




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

Posted at www.pctresources.com

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

Spring 199

Volume 2

Number 2

Front cover

Closed Loop

Threads from CSGNet

Spring 1992 Volume 2 Number 2

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

CONTENTS

<i>Epistemology</i>	1
Avery Andrews, Gary Cziko, Wayne Hershberger, Joel Judd, Rick Marken, Bruce Nevin, Bill Powers, Martin Taylor	
<i>Levels of Perception</i>	60
Mark Olson, Bill Powers	
<i>On Modeling</i>	67
Bill Powers	

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

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

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

Each contribution to this issue of *Closed Loop* is Copyright © 1992 by its respective author, All Rights Reserved.

Inside front cover

Epistemology

Bill Powers: The aim of scientific objectivity, it is said, is to remove as far as possible all subjective bias on observations of the real world. In the physical sciences, this is done through the use of instruments, reducing observations to simple judgments of coincidences. But I want to put off that part of the subject and look more closely at the concept of observing without bias. The model of perception that is assumed makes a great deal of difference in the meaning of “observing” and of “bias.”

The model I assume is this: the world we experience consists of signals in the brain created by the interaction of the nervous system with the world outside it. This means that neural signals are not *about* the world of experience; they *are* the world of experience. What they are about is another matter that calls for considerable investigation.

If the world we experience exists in the brain, we must then ask what objectivity could possibly mean. I think it means a certain attitude toward experiences, toward perceptions.

If you see a man carrying a briefcase hurrying along under an umbrella through the rain, you can interpret what you see in different ways. You might see a man trying to get to work on time, or someone late for an appointment, or a thief who has just stolen a briefcase and an umbrella. To see these things, you clearly have to add imagined information to what you are actually observing. The same would apply if you saw a man who seemed anxious, or angry, or oblivious to the world. The most objective way of reporting what you see would eliminate all imagined information, all that is not actually in the scene before you.

To be even more objective, you would have to examine the details of what you are seeing. The man seems to be hurrying, but all you are really observing is that he moves more rapidly than others. “Hurrying” is a characterization added to what you see. He seems to be carrying a briefcase, but it could be some other object. “Briefcase” is an interpretation of the shape you see. He seems to be carrying it, but perhaps it is shackled to his wrist. “Carrying” is an interpretation of the relationship between his hand and the object. He seems to be under an umbrella that shields him from the rain, but perhaps the umbrella is a signal to someone he is to meet and isn’t being used to keep him dry.

To be most objective of all, you have to ignore all these characterizations, because no matter how you characterize what you see, the characterization always goes beyond the perception; a different char-

acterization is always possible. To be completely objective, one must simply observe without the accompaniment of an internal explanation or characterization of the observation.

This is almost impossible to do. It is possible, however, to broaden the scope of what one thinks of as observation to include not only the scene being observed, but the internal explanations-interpretations-characterizations that come along with it. If one observes both, then it is clear which set of experiences is the interpretation added to the observation. Or at least it becomes more clear.

So the most completely objective observation is that which is totally subjective and silent. It is simply attending to appearances as they actually present themselves, without any attempt to add to them or manipulate them rationally, without saying that they are real or unreal, without theorizing, associating, or explaining. Doing this to the extent that is possible takes practice and discipline leading to a state of mind much like what Zen practitioners seek through meditation.

Now we can reintroduce the subject of instruments, asking on the way why it is that such instruments are used.

The object of scientific explanation is to explain experience. More exactly, it is to explain why some parts of the experienced world are related to other parts as they are. Why is it that when there is a flash of lightning, there is quite often, after a delay, a roll of thunder? All we experience is the sequence of events; any other relationship between them is hidden. Science is an attempt to guess the connections between the flash of light and the sound, to explain the sound as a natural or necessary result of the process that created the flash of light.

Past experience tells us that the world does not appear exactly the same to everyone; furthermore, observations are inevitably tinged by explanation and interpretation, which creep in under the cover of innocuous words like "hurrying." So to eliminate these subjective differences, science employs instruments.

To measure the flash of light, a scientist would use a photoelectric cell, which responds to light by generating a small current that can be indicated on a meter or recorded on magnetic tape. But what does it mean to say that the photoelectric cell "responds to light"? It means that when the photoelectric cell shows a response, a human observer sees light. When we examine the two things being compared here, we see that they are very similar: both the perception and the meter reading are outcomes of receptor processes, one organic and one inorganic. Both outcomes depend on something else, but the human outcome is a brain signal measured in impulses per second, and the light-meter outcome is the angle of deflection of a needle. Neither outcome is in units of "light."

Instruments, therefore, provide us with consistent indications of

something going on on the other side of the instrument, but they don't identify to us what it is that is being measured. Instrument readings are more objective than eyeball observations only because they are more repeatable and are not influenced by interpretation prior to the reading. They are not more objective in the sense of bringing us closer to a pure description of reality itself. The basic correlate of an instrument reading is not some real physical variable, but either another instrument reading or a human perception.

Now what about the claim that instrument readings are more objective because they reduce observations to a simple discrimination of coincidence? The claim would be that the photocell measurement of light intensity is more objective than the visual estimate because the photocell always responds the same way to the same light intensity. All that the human observer has to do is read the meter face carefully, or, for a digital instrument, write down the number on the display.

But what do we have then? Suppose the reading is 12.5678241. Just writing down that number is reminiscent of the joke that goes, "We interrupt this program to bring you a late score: six to nothing." The number by itself is meaningless. At the very least, you have to know that it is the reading from a photocell, not from a thermometer. To use it in relation to any other meter reading, you must know how the meter is calibrated: what are the units of this number? Foot-candles? Lumens? Ergs per second? What is the spectral range being measured? And to use this reading in the context of science, you must also explain what it is that is being measured: the absorbed part of a flux of photons, a flow of energy, a squared amplitude of magnetic and electrical vibrations at a certain frequency or with a certain wavelength. You must, in short, reveal the complete model of what the meter supposedly measures.

There is no way, in fact, to reduce an observation to a coincidence of a meter needle with a mark. If it is reduced that far, it ceases to mean anything.

What gives meaning to the meter reading is exactly the same thing that gives meaning to an uninterpreted human perceptual experience. It is the structure of interpretations and theories that depicts a world on the other side of the receptors. That world does not exist in unadorned, uncommented observation. It exists only in the adornments and comments added by human intellectual processes.

When we try to understand human perception scientifically, we automatically introduce something other than direct experience. We introduce a term "perception," indicating that there is a perceiver, a consequence of perceiving, and something to be perceived. This is like introducing a photocell, a photocell reading, and light-energy. There is an "inside" component and an "outside" component linked by a physical device. All three of these components are theoretical entities,

part of a model of processes that underlie direct experience.

The popular conception of a light-meter is that the light is what exists, while the meter reading is only an indication of it. But considering how scientific modeling actually works, the priorities must be reversed. It is the meter reading that is given; what it indicates and how the indication is derived are matters for theory and conjecture. We must reason backward from the meter reading, taking into account the theoretical properties of the meter and the photocell (and doing the same for many other kinds of meters), to deduce what lies at the origin of the reading: reality, the world.

In a model of perception itself, we must do the same thing. The physical device we place between reality and perception is the nervous system. The properties of the nervous system are a matter of theory and interpretation of observations. The given is the perception, the experienced part of the process: the way the world appears. What remains to be deduced, by reasoning backward through the assumed properties of the perceiving device, is the external world.

It is possible to observe objectively only the outcome of this theoretical perceptual process. We can see whether the model consisting of a nervous system with its properties and a physical world with its properties (derived through studies with instruments) can be made to produce an outcome that matches what is in fact experienced. Objectivity then consists in observing what the model actually does, with as little interpretation as possible, and how the world actually appears, also with as little interpretation as possible. But objectivity has nothing to do with reporting on the world that is represented by our models.

Wayne Herschberger: Bill says that “neural signals are not *about* the world of experience; they *are* the world of experience.” Almost, but not quite. Any claim that neural signals are *about* the world of experience is undoubtedly false, as you say, precisely because the reverse is the case: my world of experience is *about* my neural signals—and about the optic array in the ambient light, and about the stuff with resting mass which is said to give structure to the ambient light. It is not true that my experiences comprise neural signals. Rather, to quote an authority you might recognize, “this world presents itself... in three dimensions, stereo sound, and living color, chock full from edge to edge of continuously-present smoothly changing noise-free colors, shades, objects, motions, relationships, and operations in progress.”

According to a coherence theory of truth, the perceptual objects comprising our experience (phenomenal world) may be said to depend upon (be about) the conceptual objects we “construct” (neurons, photons, electrons, input functions) to the degree and only to the degree that these conceptions (models) account parsimoniously for the

perceptions in question.

Perception is not imagination.

Bill Powers: Wayne, I was not describing what the world *is*, I was describing how it *appears*. This appearance is the world we directly apprehend. The control-theory model, based on the appearances of neurology and physics, leads me to conclude that this world is, physically, a collection of neural signals, although it does not look that way (that is, it does not look the way neural signals do when we visualize them on an oscillogram using electronic means of sampling limited aspects of neural activity).

I know that you insist that there is no other world than the world of appearances (or that we don't need to consider one in explaining perception). When we aren't concerned with explaining, but only with experiencing and living (no theoreticians required), I agree with you. Reality is precisely what we experience. But to add "and nothing more" is to assert what we can't know directly, and to deny, for no good reason, the implications of physics, chemistry, and neurology, all of which claim to represent a world of immense detail that is inaccessible to our senses.

The crucial difference between our views, as nearly as I can comprehend it, is in my assumption of another reality that is not part of the world of direct experience. If there is no such separate reality, then of course all that exists in nature is experience as it appears to us. Neural signals, physics, chemistry, and all such conceptions are just that: conceptions and nothing more. They are simply ways of ordering our experiences and have no significance beyond that. They do not refer to anything unexperienced.

I can understand that conclusion as a conclusion, but I can't accept it as a fact. It is simply another conjecture. Its truth, then, comes down to the evidence we have for and against it, and to how we reason about that evidence.

I think that control phenomena provide us with evidence that there is a universe beyond the limits of human perception; that this universe imposes its properties between our actions and their perceived results; that there are independent agencies in this universe that are capable of disturbing our control actions without our being able to detect the causes of the disturbances. I think we learn about these properties only indirectly, and as conditioned by the kinds of perceptual systems we have and do not have. I think we have to infer the nature of the disturbing agencies and the properties of the world, by building models that would, if they were true representations of the unseen world, explain how our experiences are related to each other. I do not think it is likely that we have arrived at models that just happen to capture every

significant entity outside us, every significant functional relationship among those entities. And I do not think it is likely that the world of direct experience exhausts the degrees of freedom that really exist in the universe around us.

So that is my basis for accepting, as the most reasonable hypothesis, the existence of a real universe apart from our perception of it, and for denying, on the basis of the same reasoning, that our perceptions are likely to be veridical renditions of that universe.

I agree that perception is not imagination. Imagination is, however, a subset of perception. Some of our experiences are generated inside the brain and do not depend on the current external state of affairs, even though they might sometimes give a convincing imitation. But the rest do depend on something outside. In neither case, however, do perceptions without the aid of reason give us a picture of what is really causing them—however inadequate the picture.

Wayne Hershberger: Bill, it would be a comfort were you to agree with me, but I do not agree with the statement, "Reality is precisely what we experience." Of course, I agree even less with the obverse idea.

That is, as an empiricist, I do not endorse the idea that reality is precisely what *cannot* be accessed empirically. Hence, I cannot imagine ever saying that physical, chemical, or neurological phenomena are epistemically inaccessible. In fact, I have been championing the antithesis.

I guess I am not making myself clear.

Let's go to square one. I claim that the epistemological challenge is not to explain how the truly inaccessible can be accessed (a logical impossibility), but rather to explain how the truly accessible can have appeared to be inaccessible (a logical possibility). Do you agree?

Bill Powers: Square one it is. I agree that the truly inaccessible can't be accessed. That's a definition, not a proposition. The truly inaccessible is that to which we have no access at all, either direct or indirect.

The second part of your claim is not a definition: it asserts that the accessible sometimes appears inaccessible. I think this is an attempt to create a two-valued situation (either something is accessible or it is not) out of one that has more than two possible values.

I see accessibility of the workings of nature to be a matter of degree, with the maximum degree falling short of 100%. To explain my view, I will resort to a thought-experiment.

Suppose we have before us a black box which we have no means of opening. Let's call this box, to humor me, the Reality Box. On its surface are numerous buttons and lights. The buttons and the lights are undoubtedly real, because we experience them directly and unequivocally.

cally: they are totally accessible and cannot be mistaken as being inaccessible to our observation. So the box, the buttons, and the lights are not an epistemological problem.

When we press various buttons, we find that certain lights and combinations of lights turn on or off. With sufficient experimentation and record-keeping, we can discover consistent effects of the buttons on the lights. As our experience grows, we can discover that some buttons alter the effects of other buttons on lights, or make certain lights come to depend for their state on new combinations of buttons. We can discover that only certain sequences of button-pressing will have predictable effects on one or more lights. We can find that certain lights have mutually exclusive states; if a member of one set is on, another set is always off, and vice versa. We could uncover logical relationships, relationships corresponding to arithmetic operations, and so on. We could even develop heuristics: some ways of turning lights on work best under one set of circumstances (combinations of lit lights), usually, than other ways.

Thus, we arrive eventually at a sophisticated empirical understanding of the Reality Box. At no point have we asked what is inside the box. We have simply observed, recorded, and tried to recognize consistencies.

Perhaps I should motivate this investigation by saying that for reasons we only vaguely understand, certain of the lights on the Reality Box have extraordinary value to us; indeed, their states of illumination seem to us to be a matter of life and death, or at least make the difference between enjoyment and disappointment. So we have an interest in pushing the buttons to keep the most important lights in the states that seem the most desirable, especially as they will not stay in those states without the button-pushing. To be crass and less mysterious, we could say that each time we succeed in maintaining the critical lights in the critical state for one minute, we receive \$5—that's \$300 per hour for this job if we can learn to do it perfectly.

I think I have now described the state of human understanding of nature in the pre-Galilean era.

In fact, we find that we are a long way from making \$300 per hour—the actual payoff isn't nil, but it's just barely a minimum wage. It would be greatly to our advantage if we were allowed to cheat: to open the box and trace out the circuits (or talk to the little men, or analyze the chemicals, or take whatever action is appropriate to what we find in the box). If we knew *why* the buttons affect the lights as they do, we could abandon the trial-and-error empirical approach and simply deduce the actions that would have the effect we want.

I am now describing the advent of the physical sciences.

We are not, however, allowed to cheat. Nobody knows how to open

the box.

Nevertheless, once we get the idea of explaining the dependence of lights on buttons rather than just observing it, we might well decide that even a good guess about what is in the box might be more valuable than random experimentation. So we begin to construct a model of the internal workings of the Reality Box, trying to outguess its designer.

This project turns out to be extraordinarily successful. By imagining circuits and functional devices inside the box, we succeed immediately in explaining why some buttons cause some lights to change their brightnesses. Numerous revisions of the model are required, however, because just when we think we have the right connections, an anomaly turns up and we have to modify the design of the hidden devices or the connectivity between them. But by demanding that the model *always* work, no matter what combinations of buttons we press, we eventually get this model to the point where it never fails in any way we can notice.

