

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

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

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

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

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

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

When I've thought about modeling conditioning, I have done so in terms of your little stick man who reaches out as if to touch visible tar-

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

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

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

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