We now begin to believe in the reality of the model. What appears to be happening is that buttons activate lights, but what is *really* happening, we say, is that the buttons are feeding their effects into a hidden complex device that in turn operates the lights. Gradually, the status of the insides of the Reality Box changes. Those insides no longer seem hidden to us. In fact, even though they are complex, they are far simpler than our records of empirical findings are. They also permit us to predict the effects of button-pushings, even combinations never tried before, with exceeding accuracy, whereas our empirical predictions, based only on unexplained frequencies of occurrence, are wrong nearly as often as they are right, and are essentially useless in unfamiliar circumstances.

I have now described the rise and maturation of the physical sciences, and their essential difference from the purely empirical sciences.

The penultimate stage in this development arises when someone notices a fact that by now is considered a very strange fact. Those who are engaged in the exploration of the Reality Box by now feel that its devices and connections are *in the box*. The lights tell them what is happening inside the box; the buttons let them influence what is happening inside the box. What this someone says that is thought so strange is merely a reminder that, in the beginning, nobody knew what was in the box, because only the lights could actually be observed, and no effect of pressing the buttons could be seen except in the lights. This is, in fact, still the case. So the model of what is in the box must exist in the minds of those who are observing the box. It is not in the box. In fact, it is perfectly possible that what is in the box is entirely different from what is in the mental model, but is equivalent to what is in the mental model under all of the button-pushing operations so far tried. Even

what seems to be a simple direct connection through a hidden wire might actually involve a hidden modulator that converts the button-press to a radio frequency and broadcasts to a receiver whose output lights the light. That would not be a parsimonious design, of course, but it might be the one that exists.

What this upsetting stranger is doing is reminding everyone that all they can actually observe are the lights, and the only effect they can know they are having is to press the buttons. All the rest is imaginary. Therefore, we should throw away all those figments of the imagination and admit that all we know is how the buttons affect the lights, and to remain pure of heart, we should talk about nothing else.

I have now described the advent of stimulus-response theory, behaviorism, biology, empirical psychology, and so on.

The final stage entails the epistemology of the Reality Box that I propose.

In fact, the model works much better than it should. Moreover, there is evidence in the relationship between buttons and lights that tells us something consistent is going on independently in the Reality Box. The lights that we can affect with our buttons sometimes turn on and off when we aren't pushing anything. Very often, we have to change which buttons we push in order to reproduce the same state of the lights, and there seems to be no way to predict when, by how much, or in what direction we will have to make these changes. Something else is interfering with the effects that the buttons have. This something else can be inferred, to some extent, because it might occur regularly, or in some regular pattern, as we can tell by watching what different buttons we have to press to reproduce the same effect, and when we have to do this.

So we are led, in the end, to recognize three major facts. First, our mental models of what is in the Reality Box have an unknown relationship to what is actually there. Second, the regularities implied by the model actually do occur, even though we can't know that they occur for the reasons we propose. And third, there is something in the Reality Box that can act independently of us. So we can say that in some regards, what is in the box is accessible to us, but we must also admit that our interpretation of its inner workings is not necessarily the only one that would be as good at explaining what happens.

The lights, of course, are our perceptions, and the buttons are our actions. The Reality Box itself is invisible; we experience only the input and output devices mounted on its surface. We conjecture that the buttons do something that we don't observe. We conjecture that the lights indicate something that is also not observed—if only the presence of a wire from the button to the light, and an invisible power supply.

So what does “accessible” mean? Does it mean that we observe

Reality exactly as it is, or that there is neither agency nor order other than what is evident to us in our sensory experiences? I feel that such questions are not matters to be deduced logically, so that we can know once and for all the truth about experience. I think that they are matters to be settled as we settle all factual questions: by the examination of evidence, and by settling for the inference from the evidence that seems most supportable by all the rest of our experiences and knowledge. Pure philosophy can't provide that sort of conclusion: it demands an end-point, a certainty. That, I think, is definitely inaccessible.

Wayne Herschberger: Bill, while agreeing that the truly inaccessible is inaccessible, you assert that much of that which is accessible is not all that accessible. Here you are changing the subject (I'm assuming that you are not simply contradicting yourself). That is, you are using the word accessible to refer not to the possibility of epistemic access, but to the amount of X that is accessible, or to the difficulty of achieving access to X, which could vary, of course, with the directness or complexity of the epistemic process. This is as misleading as referring to the length of a pregnancy as a degree of pregnancy (e.g., the unwed mother who claims she is "just a little pregnant"). More specifically, if 90% of X is truly accessible and 10% is truly inaccessible, it does not follow that X (i.e., every bit of it) is 90% accessible.

You say: "... the box, the buttons, and the lights are not an epistemological problem." On the contrary, the epistemological problem is usually stated in just such terms; for instance, it might be said that although the buttons "appear" to be solid and stationary, they "really" comprise a swarm of whirling dervishes known as atoms. Or alternatively, it might be said, as you are wont to do, that they really comprise a collection of neural signals. In fact, the buttons, lights, etc. are the phenomenal objects that the empirical process we call perception provides us, and the epistemological question concerns whether or not these phenomenal objects are as objective as the label "object" implies. The box, the buttons, and the lights pose the epistemological problem! Your elaborate example begs the question.

To suppose that the phenomenal objects are accessible appearances comprising indirect representations of an inaccessible reality is to embrace a radical skepticism, because there is no way to assess the fidelity of the representation without accessing the inaccessible (i.e., in order to test the correspondence between the reality and the appearance, one needs access to both, and that, by definition, is not possible). This question of correspondence between what is accessible and what is inaccessible (i.e., between what is internal to and what is external to the limits of experience) is readily confused with the correspondence between what is internal to and what is external to the nervous sys-

tem. But whereas the former type of correspondence is impossible (by definition), the latter type of correspondence is easily determined—neurophysiologists do it all the time.

You say, “What this upsetting stranger is doing is reminding everyone that all they can actually observe are the lights, and the only effect they can know they are having is to press the buttons. All the rest is imaginary. Therefore we should throw away all those figments of the imagination, and admit that all we know is how the buttons affect the lights, and to remain pure of heart, we should talk about nothing else.” If I get your meaning, that stranger is no stranger. His name is George Berkeley. Later, Johannes Muller echoed Berkeley’s Subjective Idealism in his doctrine of specific nerve energies; as Muller put it (in his article 5), the sensorium is aware not of the external object, but of the state of the nerves only. Having said that, Muller then seemed to recognize belatedly that Berkeley’s thesis makes no sense expressed in physiological terms, because in article 8 he said that the sensorium is aware not merely of the state of the nerves, but of the external causes as well. (There’s nothing like having your cake and eating it too.) Muller was confusing the two correspondence questions described above. Berkeley’s philosophy concerns only the former type, as Kant observed; that is, the only objects we experience are phenomenal, not noumenal, things.

I share your concern with the essential nature of nature. That is, when an experiment asks a question of nature, “someone” answers. But I see no reason to exclude this final arbiter of empirical truth from the phenomenal domain. Banishing this arbiter to an inaccessible realm from which it creates accessible appearances (like the Wizard of Oz) makes about as much sense to me as claiming that today is but a representation of yesterday’s *real tomorrow*. What does it buy one, but a big headache? Who needs it? What’s wrong with immanent truth, as reflected in phenomenal coherence?

I am interested in your staking a claim to what I see to be the epistemological high ground (in my view, solipsism is *not* the high ground). My motivation is selfish. Because you are the principal champion of psychomodular control theory, I have a vested interest in your being in the best position to defend both your psychomodular theory and your epistemology. And since I do not see your psychomodular control theory as implying a solipsistic epistemology, I see no reason for you to defend that indefensible epistemological position. My inability to persuade you to give solipsism a wider berth than you do leaves me ambivalent about my efforts to that effect, because I am not interested in being a mere disturbance.

Bill Powers: A model might be epistemically correct, but we will never be able to prove that. I was not saying that we have complete epistemic

success part of the time, and incomplete success the rest of the time. I was saying that we do not *ever* know whether our models of reality are successful or not, because the only way to check for that success is to repeat the process that led to the models in the first place: there is no independent check.

Perhaps we could get to the nub of this matter sooner if you would give me one example—any example—of a case in which we have complete epistemic success in verifying that there is a real counterpart of any perception. We apparently agree that there are some cases in which uncertainty remains, so there is no point in dealing with them. My claim is that there is *no* case in which we have reached certainty, so you should be able to demolish my claim with a single counterexample.

Wayne Hershberger: I seem not to be making myself clear; you are looking for me in the wrong direction. My argument (actually, Hume's) is not that any particular case affords certainty, but rather that *every case is entirely uncertain* (i.e., "verifying that there is a *real* counterpart of a perception" is absolutely impossible, even as a matter of degree). Hume's arguments to this effect are called Radical Skepticism for good reason. Consequently, modern science uses a coherence, rather than a correspondence, theory of truth—where reality has no ontological status.

When asked how he discovered the laws of chemical compounding, Linus Pauling replied, "I made them up." Pauling avoided any claim of having gained epistemic access to Reality—because such a claim would be gratuitous (God's Reality is a matter of faith, and serves no scientific purpose). Rather, Pauling made up a parsimonious model which provides a very *coherent* account of the chemical phenomena in question. Any claim that such human-made models correspond, in varying degrees, to some divine original is epistemically empty.

You have made up a parsimonious model of living control systems which control the value of inputs from their environments. Sometimes you have used the word "virtual" to refer to these controlled variables, because they are defined by the input functions which process the input. But it would overstate the case to claim that the environment contributes nothing to the values comprising these virtual variables. That is, only by overstating the case is one misled to suppose that your model implies solipsism. Your model addresses questions of correspondence between what is inside the brain and what is outside the brain, but that is physiology, not philosophy.

Bill Powers: Comes the dawn. I feel like a wrestler who has converted an advantage to a position flat on his back. Your previous arguments have given me the impression that you believe there is a reality outside

of perception *and* that perception somehow manages to represent it veridically. Now it seems that you're saying that human perception bears no verifiable relationship to any universe "behind" or "beyond" perception, which is, of course, the position I have also been taking.

Unfortunately, your language still leaves me wondering what precisely is your position on the constraints we can detect between actions and perceptions, and on the significance of models. You say, "Any claim that such human-made models correspond, in varying degrees, to some divine original is epistemically empty."

If it weren't for that word "divine," I would be more sure of how to take what you mean—I trust you're not accusing me of religious fervor. Would you still allow for a correspondence to a "non-divine original"—i.e., a natural universe that exists independently of our perceptions of it? In other words, are you opposing a religious view of reality, or any view that there is (or could be) a reality more inclusive than what is perceived, whether or not we can be certain about its nature? More on this at the end.

The main conundrum comes up when you say, "Your model addresses questions of correspondence between what is inside the brain and what is outside the brain, but that is physiology, not philosophy." A problem is created by talking about what is inside the brain and what is outside the brain. The problem arises when we assume that we, as conscious beings, are conscious because of activities in a brain. Allow me to elaborate.

In order for my model to be consistent with the physical model of reality (both, I quite agree, being "made up," so that physiology, too, is "made up"), there are certain relationships between physical-model variables and neural-model variables that must be assumed. The physiological neural model allows for no way of getting information from physical stimulation other than through interaction of physical variables with neural sensors. For example, in the physical model there is a made-up entity called the photon. The signals in the neural model's retina supposedly arise from absorption of photons. However, the neural signals carry no information about the origins of those photons; furthermore, there is an infinite number of different photon energies and fluxes arriving from an infinite number of directions that will yield exactly the same neural signal in any given receptor.

Given a model of physical optics and observations of reflection or emission sources, we can construct a model of the origins of the photons and show that this model is consistent with an array of neural signals that amounts to a map of the scene toward which the eye's lens is directed. So far, so good. But if we then look at the basis for accepting the physical model—which includes things like "lenses" and "objects" and "light rays" going through something called "three dimensional

space”—we find that there can be no basis but observations made by the same means by which an “observation” of a photon is made. We identify objects by looking at them with our eyes; we verify that there is a photon flux by interposing a light-meter (which we see) and reading—with our eyes—its indicator. So, from the standpoint of the neural model, the physical observations we are using to assign an external source to the visual neural signals arrive in the brain by exactly the same means as the signals we are trying to explain.

This is not a problem if we adopt a point of view from which we can see both models, the model constituting an exterior physical world and the model constituting an interior world of signals in a brain. It is not a problem if we add to the other two models a model of a non-neural conscious observer which is not confined to a brain. It becomes a problem only when we decide that the model of the brain must be a model of ourselves, the observers and thinkers.

When we adopt that view, as I do, we can no longer take the third-party omniscient view. The hypothesis is that we, who are thinking about perception, are brains like those in the model. Therefore, we must be dealing with the external world (represented by the physics model) by the same means just proposed: through neural signals. If this is true, then the physical model is *not* outside the brain. It must be located inside the brain-model, as part of that model. It is a construction existing as patterns of neural signals related not by physical constraints outside us, but by abstract rules and computational processes taking place in our heads.

When we apply this reasoning to purely physical models, there isn't much difficulty except with people who insist on reifying photons, electrons, quanta, phlogiston, and so on. The real difficulty arises when the external world we are thinking about is the world of subjective reality: the world we experience directly. This is clearly not the world of the physical models. Between the physical models and this world of direct experience, there are few points of contact. For the most part, physical models consist of entities and relationships that are not evident in direct experience. Here and there are points where, usually through the use of instruments, but not always, a physical variable corresponds to an experientiable variable. With the unaided eye, we can perceive an approximation to what a physicist calls “distance,” although by using instrumentation like radar or optical range-finders, we can arrive at meter-readings much more consistent with physical theory than is the direct apprehension of distance. But when it comes to functions of distance such as gravitational acceleration or potential energy, direct experience remains blind.

So where do we put this world of direct experience, with all its objects and sounds and smells and relationships and people? It is not

represented in either the physical model or the (physiological) neural model. I have elected to put it into the neural model, but not in the form of neurons. It exists in the brain as a weightless, massless organization of neural signals, the appearance and behavior of which is precisely the appearance and behavior of the world we experience. Certainly this assumption creates a mystery; more than one. The main one is who or what is it that apprehends this collection of neural signals in such a way that it takes on the appearance that we experience?

The most obvious error to be made at this point is to say that this mysterious observer is the agent who imposes interpretations on the neural signals so they become objects, relationships, processes, concepts, and so on. But as everyone knows by now, that simply requires expanding the model to explain how these interpretations are made. My way of avoiding this error has been to propose the levels of perception in my model. By looking for classes of perception in the apparently real world around me, I attempted to show how neural processes can themselves create signals which contain the interpretations that are needed. While this initially seems to rob experience of some vital qualities, a close examination of any particular example of this problem shows that it does not exist. These vital qualities can't be pinned down by direct inspection, either. When one attempts to isolate them for a close look, they lose any special quality and become just an amount of something that can be more or less present. Just like a neural signal. The only specialness that there is exists in the entire collection of neural signals, each behaving in the context of all the others.

The other function of the levels is to enumerate and classify types of perceptions ranging between what have been considered "concrete" and "abstract" perceptions. By showing how successive levels of interpretation can form a link between the concrete and the abstract, the model removes the necessity for explaining these interpretations by assigning them to a homunculus. As each new level is considered, the subject-matter with which it deals is stripped out of the homunculus and returned to the physical brain. In the end, the homunculus contains only those functions of observation that are not accomplished by the brain model. And all that is left is awareness.

We now seem to have a model, itself a neural model, that contains a physical model in a nervous system. But the nervous system is basically a physical conception: it is a subset of the physical model. Logically, if one model is contained by and contains another model, there can be only one model. But there is another answer: it is that both models are contained in direct experience, and they interact with each other..

The ultimate reality, therefore, is direct experience. That is the super-class within which models exist. This leads us, finally, to the ultimate mystery.

We can divide direct experience into things we do and things that happen. Many things that happen proceed without any need for our action. Among such things, we can detect consistencies and dependencies. This leads us to formulate expectancies which, when formalized, we call laws of nature.

We can also take actions, which are the set of all those experiences that we can influence by an act of will. We find that these actions, themselves capable of being experienced, affect other experiences. We can learn to create some experiences that are not directly subject to acts of will by varying those experiences that are directly willable: the whole is an act of control. Through long experience with this kind of act of control, we have found regularities that show how we must act in order to control many kinds of experiences. The reasons for these regularities are not evident in experience—there is no a priori basis for expecting any particular act to have any particular effect on something else.

This is where we get the idea of a natural world of regularity that lies outside the boundaries of experience. And this is why we build models, both physical and neural. With models, we hope to probe into that mystery that is hinted at by these unexplained regularities. We hope to reduce the complexities of these piecemeal regularities by finding underlying simplicity; I think this is what we mean by “mechanism.” In physics, simplicity is attained by imagining a hidden world of fields and particles, energy and momentum and entropy. The few kinds of variables in this world lead to the vast multiplicity of different-seeming phenomena in the world of direct experience.

The question, Wayne, that you and I have not brought out into the open and resolved between us, is whether these models constitute increasingly good approximations to something beyond experience, or whether they are simply “summaries of observations.” The complex picture I have tried to lay out here should indicate my view. Clearly, I don’t think that either physics or neurology is as good an approximation as is usually assumed. There is too much of the human observer entwined, unanalyzed, in all our models. The very name “particle” in physics shows this. But I think that there is evidence of agency outside us (other people, for example), and evidence of relationships imposed by unseen means (e.g., other people’s intentions). I think that there is structure inside the Reality Box, and that while we can never arrive at a unique representation of it, we can arrive at an equivalent representation, equivalent in the sense that our models show one way it could be constructed inside, functionally equivalent to the way it is constructed. I see no contradiction in saying that all we will ever know for certain is what our own brains present to awareness, while maintaining that uncertain knowledge is not empty.

Is the remaining problem, perhaps, what is meant by “epistemic”? If

“epistemic” knowledge is certain knowledge, then the argument resolves itself: there is no such thing outside direct experience. But to say that a proposition is epistemically empty does not then mean that it is incorrect or empty of significance, because that would say that all of experience is, with complete certainty, incorrect or empty of significance beyond itself—an epistemic fact which, of course, we can never verify.

Martin Taylor: Bill, in answering Wayne, says: “But if we then look at the basis for accepting the physical model—which includes things like ‘lenses’ and ‘objects’ and ‘light rays’ going through something called ‘three-dimensional space’—we find that there can be no basis but observations made by the same means by which an “observation” of a photon is made. We identify objects by looking at them with our eyes; we verify that there is a photon flux by interposing a light-meter (which we see) and reading—with our eyes—its indicator. So, from the standpoint of the neural model, the physical observations we are using to assign an external source to the visual neural signals arrive in the brain by exactly the same means as the signals we are trying to explain.” But one actually can test the “existence” of the things detected through photons by using other senses—the acoustic effects and so forth. These form a set of converging operations that help to reduce the set of possibilities for interpreting the perceptions obtained through one sensory system.

If that were all there was to it, the same argument could be made, but extending the notion from “photon” to “physical energy exchange phenomena” or some such. But that is not all there is to it. There is the volitional aspect of what and how we choose to observe.

Let us presume a deceitful Nature, and a passive (multi-sensory) observer. This Nature could present us with any of an infinite number of sources for PEEPs (Physical Energy Exchange Phenomenon, plural) that had the same effect on our sensory organs. But when *we* choose which aspects of the universe to test, and in what way, the deceit becomes much harder to sustain. That’s the fundamental difference between an observational science like astronomy and an experimental one like physics. Psychology is somewhere in between.

The difference between active observation and passive observation was clear to Gibson, who distinguished “haptic” from “tactile” perception. You can try it yourself. Have a bunch of objects available, and a friend. Close your eyes; have your friend take one of the objects and touch it to your open hand in various orientations and ways. What you perceive is a set of touches, some soft, some warm, some sharp, and so forth. Now have the friend place the object in your hand for you to manipulate. What you perceive is not a set of touches, but an object.

In either situation, you might be able to determine which object from the set was touching your hand, but when you yourself choose where and how the object contacts (or fails to contact) your hand, it has a completely different subjective quality.

Wayne Hershberger: Bill says, "... it seems that you're saying that human perception bears no verifiable relationship to any universe "behind" or "beyond" perception, which is, of course, the position I have also been taking."

Yes.

And that, I believe, is what Gibson was saying as well, or at least trying to say. And Kant, too. Your prose is more lucid than most of theirs, but as I read you (Hume, Kant, Gibson, and yourself), you all seem to be motivated by this same epistemological insight. However, the other three believed that this insight also implies that this hypothetical Reality has no empirical basis. So did Plato, who claimed that Reality's basis must, therefore, be rational.

You, on the other hand, seem to accept a Reality appearing to require neither basis. That is, the Reality to which you persistently refer, despite your above remark, appears to be neither an induction nor a deduction, but rather an abduction; you seem to pluck it out of thin air. For instance, while admitting that a perception cannot be proven to be a veridical representation of Reality, you are wont to claim that neither can it be proven that perceptions do not approximate Reality to some degree. This begs the question of the Reality itself!

I would encourage you to accept the harsh implications of your own epistemological insight and not backslide, admitting through a back door what you have banished from the front. The challenge, remaining to be addressed, is what exactly is meant by the term real or true, that a perception may be identified as veridical, as opposed to illusory, *and* at the same time *not* be regarded as a representation of Reality? That is, what is it that distinguishes a veridical from an illusory perception if *not* the perception's degree of correspondence to some transcendent reality—something you claim to eschew as an arbiter of truth?

In practice, we seem to use a coherence theory of truth. Laymen and scientists alike regard a perception that cannot be replicated as illusory. A perception that does not survive the layman's double take is an illusion. An empirical observation that science cannot replicate is no fact.

You say, "So, from the standpoint of the neural model, the physical observations we are using to assign an external source to the visual neural signals arrive in the brain by exactly the same means as the signals we are trying to explain." Although this is often the case, it is not necessarily the case. Let me address the flip side of this question—to

which you also allude: is there “a natural universe that exists independently of our perceptions of it?”

I submit that the ability to register luminous flux with virtually any retina or photomultiplier tube provides the very sort of independence referred to here. Only those perceptions which are demonstrably replicable across observers are objective perceptions, or “objects,” as we are wont to say, for short. This is the type of independence required of the objects comprising our *natural universe*. When it does not matter who or what makes the observation (i.e., the results are independent of the particular observer), the perception is said to be objective, or to be an object. The natural order is immanent in experience and not to be confused with some hypothetical Reality that transcends all experience.

You ask, “So where do we put this world of direct experience, with all its objects and sounds and smells and relationships and people?” In the phenomenal world of time and space—which Kant recognized as intuitions (meta-models?). “Silicon Babies,” an article in the December 1991 issue of *Scientific American* said something relevant here. Speaking of robots as Rodney A. Brooks conceives them (which is similar to the way you conceive them), the author of the piece said: “Subsumption architecture relies largely on the nature of the outside world rather than sophisticated reasoning to structure the robot’s actions. For example, if the robot encounters an obstacle, the important thing is to go around it... The robot may not need even to remember that the object is there—after all, it will detect the obstacle perfectly well the next time it approaches it. (p. 128) The expression “outside world” in this passage obviously refers to the robot’s environment. This world outside the robot is not outside the robot’s realm of experience. Neither is the robot’s world a re-presentation (copy or memory) in the robot of a world actually transcending its experience such as our model of its environment. The robot merely *registers* its environment in its “inimitable way.” Call it modeling the environment. Of course, the way a robot registers/models its environment is not actually inimitable; identical robots would register/model in the same way.

Similarly, if your expression “a natural universe that exists independently of our perceptions of it” refers to something other than the natural order immanent in the psychophysical flux we call experience, it is surely a reference to the perceptual/conceptual models that are constructed out of that flux registering that immanent order.

“So where do we put this world of direct experience, with all its objects and sounds and smells and relationships and people? It is not represented in either the physical model or the (physiological) neural model.” Right you are. The objects of direct experience are not part of either of these scientific models (physics or physiology), because these objects of direct experience are themselves models—empirical, if not sci-

entific. They are the layman's perceptual models, analogous to science's conceptual models; both types seem to involve a lot of neural processing, just as your own theoretical/scientific model says. Both types of models are modeling the same natural order. They are twin born of experience. One type of model is not modeling the other type of model. (Only psychological theories such as your model are reflective, modeling the process of modeling itself.) Neither is the basis of the other. The basis of both is the natural order which, as you say, "exists in the entire collection of neural signals, each behaving in the context of all the others." Strike the word "neural," and I think I could buy it. That is, if this psychophysical flux has any essential characteristic, surely it involves informing and being informed—in a word, signaling. Matter which can neither influence nor be influenced by other matter, doesn't matter. But the signaling does not begin and end in the nervous system.

Bill Powers: Wayne, you are telling me that Hume, Kant, Gibson, and I all seem to be "motivated by this same epistemological insight"—yet "... the other three believed that this insight also implies that this hypothetical Reality has no empirical basis." They might not have had the same attitude toward empiricism that I have. For instance, they might have been of the opinion that only empirical facts can be true and real. This would put them in good company, but it would be the company of those who customarily elevate statistical preponderances to universal certainties. I have met very few empirical facts that did not contain easily discernible uncertainty; certainty is achieved by ignoring the actual data and plunking a dot down in the middle of the scatter: that's the *real* value. Most of the time, there is a background of approximations, arbitrary assumptions, and interpretations without which empirical data would have no meaning. These assumptions, and the fact that someone is interpreting, are not mentioned in polite company.

So to say that a hypothetical reality has no empirical basis is not the indictment it might be if there were such a thing as pure empiricism untainted by human imagination and interpretation.

Even when we confine our observations to the omnipresent psychophysical flux, we see things that are contradictory. Our judgments of width and height do not agree with readings from calipers. Our judgments of straightness do not agree with straight-edges. Our judgments of relative temperature do not agree with thermometers. Our judgments of relative brightness do not agree with photometers, and our judgments of relative color and, especially, color composition do not agree with spectrographs. In realms of more complex observations, we do not agree with each other about palatability, difficulty, comprehensibility, spelling, grammar, or miracles. We observe nonexistent phantom arrays created by what we know to be a single stationary flashing light. We

don't even agree on a color like "green." When we do agree, the spectrograph can tell us we are looking at different colors; when we disagree, the spectrograph can tell us we are looking at the same color.

You say, "An empirical observation that science cannot replicate is no fact." Is there *any* empirical observation that science can, literally, "replicate"? Replication never in fact occurs: perhaps that is the *only* replicable fact. What happens is that we make a series of meter readings that disagree with each other, and then we say that there is a Real value that lies somewhere within the range of the readings. We can't actually "replicate" a reading (in fact, if we get exactly the same reading twice in a row, we tap the meter from then on). We can't even replicate the scatter in a series of readings. We replace the scattered, variable, inexplicable, individual observations with an idealization that we conceive of as the real observation. In doing this, we create precisely the reality I am talking about: a reality that we do *not* observe, but accept because it makes sense of experience. Empiricism itself leads to acceptance of a reality that underlies observation, but is not the same as what we observe.

Much of what seems to be replication is a product of the human capacity for categorization. We can make observations that vary widely, yet make them appear to be the same by classing them together. If we ask 100 people, "Are you in favor of abortion?" and 60 percent of them say "no," we lump the 60 "no" answers together and say that they indicate the same opinion about abortion. In fact, we don't know what question the respondents were actually answering; all we know is the question we heard ourselves asking, and what it implies to ourselves. Some were thinking, "No, not even to save the life of the mother," while others were thinking, "If it's necessary to save the life of the mother, but in general, no." Those are both "no" answers, aren't they? "No" is "No," isn't it?.

'Empiricism fails as soon as you go beyond a description of a snapshot of the psychophysical flux. Take as simple a thing as a lever with the fulcrum in the middle. Pushing down on one end, you observe that the other end goes up, while the lever itself tilts. You can easily satisfy yourself that if you do push one end down, the other end will rise, and if you don't push that end down, the other end will stay where it is. There is little more to be determined, empirically, about the behavior of this lever. Now, is there any connection between your pushing on one end of the lever and the subsequent tilt of the lever and the rise of the other end? Have you given an adequate account of the lever by reporting just the facts of what happens in the psychophysical flux? Obviously not. You have reported three facts: pushing down, tilting, and rising. Is the first fact directly influencing the other two facts? Of course not: at least one physical property of the lever, its rigidity, is

needed in order for the one fact to lead to the others. Facts do not influence other facts just by existing. Given two identical-looking levers, one may behave as you expect, while the other simply bends. The difference that explains the difference in behavior is not among the empirically observed facts. It is an imagined property of the lever, deduced from its observed behavior. That property is part of the unseen reality of which I speak. Does this “rigidity” actually exist? It might. But it might also be a consequence of unseen factors such as intermolecular forces, none of which itself is “rigidity,” but which together have the consequence of imparting rigidity to the lever.

The “physical” part of the “psychophysical flow” is imagined or deduced (it is deduced, then imagined). It is not observed. There is no meter that measures rigidity, or whatever factors give rise to it. And if there were a meter measuring the straightness of the lever, it would not be measuring the causes of rigidity.

Or consider another case that tells us even more about our relationship to reality: a wall switch that operates a light in the ceiling. We can easily determine that when the switch is up, the light is on, and when it is down, the light is off. We can manipulate the switch and reliably observe—for a while—that the state of the light obediently changes. If all we care about is operating the light, we are finished.

But some of us assume that there is some connection between the switch and the light. We can’t observe this connection without destroying the wall and ceiling, yet we have little doubt that there is one. This imagined connection has no basis in our empirical observation of *this* switch and *this* light, yet as a matter of faith we accept the existence of the connection. We do not accept effects at a distance, in most cases.

Does it now shake our faith if we flip the switch up and the light fails to come on? Not at all; we deduce that the switch has failed, the bulb has burned out, or there has been a power failure. We don’t observe those explanations empirically, either—although we would like to check them out by some indirect means, like turning on a floor-lamp to eliminate the general power-failure explanation. We don’t actually need to visit the power plant.

We begin to suspect the switch when we wiggle it up and down and observe that the light comes on again. But now we observe an odd thing: the light now comes on when the switch is *down* rather than up. We can still toggle the state of the light by moving the switch, but the relationship has reversed. Have the innards of the switch suddenly turned upside-down? That seems ridiculous.

Then we remember that high-school puzzle, the two-way switch. There is, we realize, *another switch* somewhere else that also controls this light. There is someone fiddling with that other switch!

Is this an empirical observation? No, it is a memory-based guess

about a hidden reality. Can we be certain that there is another person fiddling with the other switch? No. How could we be? We haven't seen the other person or the other switch yet, and if we have to catch a plane, we might never do so. It's possible that the spring in the other switch broke and let the switch flip down, with nobody operating it. It's possible that there's a relay in the circuit that short-circuited. Yet there is value even in our wrong guesses, because they are possible explanations and in other circumstances might be the correct ones. These possibilities relate not to the empirical world, the psychophysical flux, but to a world beyond what we are sensing: inside the wall, in another room we haven't visited, in a power plant we have never seen.

Most of the world within which an individual human being makes empirical observations is outside the scope of that person's perceptions, yet its imagined state forms the context within which what is observed is interpreted.

There are some aspects of the hidden reality that we strongly suspect to exist, but which we will never be able to verify. Is there really an electromagnetic flux propagating through empty space, a flux that we call "light"? There is no way to check this. We can only say that when some sort of receptor *stops* the imagined propagation, we get some sort of meter reading. There is simply no way to detect light in flight. Human reason screams at us that *of course* light has to propagate through space in order to reach our detectors—but that is not and never will be an empirical fact.

Nearly all of our meter readings interfere to some extent with what is being measured; meter readings lie. An ordinary volt-ohm meter draws current when it measures voltage. That current causes the actual voltage to drop a little—sometimes a lot. In electronics, we learn to measure the meter's resistance and calculate that of the circuit, and correct the reading to the "true" value. So the meter reading we see has to be corrected to indicate the voltage that really exists.

In order to estimate how hard a suitcase is pressing down on the rug, you have to lift it. Data from polls has to be corrected to show what the true opinions would have been if everyone had been telling the truth. When we bargain with another, we try to estimate from the offers the other is making how much that person is really willing to pay. When we see a car in the convex right-hand outside rear-view mirror, we see a • label saying, "Objects appear farther than they are." When a PhD candidate fumbles a question, we make allowances for her nervousness. When an agent says, "I'll give you a call if anything comes up," the actor ceases to expect a call. We are always making adjustments to observations, denying the validity of empirical data, to bring our actions and expectations more in line with a world that underlies appearances. We are better off doing so than not doing so, even though

we are sometimes mistaken in not taking literally what is before our eyes. Sometimes the phone rings and it is the agent telling the actor to report for rehearsals.

There are really no justifications for denying the existence of a reality that is different from the one we experience, even the one we experience through the use of scientific instruments. The scientific instruments themselves shout at us that there is something going on that is invisible to us. If we were strict empiricists, we would report analog meter readings in radians, not volts or pH or counts per second or RPM or pounds per square inch or quarks per cubic meter. But we do not: doing so would leave us with a world that made no sense. We have constructed an elaborate network of imagined entities and relationships that purport to live in the world on the other side of the meter readings. While this conceptual world might miss the mark and might describe only a projection of a much larger space onto the dimensions to which the meters are sensitive, it might be correct in some respects, particularly respects having to do with derivative notions like conservation of energy or control. We will never know, of course; our meters and our sensors stand between the observer and the reality. If information is coming in to us through these channels, we still don't know what it is about. The incoming information carries no identifying labels.

Still, it pays to guess, as long as we are alert for the evidence that says we should change the guess. But today's empiricism is tomorrow's illusion, and often today's.

Wayne Herschberger: Bill, you use the word reality very differently in the following two passages. The first usage refers to a created reality; the other refers to a hidden reality (my italics below), which you say is unverifiable. I find the former concept very useful, but not the latter. The latter is essentially a contradiction of terms, a paradox. This paradox is readily resolved, however, using your own very cogent arguments. That is, if the words "created reality" used in the first passage below are substituted for the words "hidden reality" in the second passage below, what you then say makes sense, don't you think?

Passage 1: "We replace the scattered, variable, inexplicable, individual observations with an idealization that we conceive of as the real observation. In doing this, we create precisely the reality I am talking about: a reality that we do *not* observe, but accept because it makes sense of experience. Empiricism itself leads to acceptance of a reality that underlies observation, but is not the same as what we observe."

Passage 2: "Most of the world within which an individual human being makes empirical observations is outside the scope of that person's perceptions, yet its imagined state forms the context within which what is observed is interpreted. There are some aspects of the *hidden*

reality that we strongly suspect to exist, but which we will never be able to verify. Is there really an electromagnetic flux propagating through empty space, a flux that we call 'light'?"

There is not just one created reality.

There are many created realities, some comprising perceptual "objects" (e.g., a laser beam), some comprising conceptual stuff (electromagnetic flux), both dealing with the same phenomena. When we attempt to order these created realities by their truth value, we find ourselves using parsimony and replicability as our criteria.


However, I am not sure that the conceptual stuff is any more or less empirical than the perceptual stuff. The mean of a set of data is a datum too. Further, although the arithmetic mean is not the only measure of central tendency which could be defined upon the raw data, it is very precisely constrained by its raw data.

It seems to me that our created realities transcend the raw psychophysical flux in essentially the same way that a mean transcends its data; that is, they are precisely constrained by the psychophysical flux—barring miscalculation. The input functions in your psychomodular control-system model create perceptual realities in this constrained way. That is why your theory is not solipsistic.

On another matter: "Empiricism fails as soon as you go beyond a description of a snapshot of the psychophysical flux." On the contrary, empiricism fails when you artificially restrict it to snapshots of the psychophysical flux.

Bill Powers: Wayne says: "There are many created realities, some comprising perceptual 'objects' (e.g., a laser beam), some comprising conceptual stuff (electromagnetic flux), both dealing with the same phenomena."

Is this a correct diagram of what the above says?



In other words,

"The same phenomena" are contained in
Perceptual stuff, which is contained in
perceptual created reality, and
Conceptual stuff, which is contained in
conceptual created reality.

Or, is it more like

Phenomena→

Perceptual stuff→

Conceptual stuff?

Wayne Hershberger: Neither alternative is quite correct, but the first is the closer of the two.

I would say that the natural order immanent in the psychophysical flux is realized both in the form of perceptual stuff and in the form of conceptual stuff. And when these two types of realization deal with the same phenomena, they often seem to depict a contradictory nature. As perceptual stuff, the desk at which I am sitting (an objective perception) is a static and solid object, but I can simultaneously regard it conceptually as a collection of whirling dervishes hurtling about within a confined space virtually as empty as an inflated balloon. The two types of stuff (i.e., the two types of reality) are, as I've said before, twin-born of the psychophysical flux (there is nothing like reason to make sense of something). This is not to say that objective perceptions (e.g., meter readings) are not useful in testing hypotheses derived from a theory representing a putative conceptual reality, but only that the perceptions involved in such tests should *never* directly involve the phenomenon being conceptualized. That is, you do not settle the geocentric vs. heliocentric world-view issue by watching a sunrise.

Further, I do not mean to say that there are only two levels or types of reality. For example, within the sphere I am calling perceptual stuff are to be found such things as real illusions (e.g., Ames' window, colored shadows, etc.), which are as different from the other perceptual stuff called real objects as the real perceptual objects differ from the conceptual stuff.

Furthermore, speaking of realities, as if there were a limited number of types, is misleading. Reality is but a dimension upon which we order the truth value of our countless epistemic creations—using replicability and parsimony as our criteria.

Bruce Nevin: An observation about the ongoing discussion of epistemology: I believe that in your language, Bill, you assume for yourself the perspective described by your (our) theory. This is an excellent way to test its adequacy, but perhaps deserves calling out for notice.

The theory or model requires there to be a "boss reality" in the environment to complete the feedback loop. Without it, perceptual control is impossible. So, from its perspective, there has to be a reality there.

However, the theory does not describe this reality or prescribe any attributes, other than that it be present and in at least some respects

stable so it is capable of being modeled within the perceptual control hierarchy.

In an important sense, this environmental reality is hidden from the perceptual hierarchy. Its only access to it is proximal stimulation of intensity sensors.

In an important sense, this environmental reality is not hidden from the perceptual hierarchy. Its model of it is presumed reasonably veridical because it in fact accomplishes perceptual control requiring feedback through the environment.

There are two senses of “model” possible here. The perceptual hierarchy may create a model of reality at higher levels of control. It is likely to use language to do this. Like any model, it is imperfect and requires periodic amendment. Because it is maintained at a high level of the control hierarchy, response to conflict is slow. The capacity for amendment is in the control hierarchy that holds the theory, not in the theory itself.

The second sense is that there is a model immanent in the perceptual control hierarchy as a whole. In its capacity to control perceptions, whatever the feedback through the environment might be, the perceptual control hierarchy is a kind of reflection of the environment. Like any model, it is imperfect and requires periodic amendment. Unlike many models, it includes this capacity for amendment in itself. Speed of response to conflict varies with the level of the conflict. Call this model 1 and the other sense of “theory” model 2.

I think it was the more primitive sense of the implicit, immanent model 1 that applied when I said “capable of being modeled within the perceptual control hierarchy.” Must “boss reality” be capable of being modeled in both senses? Does it have to be able to sit still for its picture to be taken, so to speak?

Bill, you are assuming the perceptual-control-theory model 2 (theory) as your perspective in talking about knowability. Anything not countenanced in that model you suspect is illusory. And you are using your model 1 to test the model 2, as indeed are we all.

We have a primitive sense that what appears to be there in the environment is real (naïve realism). We can talk ourselves out of this if our model 2 (theory) calls for it. We can also ignore perceptual signals if our expectations say they are not there—if associated error at higher levels is not significant. The two cases seem to me entirely alike.

Suppose there were a physical, mechanistic basis for our primitive sense of the reality of our perceptions. This doesn’t entail that this sense be articulate enough for us reliably to distinguish illusion and hallucination, something that we appear not to do. I would base this sense in the continuum of the environment outside the skin with the biochemical and biomechanical environment inside it. Awareness of this would

probably have an intuitive and emotional quality rather different from the attention to perceptual signals that usually concerns us. At its peak, an awareness of being part of a larger unity, perhaps.

But anything not countenanced in our model 2 or theory (the one we have adopted as our perspective) we suspect is illusory. For good reason: naive realism runs into well-known difficulties. And these intuitive apprehensions are global rather than particular. I think, too, that this sort of apprehension of the reality of reality is undemonstrable within the perceptual hierarchy. It is only apprehensible to it. Does that mean it is part of the environmental feedback for the control hierarchy? I don't know.

As we all know, "boss reality" doesn't really sit still for its picture to be taken. A model 2 or theory is possible only by categorization, subsumption with neglect, conventionalization, language. Only a model 1 with its continuous, live tracking can be veridical, and that only in a limited and local sense.

Theories are models 2 of perceptions in our model 1, which is a model or reflection of reality. All but the specific environmental feedback being tracked and controlled for is hidden from the model 1. Potentially, nothing is hidden from models 2, but their precision and accuracy are suspect. Partly this is because they are constructed using conventionalized social products such as verbalized categories. Partly it is because their responsiveness in the face of *aniccha*, impermanence, is too slow.

The obvious generalization is to speak of a level -1 model as most local and most accurate, a level -2 model, and so on, up to the models 2 of the system-concept level (some idealization here about the sequential separation of levels, as we know). Assuming of course that our model 2 of perceptual control is veridical.

As students, we take your verbalizations about error signals from your comparators as indications for setting reference values in our own. Which we might do, or we might verbalize error signals in turn. A reciprocal process called communication, of course.

This is a test. This is only a test.

Martin Taylor: Bruce, I would like to raise a flag to signal my objection to your claim that model 1 can be veridical "only in a limited and local sense." It is true that mathematical theories of physical "reality" take us a lot further than intuitive physics in predicting the behavior of the world, but it is not so clear that this is true for the less simple sciences. Physics is, after all, the only science so simple that the most intelligent humans have a reasonable hope of understanding some of it.

But we do behave reasonably successfully in the much more complex world of nutritious and poisonous foods, friends and enemies,


and so on, for which linguistically (e.g., mathematically) based models do a lousy job.

Bill Powers: Wayne: "I would say that the natural order immanent in the psychophysical flux is realized both in the form of perceptual stuff and in the form of conceptual stuff."

You have picked the first of my representations.


So does your model looks like this?

A:



Or would you draw it this way?


B:



(Psychophysical flux includes more than perception and conception together.)

Or this way?


C:



(Or with perceptual and conceptual interchanged.)


Or this way?

D:



Or this way?

E:



Or some other way?

Bruce: Isn't the concept of models 1 part of models 2?

What we know is the world of direct experience. If this is what Wayne means by the "psychophysical flux," then I would agree with him, but I would not use that term. To say either "psycho-" or "-physical" is to introduce description, characterization, and classification, which are added to what we observe and go beyond what we observe. The term "psychophysical" introduces a model and a theory right at the beginning of the discussion, which for me is too early. I want to explain the psychophysical flux, not assume it as a premise.

I do take the viewpoint of the hierarchical-control-theory model in explaining the structure of subjective experience. I take the viewpoint of the physics model in explaining the structure of the environment and brain. The purpose behind doing either one is to explain the way the world seems to me, as it is directly experienced. The ultimate criterion of truth for me is not any principle of philosophy, physics, neurology, or logic. It is simply whether these viewpoints, considered together, explain in an honest, testable, and self-consistent way what is puzzling to me about the world that I experience directly. I am not puzzled about its existence. I am puzzled about why it works as it does.

Why do I prefer honesty, testability, and self-consistency? Because

I like such things. Everything seems to work better when such principles are accepted as constraints. Highly recommended. Best buy. But not proven.

Model 1 is an assertion about the brain's built-in abilities to generate a perceptual world based on an external physical world. So it already contains an assertion about an independent physical world and a world of brain function. Model 2 extends model 1 to higher levels of brain function, by explicitly introducing modeling in terms of symbolic processes. But aren't these two models simply ways of classifying subsets in the general hierarchical-control-theory model?

Is either one more "real" than direct experience?

All of science, in my view, revolves around direct experience. We don't require models to explain each other: they are all required to explain what a human observer can experience, and how human actions and spontaneous changes in the observed world affect human experience. A theory or a model must bring order into the relationship between actions and perceptions, where "perceptions" includes both unaided human observation and observations of the readings generated by instruments (there's no fundamental difference).

To me, it is simply a fact that I don't experience anything but the surface of the world. I don't understand how anyone can claim that this is all there is. All you have to do is dig a hole, and you will see that the surface of the ground is held up by something else. That something else is hidden from the senses until you dig the hole. What holds *it* up is hidden until you deepen the hole. I see no hope of ever seeing what holds it all up, at the center of the Earth: long before we could get there, our shovels would melt. I'm willing to entertain the possibility that there is really a nickel-iron core in the center of the Earth. I do not, however, confuse that possibility with an actual experience of the Earth's core, the only incontrovertible verification of the possibility. Another theory might claim that there is a black hole at the center. The universe is not expanding: we are shrinking into the event horizon of the black hole. How would we verify that?

Hierarchical control theory is verifiable to the extent that it predicts classes of perceptions that we can actually experience and control. Control theory is verifiable to the extent that it predicts relationships among actions and perceptions that we can actually experience. In neither case, however, can we verify intermediate processes required to make the model work but which don't themselves correspond directly to aspects of direct experience. None of us, for example, can verify that these processes take place in a brain. That is conjecture. We will only know that these processes take place in a brain when we can link each process to a perception or measurement of activities and relationships in a brain and show beyond doubt that affecting each process as mea-

sured affects direct experience exactly as predicted. On the way to doing this, our conceptions of the intermediate processes in the model will undoubtedly change, and radically. The only things that must not change are the correspondences between variables in the model and aspects of direct experience. They provide the anchor points in reality.

In perceptual control theory and hierarchical control theory, certain identifiable aspects of direct experience are labeled "perceptions," and they correspond one-to-one with specific signals in the model (or they would if the model were complete). This does not change direct experience. It does change what we think about direct experience. We are led to think of all discriminable aspects of the experienced world as "perceptions," not just as givens. The perceptual signals in the brain model are linked theoretically through physical properties of neurons to other signals, and eventually to variables in the physical model of the world. The physical model deals primarily with variables and relationships that do not correspond to perceptual signals: a world beyond the senses. As predicted, we do *not* experience electrons, light waves or quanta, force fields or energy. The physical world becomes directly experientiable only at contact points established by meter readings of various sorts. What we experience is a meter reading, not the physical process that gives rise to it. Processes intermediate to those contact points and the physical variables on the other side of the meter remain conjectural and unverified. Therefore, the two models together imply that what we perceive is not necessarily in direct correspondence to the entities and relationships in the world proposed in the physics model. If we choose to use both models, the viewpoint we must take is that the world of experience is derived from or dependent on another world that is not experienced, just as the surface of the ground that we can see and touch is held up by deeper layers of unknown composition that remain invisible and intangible.

This is the only viewpoint I can see that is consistent with physical models, neurological or biological models, functional models of the brain, and direct experience. What we experience is not a model. Everything we say about experience is.

Wayne Hershberger: Bruce says, "As we all know, 'boss reality' doesn't really sit still for its picture to be taken."

What!? You boggle my mind.

Your sentence implies what it denies: that is, although we cannot picture it, "we all know... 'boss reality.'"

Perhaps you meant to say that, although we can picture it, we cannot know boss reality. But, of course, such a transcendental reality as that smacks more of heaven than earth.

The relationship between a hierarchical control mechanism and its

environment is a much more mundane affair than picturing a transcendent reality. You imply in the following two paragraphs two different avenues of access:

“In an important sense, this environmental reality is hidden from the perceptual hierarchy. Its only access to it is proximal stimulation of intensity sensors.”

“In an important sense, this environmental reality is not hidden from the perceptual hierarchy. Its model of it is presumed reasonably veridical because it in fact accomplishes perceptual control requiring feedback through the environment.”

I am not sympathetic with the first point. To say that the environment is hidden by all the proximal stimuli is to paraphrase the fellow who claimed not to be able to see the forest for all the trees. (E.g., “Gee officer, I couldn’t see the fireplug; my eyeballs got in the way.”) Also, don’t forget that the “intensity sensors” are spatially arrayed and sensitive to various forms of energy — over time. Further, transducers such as radar scopes vastly expand the range of our biological transducers. More trees to obscure our view?

However, I am favorably impressed with your second point, which is very similar to one I addressed last year—before you logged on to CSGnet. At that time, I observed that sensed efferece affords a significant window to the world; that is, when an environmental variable is being controlled, sensed efferece reflects the environmental disturbance (e.g., the weight of an object is proportional to the effort required to heft it). This principle provides a basis for the ideas of J. G. Taylor. Since that time, I have come across a delightfully lucid example from physics. Some physicists (Gerd K. Binnig & Heinrich Rohrer) won a share of a 1986 Nobel Prize by capitalizing on this principle in their design of the scanning tunneling microscope, STM.

The STM operates by passing an ultrafine tungsten needle over the surface of a sample to be studied. A low voltage is applied to the needle, creating a tiny electric potential between the tip of the needle and the atoms on the surface. Although the needle and the sample never touch in the classic sense, quantum fluctuations enable electrons to “tunnel” through the intervening distance, hence the microscope’s name.

The current passing between surface and tip depends on the distance between them. A feedback mechanism continuously repositions the needle as it scans over the surface to maintain a constant voltage: the undulations of the needle are studied to re-

construct the sample's contours. (Scientific American, June 1990, p. 26)

Bill Powers: Wayne, you comment to Bruce: 'Perhaps you meant to say that, although we can picture it, we can not know boss reality. But, of course, such a transcendental reality as that smacks more of heaven than earth.' Is that "of course" an argument against the proposition, or a bit of innuendo associating Bruce with a proposition that you *do* know how to refute? You must have a better reason than that for rejecting the possibility of a boss reality. Are you arguing against the uses of imagination?

Bruce Nevin: Wayne—sorry to butt in where angels fear, etc, I was supposing (out loud) why Bill might speak of aspects of reality being hidden. If Bill assumes the point of view of a perceptual-hierarchy model, and we assume a perspective supposedly outside of both that model and that which it is modeling, then we see that the only contact that a perceptual-hierarchy model has with "boss reality" is proximal stimulation of intensity sensors.

What might lie beyond that, accessible or potentially accessible (directly or in a further mediated, i.e., inferred, way) by way of proximal stimulation of intensity sensors might be reflected or imaged or modeled in the connections, input devices, and neural signals on up the hierarchy from those initial input devices and effectors.

Is the fidelity of that reflection or image or model verifiable? We postulate that coherent, successful behavior (however we define that) as an outcome of ongoing perceptual control constitutes a demonstration of fidelity. But the existence of conflict and reorganization must then be a demonstration of less than full fidelity. Since everything is connected to everything else, I suppose it might be argued that the "representation" immanent in the control hierarchy is complete—the universe in a grain of sand. But completeness in the same sense must be accorded the control hierarchy of a turkey.

All of which is only to say, there are grounds for assurance that the world of forces and impacts is there, but not for assurance that one knows everything going on in it. This is different from saying that some knowledge of it is in principle inaccessible. I know of no basis for either affirming or denying that.

As we all know, our pictures of "boss reality" are imperfect. (Our pictures: our snapshots, portraits, models, theories.) We know this by internal inconsistencies (conflict), and the very provision of means for revision (reorganization) in the model itself indicates that coevolutionary mutual adaptation is an aspect of that which we are modeling. A moving target indeed. I think this formulation is not ambiguous so as to

allow the pernicious interpretation entailing that we “know ‘boss reality,’” an interpretation that I did not intend in the original formulation. It relies only on our own perceptions, and on the assumption that these reflect reality, etc., as above.


I wonder if it would be useful to consider Bateson’s distinction between *pleroma* and *creatura*, the old Gnostic terminology by way of Jung, in place of the mind/body dichotomy that is the usual starting place for epistemology. From this perspective, the perceptual control hierarchy in a living control system is part of a continuum of cybernetic feedback loops extending throughout “boss reality.” The fact that a control hierarchy is more strictly organized than other parts of this cybernetic soup is an important distinction as regards the control (behavioral) aspects of perception, but does not bear so strongly on the receptive (observational) aspects of perception.

A sense-intensity receptor is a difference detector, as I understand it. A perceptual signal is then news of a difference as it enters the control hierarchy by way of a receptor from some other part of *pleroma*, and as it passes up the hierarchy being transformed into other differences that make other differences in turn. The combining of signals to make a signal of a different type is unique to control hierarchies, I suppose.

Wayne Hershberger: Bruce, a cybernetic perspective is certainly very appropriate. In fact, that is exactly my point. The environment is an Integral component of cybernetic systems, and Bill’s model is no exception. To speak of the environment as being outside Bill’s model makes no cybernetic sense to me. It is OK sometimes to linguistically “zero” the environmental part of Bill’s model (he certainly has greater proprietary claims on the internal hierarchical part) just so long as we don’t forget that the loops are closed through an environment. Cybernetically, the environment is part of the epistemic system.

Bruce, you ask, “Is the fidelity of that reflection or image or model verifiable?” If I understand what you are saying, “verification” cannot possibly entail a demonstration of any correspondence between the “model” and what you are calling “boss reality” (there is no one to bring the boss). So, it seems to me that boss reality is really a gold brick: a charming fellow who is nowhere to be found just when you need him.

Bill, if I were to draw Venn diagrams, I think I would want to label them as shown below, meaning that the natural order immanent in the psychophysical flux is realized both perceptually and conceptually. Further, there is more natural order in the psychophysical flux than is currently dreamt of in our philosophies, meaning only that the subset boundaries are not fixed.



Beyond this, I am reluctant to go, because it seems that I would then be doing what I claim we should not be doing: confusing control theory with cosmology.

However, I admit that the expression “psychophysical flux” does reflect my control-theory perspective. When I think of “the psyche,” I tend to think “reference values,” and when I think of “the physical,” I tend to think “disturbances.” Each of these is an input to the canonical control loop, giving the loop psychical and physical poles. These poles are as inseparable as the poles of a magnet, making the canonical loop (incorporating the two inputs) a psychophysical whole.

The canonical control loop may be partitioned into separate arcs by a mechanism-environment interface, but the location of this interface is an accident of nature and does not separate matter from mind. The loop itself is *not* psychophysical, in the sense of comprising a mental arc plus a material arc separated by receptors and effectors.

That is what I think—I think.

Bill Powers: Wayne, Bruce has reiterated the basis in the control-system model for entertaining the concept of a boss reality. Your response basically says that if it’s impossible to find the correspondence between the boss reality and perception, why bother with the concept?

You open that comment with: “The environment is an integral component of cybernetic systems, and Bill’s model is no exception. To speak of the environment as being outside Bill’s model makes no cybernetic sense to me.”

Your comment and Bruce’s finally, maybe, perhaps, have joggled me into the right point of view for explaining my recalcitrance and possibly bringing our mysterious controversy to an end.

Yes, in my model there is always an environment and a behaving system. Neither makes sense without the other. I have always taken both into account. So follow me as I outline a chain of reasoning, and see if there is any point where you detect a weak link.

We’re being modelers now. Imagine a sheet of paper on which we

draw two boxes, an Environment on the left and an Organism on the right. We don't need to model the environment; that has already been done better than we could do by physics, chemistry, and if you want to include raw meat, anatomy and neuroanatomy. We can put physical variables into that Environment together with all the laws that express relationships among them.

What we're trying to model is the organism part. So we draw two arrows: one from the environment to the organism, representing effects the environment has on the sensors of the organism, and one representing effects the output devices of the organism have on the environment. We are sitting up here with a good view of the paper, so we can see what is in the environment and what we're putting into the organism.

The challenge is to build a model of the organism so it will interact with the environment exactly as the real organism does. This means that basically we can give the model no help other than to provide it with the functions and interconnections that will, by their operation, generate some sort of behavior. When we guess wrong, we find that the functions and interconnections do *something*, but it bears no resemblance to real behavior. We just keep fiddling with the model until it behaves correctly. This leads us to a hierarchy of control systems, and so on.

If this model is to be complete, however, it has to reproduce not just behavior, but experience. In other words, the physical environment over on the left has to appear to this model just as it does to us. If we see intensities, the model has to see intensities. Simple receptors excited by various forms of physical energy will do for that. If we distinguish sensations in which different intensities are interchangeable, the model must do so. No problem: weighted sums seem to make sensation perceptions depend on physical variables as they should.

As we go higher, the problems become tougher, but we know what we're working toward. We want the model to contain signals representing configurations, transitions, events, and the rest, because we can see the world in such terms. We can't just tell the model about such things, of course; it has to contain the equipment that will, all by itself, derive such perceptions from its inputs. At the moment, we're pretty far from being able to do that, but we can at least draw boxes into the model showing where we will put the machinery for deriving the signals once we know what it is. As we know what the signals have to correspond to in our own experience, we can label them: "event perception," "relationship perception," "category," etc., corresponding to our subjective analyses of private experience. The model has to have those same private experiences. It has to have *all* the private experiences that we can discriminate into "natural kinds." That includes

thought and reasoning.

If we now want to go far beyond where we are in the process of building this model, we may want to ask about epistemology. From our perch above this sheet of paper, we can see both the physical variables in the environment and the perceptual signals inside the organism model, the model of the person. It's perfectly clear that the perceptual signals are derived in systematic ways from energy fluxes connecting the physical variables to the sensors. As we fill in the boxes, we come to understand the details of that correspondence: just how an object in the environment, through the properties of light and optical devices, and through the photoneural receptors, comes to give rise to signals indicating its size, its distance, its shape, its orientation, and so on.

But now we come to the crux of the problem. We want to let the model figure out what there is external to it that corresponds to its perceptual signals. For example, the object it is looking at is actually a hologram, and all that actually exists in the environment is a set of wavefronts of light that don't actually originate at the surface of an object. How does the model go about checking into the reality of the object? We have no problem; we can see exactly what is going on. But how can the model figure it out without us to whisper in its ear? The model doesn't necessarily understand holograms (this has to be a model of any person).

One way is for the model to extend a limb to bring its visual image into the same region of visual space as the apparent object. If no contact is felt, the object could be considered intangible (that being what intangible means). But is it an intangible object in that position, or is there no object at all? Is this some kind of plasma object, or a less familiar trick of nature?

Solving this problem would clearly require a lot of sophistication and experience on the part of the model. It would have to compare what one set of sensors reports with what another set reports. It would have to form hypotheses and test them by performing appropriate acts. In the end, it would probably narrow the possibilities down to a small set, and on the basis of preference or niceness or some general principle, pick one of them as the answer.

Would it pick the same answer we would give from our omniscient point of view? Possibly, possibly not. In truth the model would have to know everything we know about the environment, and interpret its information exactly as we interpret it, and know what operations take place inside its own perceptual functions (which are not represented in the signals) to arrive at exactly the correct conclusion about what corresponds to any of its perceptual signals.

There is one thing we can be certain that this model can't do. It can't rise out of the plane of the paper and peer across at the environment

model to see what is going on there. We have given it no abilities that would allow it to see the environment except through the raw sensitivity to energy at its input sensors. The line separating it from the environment is a barrier that can be crossed only at the most primitive level, by physical energy.

So, for this model, as we have constructed it, *we* can know for certain how its perceptual signals correspond to what is happening in the environment model, but it can't know for certain. All it can do is entertain possibilities. One of those possibilities might be absolutely correct. But it can't know which one, if any.

So that is the epistemology of the model. Now what about our own?

If this is indeed a model of a human being, if we've got everything right, then it is a model of the observer, of ourselves. It is a model of us sitting up here and looking down at a sheet of paper on which there are diagrams of an environment and of a nervous system. The model has eyes and limbs; they are models of our eyes and limbs. The model has sensors and neural signals which are supposed to represent our own sensors and neural signals. The model, if it were looking at a sheet of paper with diagrams on it, would know of those things only in the form of neural signals inside itself. As the model can't rise out of the plane of the paper to see what is really in the other diagram, the diagram of the environment, so we can't rise in a fourth dimension out of our brains, to peer at whatever it is that is causing our neural signals. As the model can't sense the internal workings of its perceptual functions and use that information to deduce what is causing any given perception, so we can't deduce the transformations that lie between the environment and our perceptions.

The model might conclude correctly that it doesn't have access to an authoritative picture of the environment model; it could reach this conclusion simply by noticing that several plausible alternative interpretations exist. On that basis, it might decide that there is no point in guessing about a boss diagram that it now realizes it can never experience directly. It might decide that all it can do is compare one perception with another, and take that as the beginning and end of reality. The boss diagram is an unnecessary frill, a religious superstition; it is to laugh.

Of course we, sitting up here, would laugh at that, knowing what a mistake it is. There really is a diagram of the environment there, and it really does have a particular state, and the model hasn't been so far off the track as to be completely hopeless. At least it could survive in its interactions with the environment on the basis of what it has deduced. What it thinks it is controlling is at least equivalent, in the necessary ways, to what it is actually controlling. It might have omitted a conformal transformation or two here or there, but because it omits the same

transformations from perception of its own actions, the two mistakes cancel for all practical purposes. And if it gives up now, assuming that all there is to be known exists already in the perceptual world it has constructed for itself, it's going to miss most of the fun.

And what of us? We sit up here, experiencing our own perceptions, and debating whether or not they are connected to a physical world, and if so, what kind of physical world. If we believe what the model of the person seems to imply, then we are in the same fix it is in: we experience our perceptual signals, but there is nobody sitting in a higher place still who can tell us what the environment diagram really looks like. We have to figure it out on our own, each in an individual private world.

So that's where my epistemology comes from. It comes from trying to think of a model that behaves and experiences like a person, and is built the way a person is built with sensors and a nervous system and effectors. The final step, to my personal epistemology, is simply an application of the model to myself. The model contains my best understanding of how the nervous system on the right, and the environment on the left, work and interact with each other. If I now don this model and imagine that I am experiencing the world from inside it, I transform my understanding of the physical world that seems to surround me. I realize that a very plausible thing to say about it would be: it's all perception.

But it is not implausible to add "... of something else."

Wayne Hershberger: Bill, a thought has occurred to me that you might agree would prove helpful. It seems to me that the little man model might be said to know where the target is, while knowing nothing either of computers or of yourself, his creator. That is, it seems to me that if you were to rewrite your essay, putting the model in a computer where it can function as a simulation, the epistemological implications might appear more clear-cut.


Bill Powers: Wayne says, "... if you were to rewrite your essay, putting the model in a computer where it can function as a simulation, the epistemological implications might appear more clear-cut." Actually, that is the route I took to my present position, only it came from building real systems more than from simulations. In the late 1950s, for example, Bob Clark and I built an "isodose tracer" that used an analogue computer as a control system to make a tiny radiation probe move along curves of constant radiation intensity in the beam of a Cobalt-60 treatment machine (VA Research Hospital in Chicago). In the early stages, we got some strange curves, because the long stem that held the probe turned out to be radiation-sensitive. The control system was

keeping what it assumed to be sensed radiation at the probe tip constant, but it couldn't know where that radiation was being detected. The variable under control wasn't quite the one that was supposed to be under control.

I've mentioned the voltmeter effect before: the reading on the voltmeter is not the "true" voltage because the meter draws current. In my electronics ventures with radiation probes and photosensitive equipment, often incorporated into control systems, it was almost always necessary to correct the meter readings when measuring low-current high-voltage sources. Automatic control of such voltages required compensation so that the "real" voltage, not the measured voltage, was controlled.

The whole world of electronics is fraught with examples. A simple circuit board is, to the electroniker, largely imaginary. The surface appearance of the board has almost nothing to do with what is "really" going on. Every component carries in it mysterious properties like resistance, capacitance, inductance, and amplification that are never experienced directly. (Voltage is one example, but it doesn't feel like voltage. It feels like hell.) Usually, such things are known only after calculations based on the few contact points with direct experience. Yet when you assume that such things exist in some boss reality, as you must in order to make any sense of "correcting a meter reading," the result is the power to make things happen in highly predictable ways. You adjust a tuned circuit a little below resonance, so it will be exactly at resonance when you remove the capacitance of the probe you're using to measure the response. The true operation of a circuit is what you deduce would take place if you weren't measuring anything!

You've picked this diagram from the possibilities I suggested:



"Immanent order" wouldn't be a bad term for "boss reality." From my viewpoint, it has the nice implication that there can be order without our knowing what it is. But I have some more questions.

By choosing the diagram in which the immanent order extends be-

yond the boundaries of the realizations, you have agreed with me that there is more to know than meets the eye. By making the two realizations independent and non-overlapping, you have said that each has its own relationship to the immanent order independently of the other.

In this diagram, there's no connection shown between perception and conception, nor any indication of how these "realizations" might relate differently to the immanent order. You describe the figure as a Venn diagram. This implies that within the outer boundary there is some immanent order, and that it's simply marked off into regions, with the elements of the largest field being no different inside and outside the two "realizations."

As you didn't specify the difference inside and outside the realizations, there are two possibilities:

A realization is simply a noticing of something that was always there, the noticing in no way altering what was always there but merely bringing it into the field of attention.

A realization is some transformation or projection of the immanent order, so that the realization is an invention or at least an expression of the nature of the system becoming acquainted with the immanent order.

In both cases, there is an implicit relationship between a realization and the immanent order. In the first case, the realization is completely passive; it is merely recognition. In the second case, there is a difference between the realized and unrealized states of portions of the immanent order. Does either of these choices fit your conception?

I take it that the rationale for the term "immanent order" is that neither perception nor conception is random; that both reflect some orderliness that constrains them. Does this not imply some effect of the immanent order on the realizations?

Wayne Hershberger: Bill, you said: "Yes, in my model there is always an environment and a behaving system. Neither makes sense without the other. I have always taken both into account. So follow me as I outline a chain of reasoning, and see if there is any point where you detect a weak link."

That is like waving a red flag in front of a bull. How can I refuse? Here goes.

Although you said, as quoted above, that the environment is in your model, you soon spoke as if it were external to it: "But now we come to the crux of the problem. We want to let the model figure out what there is external to it that corresponds to its perceptual signals. For example, the object it is looking at is actually a hologram, and all that actually exists in the environment is a set of wavefronts of light that don't actually originate at the surface of an object."

The basic (perhaps only) problem I see is your confusing separate issues. In your elaborate analogy, there is one dichotomy (creator-created) and one dyad (organism-environment). You confuse the two, shifting from one to the other as though they were one and the same. As I have said before, this confuses physiology with metaphysics.

For example, you say that "... we can't rise in a fourth dimension out of our brains, to peer at whatever it is that is causing our neural signals." The first part of this remark, mentioning a fourth dimension, is alluding to the inability of the created dyad to assume the epistemic perspective of the creator (epistemology), whereas the last part is concerned with the relationship between the two parts of the created dyad (sensory physiology). Apples and Oranges.

Then you reverse course and switch from sensory physiology back to epistemology with the following remark, effectively by substituting the word "perceptions" for the words "neural signals" (moral: perceptions are not to be equated with perceptual signals).

"As the model can't sense the internal workings of its perceptual functions and use that information to deduce what is causing any given perception, so we can't deduce the transformations that lie between the environment and our perceptions." If a neural model could monitor its role in the perceptual process, could it deduce the nature of the transformations that lie between the neural model's signals and the neural model's environment?

You also say: "'Immanent order' wouldn't be a bad term for 'boss reality.'" From my viewpoint, it has the nice implication that there can be order without our knowing what it is." Yes, exactly. It seems to me that the expression "immanent order" (or natural order, or what have you) would be a much better term for your purposes than "boss reality," for precisely the reason you mention. The word reality connotes a verifiability you are denying to "boss reality," making the expression an oxymoron.

And: "By choosing the diagram in which the immanent order extends beyond the boundaries of the realizations, you have agreed with me that there is more to know than meets the eye." No. I think there is a point of agreement here, but not for the reasons you say. The aspect of the diagram that implies that there are some things which can be known but which do not directly meet the eye or the ear or the other sense organs (e.g., your example of voltage) are the elements which are both *in* the set labeled conceptual realizations and *not in* the set labeled perceptual realizations. In contrast, the elements that are in neither subset (neither type of realization) simply imply an immanent order which is not realized—either perceptually or conceptually. Whether this unrealized order is potentially realizable is something a static Venn diagram doesn't capture. But if one takes the view that at

least some of the immanent order unrealized at present might be realized in the future, it is presumptuous to suppose that this realization *cannot* be perceptual. Further, any immanent order which cannot possibly be realized at any time in either way is simply not to be known; it does not mean that there is more to know than can be known. I readily admit that there can be more than what-can-be-known, but I cannot agree that there is more to be known than what can be known, without contradicting myself. Nor can you. We are talking here about the limits of the epistemic process, not the limits of a man—obviously, there is more to be known than any one man will ever know.

“By making the two realizations independent and non-overlapping, you have said that each has its own relationship to the immanent order independently of the other.” I would say that one is not a subset of the other, but their intersection is not nil, meaning that the two realizations are independent of each other. Your drawings did not seem to include this alternative, so I selected the one which I thought would “suggest” independence (actually non-overlapping subsets depict mutual exclusion, a form of dependence).

“As you didn’t specify the difference inside and outside the realizations, there are two possibilities:

1. A realization is simply a noticing of something that was always there, the noticing in no way altering what was always there but merely bringing it into the field of attention.

2. A realization is some transformation or projection of the immanent order, so that the realization is an invention or at least an expression of the nature of the system becoming acquainted with the immanent order.”

Your two alternatives are not a matched set. They are not necessarily mutually exclusive.

First the latter: The expression “system becoming acquainted with the immanent order” seems to me to suggest that the “system” transcends (stands apart from) the immanent order. That would insinuate a gratuitous wild card. For me, the system becoming acquainted with the immanent order must be part and parcel of the immanent order. Perhaps they are even coextensive. The system responsible for the two types of realizations is best characterized as an ecological system (i.e., an organism-environment dipole). If we attribute the “becoming acquainted” merely with the organism pole and the immanent order merely with the environment pole, we are being arbitrarily inconsistent. Therefore, if one is to be consistent, it seems to me that the realizations would inevitably be “an expression of the nature of the system becoming acquainted with the immanent order,” because the system (the ecological dipole) becoming acquainted with the immanent order is part and parcel of the immanent order.

Now the former: Does “self-acquaintance” rule out the possibility that acquaintance is simply a registration of what is “there”? Fortunately, the question appears to be academic. If some aspects of the immanent order are hidden by the recursiveness of self-acquaintance, or whatever, I would say, so what? Call it *Noumenon*, and let the faithful worry about it, because, by definition, it is not to be known.

Finally: “I take it that the rationale for the term ‘immanent order’ is that neither perception nor conception is random; that both reflect some orderliness that constrains them. Does this not imply some effect of the immanent order on the realizations?” The relationship is not cause-effect, but *yes*, realizations of both types reflect some orderliness that constrains them.

Bill Powers: Wayne, we keep going around and around on the same points without getting anywhere. You keep saying that I am missing the distinction between modeling and metaphysics, and I keep saying that metaphysics is just one of the things a brain can do. Let’s take it from the top.

You say, “It seems to me that the expression ‘immanent order’ (or natural order, or what have you) would be a much better term for your purposes than ‘boss reality,’ for precisely the reason you mention. The word reality connotes a verifiability you are denying to ‘boss reality,’ making the expression an oxymoron.” So in your book, “reality” is identical with “verifiable reality.” It’s not, in mine. I don’t need to understand electricity to comprehend that touching certain objects is highly unpleasant. I can generate acts like touching objects, but I can’t decide what their consequences will be. That is decided for me by something I don’t sense and only partially conceptualize. I can choose whether to repeat a consequence or to avoid it, but I can’t make an act have a different consequence. In that department, something else is boss.

Are you saying that I must realize in perception or conception the connection between an act and its consequence, and *verify* the nature of that connection, before I can accept that there really is a connection? Or are you saying that it is sufficient to verify only that the consequence reliably follows the act, and never mind why? I would argue against the latter as being simply pre-Galilean empiricism, and reject it because it works so poorly in comparison to the method of modeling. The method of modeling posits an unseen reality mediating between act and consequence, and has most profitably interpreted nature in those terms. The assumption has repeatedly been vindicated. How could the purely empirical approach ever predict a new perception, and experimentally reveal the link explaining the surface appearance of a causal sequence?

Later in your post, you say, "... if one is to be consistent, it seems to me that the realizations would inevitably be 'an expression of the nature of the system becoming acquainted with the immanent order,' because the system (the ecological dipole) becoming acquainted with the immanent order *is* part and parcel of the immanent order." This would be consistent. It would also be an empty generalization, a true statement of which one can legitimately ask, "so what?" To say that all of knowledge is an expression of the immanent order (whatever that is) is meaningless: any statement that is true of everything is trivial. Even that statement and my response to it are part of the immanent order. I repeat: so what? Knowing that does not contribute to our understanding of any specific phenomenon—in fact, it seems to discourage asking questions and conjecturing. All of our useful understanding comes from discriminating one part of the immanent order from other parts, and from realizing that different parts of it have characteristics of their own unlike the characteristics of other parts. It is out of these differentiations that all knowledge comes. From these differentiations, we come to realize that organisms and environments are *not* alike. We realize that some parts of organisms function differently from other parts. We realize that brains exist.

And ultimately we are faced with a paradox, the one you and I have been arguing about. We find by experimentation that the presence of certain signals in a brain is the *sine qua non* of perception. Remove those signals and you destroy, as far as the victim is concerned, a chunk of the immanent order. Yet you don't destroy it for anyone else. What other conclusion can we reach but that perception is absolutely contingent on those signals? That puts us, as perceiving entities, inside the brain. To deny that would be to destroy the whole structure of perceptual and conceptual organization we have so painfully built up. That structure is at least as well worked out as any metaphysical argument in words, and a lot better tested experimentally.

I don't see that any philosophical conception, any combination of words, any exercise of pure reason, can be more persuasive than these simple observations. By simple and straightforward reasoning based on close attention to experiment and observation, we are led to conclude that the object of perception and thought is a world existing inside, not outside, a brain. We can see how this world of experience is related to what we conjecture to exist in a physical environment outside of us, but we can also see that the relationship is not a simple or direct one, nor is it wholly verifiable because of our peculiar circumstance of being inside the very system we model and by necessity having to perceive and think using its equipment.

Until you can come up with an equally persuasive set of observations and deductions that lead to a different conclusion, I will continue

to be satisfied with my view of the relationship between consciousness and reality. Simply reiterating your point of view without revealing and justifying each step of the way that leads to it will not win me over. I understand that if I believed as you do, all would be explained. But I do not.

Rick Marken: I have gotten behind on this epistemology debate (or maybe I just don't understand it). Could Bill or Wayne give me a short (like two-sentence) description of what is being debated? I am wondering if Wayne is arguing that there is no physical environment, or that the physical environment is an unwarranted assumption, or what?

Wayne Hershberger: Bill, we are back where we began. You say: "And ultimately we are faced with a paradox, the one you and I have been arguing about. We find by experimentation that the presence of certain signals in a brain is the sine qua non of perception. Remove those signals and you destroy, as far as the victim is concerned, a chunk of the immanent order. Yet you don't destroy it for anyone else. What other conclusion can we reach but that perception is absolutely contingent on those signals? That puts us, as perceiving entities, inside the brain. To deny that would be to destroy the whole structure of perceptual and conceptual organization we have so painfully built up. That structure is at least as well worked out as any metaphysical argument in words, and a lot better tested experimentally." And I say that ablation (to which you refer above: "Remove those signals..."), the technique pioneered in the 17th century to localize mental functions, identifies certain *necessary* components of the various functions we call mental (e.g., vision). It does *not* identify the *necessary and sufficient* components. Without the photon, there is no vision—for anyone. Ablate photons and we are all blind. This means that the *proprietary* aspect of our respective experience (my perceptions versus your perceptions) are contingent upon our respective brains. That is the argument you are making, right? But that does *not* put us in our respective brains! Our feet are too big.

So, you are right, this is where we came in. Perhaps it is time to take a different tack. Let me reciprocate by asking you what, if anything, is wrong with the following remark (using your terminology): Bill Powers' hierarchical-control-theory model models an aspect of boss reality (the organism aspect), the other aspect (the environment aspect) already having been well-modeled by contemporary physics.

Bill Powers: Wayne, what are these mythical "photons" of which you speak? I don't know anyone who has the ability to "ablate photons." We can perform various acts, like shutting our eyes or pulling

the chain on a light, that result in loss of vision, but to attribute that loss of vision to a loss of “photons” goes far beyond anything that is observable.

By accepting “photons” as necessary precursors to vision, you are leaping ahead to the conclusion you want to reach; namely, that photons actually exist just as we imagine them to exist. But I can’t accept that mode of argument: I want to know the operational basis for every critical entity you use in your proofs. I will accept that turning off the lights results in loss of vision. Those are both observables, perceptions. I do not accept that you have shown photons either to exist or to have anything to do with this phenomenon—not until you tell me your basis for knowing that

“This means that the *proprietary* aspect of our respective experience (my perceptions versus your perceptions) are contingent upon our respective brains. That is the argument you are making, right? But that does *not* put us *in* our respective brains! Our feet are too big.” Cute comment, but irrelevant. Our perceptions that we call “feet” are certainly not too big to fit into a brain: they are precisely small enough to pass through a neural fiber. All aspects of our perceptions are proprietary, including our convictions that some are not. If that were not true you would have convinced me by now. But you have nothing objective to show me to help make your case.

You ask, “... what, if anything, is wrong with the following remark (using your terminology): Bill Powers’ hierarchical-control-theory model models an aspect of boss reality (the organism aspect), the other aspect (the environment aspect) already having been well-modeled by contemporary physics.” Sensing a bear-trap, I answer cautiously. Both the hierarchical-control-theory model and the physics model purport to represent aspects of a boss reality. Both are tested (by a person, using a brain and body) by assuming the model to be correct, and predicting the effects of actions on this boss reality that have consequences we can perceive. It is the boss reality that determines whether our predictions work out as we expect, or whether different consequences occur. If the consequences are different, we modify our imagined pictures of the boss reality in a direction that promises to lessen the difference. This process converges to some minimum-error condition where we declare ourselves satisfied with the models. According to both the physics model and the hierarchical-control-theory model, this process of acting and testing takes place inside a brain. It is not necessary to assume that the model in the brain has any particular correspondence to the boss reality. It is necessary only to assume that whatever that correspondence might be, it is stable over time.

I ask a similar question of you: is it fair to say that you believe (a) that there are non-proprietary aspects of our respective experiences and

(b) that we can say unequivocally what they are?

Avery Andrews: The first wisdom of linguistics is that speakers are always wrong when they try to explain why they say what when (the stories are pathetic, and tend to fail within 30 seconds). Getting behind the descriptions to the explanations will be reverse engineering all the way (I regard current linguistics as being essentially descriptive, in spite of the presence of a lot of talk about explanation).

Bill Powers: Avery says that "... getting behind the descriptions to the explanations will be reverse engineering all the way..." Beautiful. Precisely.

Gary Cziko: I would appreciate Avery and/or Bill giving me a description and example of "reverse engineering." I have a hunch that *all* science and *all* nontrivial engineering (i.e., finding engineering solutions to new problems) is in fact reverse engineering, but I want to know more about what this term means before making this claim.

Avery Andrews: Gary: All I meant by "reverse engineering" is that there is no quick substitute for figuring out how it works on the basis of analyzing what it does.

Bill Powers: Gary, reverse engineering is a term from (I believe) the semiconductor industry. It refers to duplicating the function of someone else's integrated circuit. What with copyrights and patent laws, modern reverse engineering gets pretty complex. One team analyzes the function of the competitor's chip and prepares a specification stating the relationships between inputs and outputs (and other aspects of visible behavior) that the "unknown" chip creates. This specification is then passed on to a design team which is never given access to the chip itself, only to the specification. The design team is never allowed to communicate directly with the analysis team. From the specification alone, the design team generates a completely new chip design, from scratch, that will accomplish exactly the specified functions. I'm sure there has to be some cheating—the design team has to know that the specs describe a computer, for example, and not a sewing machine.

At any rate, the result is a new chip that can be plugged into the same socket that the original chip occupies and works exactly the same way, down to the last detail of functioning. This is the ultimate in the method of modeling.

In fact, the final chip might not accomplish the functions in exactly the same way the original did. Sometimes the new chip proves to perform some functions more efficiently than the original—in fewer steps,

or faster. Presumably, if those aspects of functioning had been part of the spec, the design team could have deliberately slowed some circuit operations and matched the slowness of the original too! But the design team, prior to releasing its product, never can know whether it has accomplished the functions in the same detailed way that the original does. In the final comparison, it is often found that some functions were reinvented exactly as in the original, while others do the same things in a different way. That is what is hoped for—what avoids a suit for patent infringement.

This is basically what I am arguing with Wayne Hershberger about. We are trying to reverse-engineer evolution (or whomever you want to blame). In doing so, we come up with a model of underlying design features constituting a system that interacts with its environment just as real organisms do. Of course, in doing this, we try to reproduce only those functions we understand, and we ignore many others, such as skin color, weight, exact lengths of appendages, and so on through a long list of “unimportant” parameters. As initial models succeed, we bring in more detailed parameters to match, even to the level of neural functions in a few cases.

But we can never know that we have accomplished something in the same way that an organism accomplishes it, in every detail. For that matter, we have no reason to think that every organism of a given species accomplishes its functions in the same way as other organisms of the same species. Judging from the very large differences in brain anatomy that exist from one person to another, in fact, it's unlikely that all people are internally organized in the same way even if they behave in roughly the same way. The brain is plastic and its organization is influenced by the experiences of a single lifetime. Our reverse engineering is fundamentally limited by this fact: no one model can ever reproduce to the last detail the inner functioning of all examples of any kind of higher organism, because the originals are not all designed in exactly the same way. We will always be limited to modeling the “general idea” behind an organism, because that is the limit of consistency in the originals. The method of modeling is primarily a method of understanding individuals, and only secondarily a way of saying general things about all individuals. Models must always contain parameters that can be adjusted to fit the “general idea” to a specific organism.

This, naturally, has some serious implications concerning the nature of scientific research into human nature. It's usually assumed that one is dealing with a standard instance of *Homo sapiens*—the very idea of assigning such a term to the whole human race is to assert that fundamentally we are all the same. In the psychology lab, great attention has been paid to using a standard animal model—the Sprague-Dawley rat, during my formative years. If you have a standard rat or a standard

person, you should get standard responses to standard stimuli. If any human being is as good an example of *Homo sapiens* as any other, you can study groups of people as interchangeable units, drawing generalizations from the data which you assume to be measures of common underlying properties fuzzed out by uncontrolled stimuli.

But what if, below some level of observation, there are no common underlying properties? Then the whole rationale of statistical studies of populations collapses. The specification team can't come up with a spec that fits all instances of the chip that is to be reverse-engineered. All they can describe, for each parameter, is the average spec. As Russell Ackoff said in a lecture that Dag Forssell has transcribed, there's no way to design the optimum human being by combining the optimum spec for each function making up the person. This would be like trying to build a perfect car by using the engine of a Rolls-Royce, the suspension of a Ferrari, the body of a Chevette, the carburetor of a Chevrolet, and so on. The functions all have to work together in a single person; the final workable form of each function depends on the final forms of all the other functions. Each part of a person is adapted to all the other parts of the same person, not to the same parts as they are manifested in other individuals. And the process of mutual interadaptation never ceases.

I use the term "generative model" as Humberto Maturana defined it (perhaps following someone else). A generative model is one that will reproduce the phenomenon of interest by operating strictly from the interplay of its own properties. A generative model of control behavior is a control system with an input function, a comparator, and an output function, in an environment that links output to input in a specific way. There is no component in a control-system model that "controls." Control is the result of operation of a system with these functions in it, connected as specified by the control-system model, and operating as dictated by the input-output properties of each component.

So, given inputs, constraints, and parameters, a generative model must always produce some kind of behavior. We can't necessarily anticipate what such a model will do, but whatever it does is rigidly set by the properties we have given it, and by the surroundings with which it interacts. We hope that the behavior of the model will resemble the phenomenon we're trying to explain. If it doesn't (and few models do, the first time they are set in motion), we have to modify the model. That's how models grow and improve.

Wayne Herschberger: Bill, you said, "Our perceptions that we call 'feet' are certainly not too big to fit into a brain: they are precisely small enough to pass through a neural fiber." Neural signals in the brain might be said to be relatively small, but the replicable perceptions

(phenomenal objects) those signals help mediate are not necessarily small. Smallness is an aspect of phenomena, and it is a mistake to suppose that the size of a phenomenal object is in any way related to the size of the neural signals which help mediate it. You have yourself been championing this sort of argument in many of your recent posts.

You also said, "All aspects of our perceptions are proprietary, including our convictions that some are not." No. A proprietary aspect is immanent in all experience, or so it seems. But this does not imply that there are no other aspects.

"This is basically what I am arguing with Wayne Hershberger about. We are trying to reverse-engineer evolution (or whomever you want to blame). In doing so, we come up with a model of underlying design features constituting a system that interacts with its environment just as real organisms do." Yes. As I see it, we are trying to reverse-engineer the phenomenal domain, and the "spec" that I think is of the first importance in this venture (also, as I think Kant was saying) is that phenomena are bipolar: in a word, psychophysical. Control theory appears to be uniquely compatible with this psychophysical specification, providing one continually recognizes both ends of the dipole—a control system *and* its environment. Perhaps we should change our language habits and speak of control subsystem, since the control system is only one part (or pole) of the system being captured by our reverse engineering.

Bill Powers: Wayne says, "Neural signals in the brain might be said to be relatively small, but the replicable perceptions (phenomenal objects) those signals help mediate are not necessarily small. Smallness is an aspect of phenomena, and it is a mistake to suppose that the size of a phenomenal object is in any way related to the size of the neural signals which help mediate it." But you assume, in order to say this, that phenomenal objects and attributes of objects are something other than neural signals. I assume they are the same thing. How do we get past that?

"As I see it, we are trying to reverse-engineer the phenomenal domain, and the 'spec' that I think is of the first importance in this venture (also, as I think Kant was saying) is that phenomena are bipolar: in a word, psychophysical." Why do you assume the "-physical" part of psychophysical? There is nothing in the physical domain that is not derived from perception and thoughts about perceptions. It seems to me that you slip your conclusion into your premises. I do not see the "psychological" aspect of experience as being on an equal footing with the "physical" part. The physical part is a set of ideas, and so is a subset of the psychological part.

I find the topology of your point of view baffling. It seems to involve

some magical way of knowing things without perceiving them, and some way of checking on the meanings of perceptions other than comparing them with other perceptions. I can't grasp the role that you give to perceptual signals, or for that matter, to the brain. I can't understand what position you're assigning to the Observer—if the observer isn't in the brain, where is it? And where, then, are the objects of observation?

Joel Judd: I'm starting to lose track of what is being claimed as individual responsibility and what's being foisted on the environment. Don't we all agree that the "world" is constituted in our perceptions?

Martin Taylor: Joel, I can't speak for Bill, but in his discussions with Wayne, I think I agree with him. I assume that there exists something outside ourselves, but it can be known only through our perceptions. Our perceptions can be constructed only through the feedback of our actions to our sensors, but we can develop internal things (which I call structures to avoid words like "simulated worlds" or "world models" or "imagined worlds") that enable us to perform as if there were certain objects and relationships in the (unknowable) world and not get into too much trouble by doing so.

Bill Powers: Mental representations, in hierarchical control theory, are identically neural signals arising from sensory receptors. Each level of signals enters a higher level of perceptual functions (neural computers), many functions acting in parallel, which re-represent subsets of the incoming signals as a new level of mental representations. There are 11 such levels in my model, covering (as far as I could) all phenomena of perception, all aspects of the experienced world, inner and outer, concrete and abstract. I refer to the mental-representation signals at all levels as "perceptions," rather than using different terms for low-level and high-level representations.

Comparison implies two things to be compared. In the hierarchical control-theory model, one of them is a mental representation, a perceptual signal, indicating the current actual state of the perceived world, or one aspect of it. The second is also a mental representation, a signal, but it represents the state of the same aspect of the perceived world as it is intended to be perceived. This is the reference signal. A comparator is simply a device that receives these two signals and emits an "error" signal indicating the difference between the two inputs to the comparator. A less pejorative term is "deviation." An error signal does not indicate a mistake. It simply indicates by how much and in what direction the current perception deviates from the current setting of the reference signal. That indication drives the corrective actions of the control system.

All that the organism can know about the environment exists in the form of mental representations, perceptual signals. The organism can't know the actual states of its physical inputs (although an intelligent enough organism can certainly make models of the external environment, and thus provide itself with a highly plausible story about what they are). When I say "an organism," I mean every human being, as well as our coevals of other species. The environment that is directly experienced by a human organism is confined to the set of all perceptual signals (although they are not all consciously experienced at once). Wayne Hershberger disagrees with me. But I agree with me.

Evaluation of behavioral-path consequences can be done through the imagination connection. A system of higher level normally acts by sending reference signals to lower-level systems. Those reference signals specify the states to which individual lower-level systems are to bring the kind of perception that each controls. Copies of the resulting perceptual signals become inputs to the perceptual function in the controlling higher-level system. When lower-level control succeeds, as it usually does, the result is that each lower-level system sends upward a perceptual signal that matches the reference signal it is receiving from the output of the higher system.

Exactly the same effect can be achieved if the higher system sends its output not to the comparator of the lower system, but back into its own perceptual function. It is just as though the lower system had succeeded perfectly and instantly. This is what I call the imagination connection. With this connection in effect, the higher system can quickly go through possible outputs (I assume a level where complex logical processes are occurring) and judge their effects on the controlled variable. Thus selection of lower-level actions (and their perceptual consequences) can be done without actually producing any actions.

This process of mental planning is undoubtedly more complex than I make it here. Modeling must be involved, in the imagination path, because the properties of the outside world (which includes all lower-level control systems) must be taken into account. But the basic picture of how imagination works seems to explain the broad outlines of planning of all kinds—not just behavior-path planning.

Behavior always follows some path. The question is whether the paths are in fact always planned, or whether they are simply the result of the way a control system gets from a state of error to a state of no error. Planning of behavior paths is not necessary in all cases—in fact, it is necessary in very few cases. To see whether a path is planned, one can introduce disturbances and see if their effects on the path are corrected, or if the organism simply accepts the deviated path and reaches the goal anyway. The latter is probably the more likely outcome. Paths would be planned in advance only when they make a difference to the

organism. Control systems do not have to precalculate behavior paths.

You can say that an apple is redder than an orange, or cheaper, or better-tasting. But the control process is separate from that comparison, which is really a judgment of relationship. Given the perceptual comparison, you must still specify what the goal is: are you going to paint the orange to make it as red as the apple, or is the difference in redness OK with you? Are you going to raise the price on apples, or inject something in the orange to make it taste better? The goal has to be stated if control is to be involved. And then the comparator—an element of the model—must take the perceived relationship between apple and orange, compare it against the desired relationship, and judge it as being not sufficient, just right, or overdone—relative to the preferred state.

We determine what the goal will be; the environment doesn't. The environment may provide us with a selection of experiences from which to pick feasible goals, but it doesn't do the picking. The environment determines what we must do in order to have the desired effect on experience. If we can't do what it requires, or if the desired result is impossible, then we fail to control.

Neurologists tell us that human beings are basically a set of neural connections. Biochemists tell us that behavior is controlled by interactions among molecules. Sociobiologists tell us that it is genetic fitness to reproduce that determines how we shall act. Physicists tell us that thermodynamics and quantum uncertainty are the keys. Radical behaviorists tell us that schedules of reinforcement are what do the trick. Personality psychologists tell us that traits and attitudes and feelings and aspirations account for behavior. Sociologists tell us that the individual is simply an expression of the society. Existentialists tell us that individual being is at the core of it all.

Doesn't this strike you as a bit suspicious? All these answers, and they all show that the particular interests of the explainer just happen to contain the correct solution to it all. But when you ask any of these explainers how their explanations work, you run into a blank stare. The explanations *are* how it works. They don't ask what lies beneath the explanation. They don't try to link their own field of study to the fields of study of others. It's all extremely provincial and, aside from the specialized expertise involved, superficial.

Control theory crosses all these boundaries because it is concerned with the how of behavior more than the what. It has nothing specifically to do with society, or even with any particular individual behavior. All examples of behavior, all aspects of behavior in any discipline, are grist for its mill. The world it addresses is larger than that of any existing discipline.

Wayne Herschberger I have a nagging itch demanding to be scratched. Bill thinks I wants physics to be part of the immanent order. No. I would say that physics is a science, involving conceptual modeling, as I imagine Bill might say. I would say that there is order immanent in the phenomenal domain that is modeled by physics. I use Bill's word *model* to refer to the intellectual achievements of physicists. That is, I use the word *model* to denote something human-made. Unfortunately, the word *model* has another, unintended, connotation: a replica of an original. Like Linus Pauling, I do not regard scientific models as being replicas of divine (Noumenal) originals. Theoretical physics does not involve "reading God's mind." I view Einstein's saying that it did as a metaphor.

Bill says, "... you assume... that phenomenal objects and attributes of objects are something other than neural signals. I assume they are the same thing. How do we get past that?" As I see it, the issue is a difference between what your theory assumes, and what you say your theory assumes (or implies). I seriously doubt that your hierarchical control theory necessarily implies (or assumes) that phenomenal objects are neural signals. In claiming that your theory is not solipsistic, I find myself in the paradoxical position of arguing that your theory is better than you say it is. That is a sort of disagreement, but one that I think belies a fundamental agreement.

Let me say some things about phenomenal objects, because such descriptions comprise the specifications which we are attempting to reverse-engineer. Please understand that what I say is not presented as an alternative to your theoretical model. What I am trying to do is describe some of the specs that all our psychological models must be able to realize.

Phenomenal objects are simply the particulars of experience. They are the constituents of the empirical world that we are wont to call things. The layman calls them objects or physical objects, and supposes that their substance is essentially material. In contrast, philosophers such as Bishop Berkeley called them perceptions and supposed that their substance is essentially mental.

It seems to me that arguing whether phenomena are substantially mental or material is much the same as arguing whether a magnet is essentially a north or a south pole. The argument makes no sense to me, because phenomena, like magnets, appear to be bipolar, with each instance involving an observer-observed (knower-known) dipole. A dipole, not a dichotomy. For instance, the visible surface of every phenomenal object in my study is the one facing that ubiquitous phenomenal object I have learned to call myself. Inasmuch as this personal "perspective" inheres in every phenomenal object, there is more of me to be found in the phenomenal world than is to be found in the phe-

nominal object I call myself.

This widely distributed aspect of myself which permeates the phenomenal world lends a proprietary aspect to the phenomenal world, making it mine, as it were. That is, the phenomenal world presents itself as a personal "perspective" with that unique point of view being tied to the phenomenal object I call myself. (I put the term "perspective" in quotation marks to signify the observer-observed relationship noted above: a dipole, not a dichotomy.)

Locating oneself is an empirical matter, and does not involve merely locating one's brain, as Dennett, for one, has nicely illustrated in his delightfully humorous essay, "Where am I?" Locating oneself involves a determination of the spatial relationship obtaining between what might be called the sentient self and the sensed self, or what William James would have called the relationship between I and Me. In my own case, Me is the human male residing (i.e., located) at 436 Gayle Avenue in DeKalb, Illinois. I, on the other hand, am distributed throughout my phenomenal world. If I am to be assigned a single spatial location, it must be in terms of an interpolated personal station point, or personal point of regard, defined by the personal perspective immanent in the phenomenal world called mine. Normally, my personal point of regard (i.e., I), appears to coincide with Me, particularly Me's head.

When persons are asked to point directly at themselves, they tend to point at the bridge of their nose (i.e., at Hering's virtual cyclopan eye). The fact that they are then pointing at their brain is accidental. Imagine a set of Siamese twins in which the brain in head X is connected to the nerves of the body attached to head Y, and vice versa. If a flash card bearing the request, "please point at yourself" is presented only to the eyes in head X, at which head would the pointing arm likely point? At X, surely. And if the request were "please point at your brain," at which head would the arm likely point? Might there not be a different reply to the two questions? And if the hand points at heads X and Y, respectively, in response to these two requests, who would have the authority to question those answers? (By the way, I see none of this as being inconsistent with hierarchical control theory.)

Bill, the same can be argued about the relationship between the We and the Us. The two of Us, You and Me, are in Durango and DeKalb, respectively, but We, You and I, have come together in a dialogue, searching for a common perspective, point of view, or parsing of the world. That is, the proprietary aspect of the phenomenal world includes Our as well as Mine. For one thing, I can imagine (project) my phenomenal world as if from various points in phenomenal space, including those that are currently occupied by other individuals. More importantly, I escape an exclusively personal perspective simply to the degree that I demonstrably share a common perspective with others.

That is, I escape epistemic isolation (solipsism) not by dint of effort, but simply by default. I cannot imagine how colors look to a dichromat (they sort pigmented chips differently than I), but I've got an excellent idea about the trichromat's phenomenal world of colors, without even trying—because we judge (see/sort) pigmented chips alike. Claiming that people who sort all possible pigmented chips perfectly alike do not necessarily see colors alike, as some mischievous philosophers are wont to say, presupposes a fictitious absolute standard of comparison (Noumenal color), because the claim of a difference without a super-ordinate frame of reference is totally meaningless; further, if such a fictional frame of reference is assumed, for sake of argument, in order to allow the claim to acquire a certain syntactical sense (as does the statement “all invisible things are red”), it still is devoid of empirical meaning. I submit that a putative difference that makes no difference in phenomenal fact, is in fact no difference.

Whereas it is easy to escape epistemic isolation from others, it appears to be impossible to transcend the phenomenal world itself except metaphorically, that is, by a leap of intellect. We might *imagine* a noumenal world of “things in themselves” that transcends all experience, but that is not what science does or should be doing, according to the likes of Pauling and Bridgeman. The theoretical models that scientists conceive must be able to generate precise predictions in the phenomenal domain, because that is where the truth value of the models must be tested.

Science models the order that is immanent in the phenomenal domain. Physics is the branch of science that models the aspect of the phenomenal domain that we call the environment. Physiology models the aspect of the phenomenal domain that we call organisms. That is, physiology and physics conceptually model those aspects of the phenomenal world laymen perceptually model as Me/Us and The Environment, respectively. In contrast, Psychology is a science concerned with the conceptual modeling of the I and the We. The psychology of perception is that branch of the science concerned with the problem of modeling the observer-observed dipole as such. That is, when one models the putative process said to underlie the perceptual aspects of phenomena, one may be said to be modeling a modeling process. In other words, you and I are here involved with conceiving perceiving, or of conceptually modeling perceptual modeling.

When I try to imagine phenomena's substance from a psychological perspective (i.e., the essential substance of the epistemological dipole) I find myself coming up with words like immanent order or detectable structure or information—all of which are compatible with physics and physiology. It does not appear inappropriate to call such information “signals,” but it does appear inappropriate to call them “neural

signals,” thereby excluding all other signal types, because that is to forget the bipolar nature of the phenomena. *The epistemic unit is the dipole*. For example, I comprise a dipole characterized as me *and* my environment. In your model, this epistemic unit takes the form of an ecological control loop having two poles, characterized as a unique organism and its environment, including all other organisms. Because there are as many dipoles as there are organisms, with each organism being part of many dipoles, your control theory model is not necessarily solipsistic.

Because a single organism plays a unique role in each of these dipoles, it is tempting to suppose that the dipole is within that unique organism. That is, it is tempting to suppose that I am in my head, but that notion is not only illogical, it is also contraindicated by the fact that my phenomenal head is in my phenomenal world—along with a bunch of other phenomenal heads. Therefore, whenever I use the word “perception” to denote this personal aspect of phenomena, I try to remember that I am referring to a personal perspective or point of view rather than to a personal replica.

The *bipolar* nature of *objective* phenomena is what our reverse engineering must explain. Your hierarchical-control-theory model accounts for both of these in terms of interacting control loops. As far as I can see, your model poses no epistemological problems, and it disturbs me to hear you imply, sometimes, that it does. If anything, your model promises to resolve epistemological problems, not create them. That’s the way I see it.

Bill Powers: Wayne, I’m going to avoid the temptation to get back into the epistemological argument; I’ll let you have the last word. I like your exposition considerably, but there are still problems to work out—like what we should say neural signals are for, in our models.

Levels of Perception

Mark Olson: (To Bill Powers:) I'm having difficulty understanding the quality of the input signals for higher-level systems. I know that an input signal for one system is an integration of () from lower systems—what's in the ()? Are the signals that become integrated simply the same as the signals the first level receives? Or are these signals somehow adapted? All the diagrams I've seen make it seem as if they are not adapted, that they simply go all the way up. But if they are integrated, then something is different. I'm close but I don't quite have it.

Bill Powers: Mark, you ask about the functions relating one level of perception to another. This is indeed the question that hierarchical perceptual control theory (HPCT) poses—but doesn't answer. What lies behind HPCT is not any proposal as to how each level of perception is derived from the one below it, but a proposal as to what the levels of perception are and how they are related. This is the phenomenon that any model must in the end explain.

The "H" part of HPCT can be taken in two ways: first, as a general sketch of a hierarchy of control in the abstract, with the communication between levels consisting of a series of perceptual re-representations of reality and a corresponding set of reference signals used to control lower levels; second, as a series of proposed levels of perception (and control) based directly on an analysis of experience with the hierarchical-control concept as a guide. This is a beginning model; there might well be other modes of communication between levels, but the basic one is probably valid.

The definitions of levels define the modeling problem. We can see that the sensation level is probably derived by weighted summations of intensity signals, the weights defining a vector in a perceptual space having fewer dimensions than there are different sources of intensity signals. But that answer to the modeling problem comes after noticing that sensations seem to depend on intensities in a particular way, a way that could be modeled as weighted summation. The phenomenon to be modeled comes before the model.

And that's as far as I can go. I don't know how configurations are derived from sensations—how it is that we can get the sense of, say, a particular person's face over a range of distances and orientations and expressions. If signals standing for the dimensions of a face existed, then it's possible to make a rough guess that transitions of the face from one state to another would be sensed using time functions and

partial derivatives; that's a feeble start toward a functional model that you could run on a computer. As to the rest of the levels, the kinds of computations involved are mostly a mystery to me. The few guesses we have come up with are strictly stabs in the dark. You can use words like "integration" to describe how some kinds of perceptions are put together to create others, but the word is just a noise. It doesn't tell us anything about the processes involved.

Behind this exploration of perception lies a fundamental postulate; if you don't internalize it, I don't think you can even get started on the problem of modeling the brain's perceptual systems, or, for that matter, in understanding HPCT. The postulate, simply put, is this: it's all perception. By that, I mean that no matter what you attend to in the world of experience, whether you refer to inner or outer experiences, concrete or abstract, verbal or nonverbal, the object of your attention is a perception. You are looking at or otherwise experiencing the brain's perceptual activities, not the objective world itself.

Vision is the most important sense to understand this way if you're sighted; understand vision and the rest (touch, taste, sound, etc.) will follow. The world you see begins as pixels (individual picture elements). The pixels are so close together that you see no spaces between them, although the sensory nerves do not overlap and in fact do not completely fill the retina. There's a world between the pixels, but we don't see it unless the view shifts slightly—and then what we had been seeing disappears into the cracks between the pixels. This is invisible to direct experience; the world seems continuous over the whole visual field. We get a sense of seeing the world at infinite resolution, and can't imagine what the whole field would look like if we had, say, ten times as many retinal receptors and the optical acuity and brain power to take advantage of them. This would be like seeing the world through a magnifying lens, except that the whole world would look that way, not just one little part of it (which we still see at human resolution). The only way to imagine this is to go the other way: view the world at a lower resolution, as in a halftone photograph or a television screen seen close up, and imagine that the result is the only world you can ever see. That's how our picture of the world would look to a different organism with higher visual resolution. But we experience it as having continuous detail right down to the level where it appears smooth. I suppose the fly sees the world in the same way. But its world is smoother than ours.

Building up definitions of the rest of the levels in the hierarchy is then a matter of noticing persistent types of structure in this world of picture elements. The first level above the pixels themselves is sensation, a type of perception that can't be analyzed in any way except into variations of intensity. Color is a sensation, as is shading.

Perhaps things like edges are sensations, derived in one step from the pixel distributions. When analyzing perceptions, however, don't use any data but your own experience. Theory and neural data will tell you that in the visual field, in the retina itself, all edges are enhanced, so that there is a strong outlining effect. But look at the edge of a sheet of paper on a dark tabletop. There is no outline. The closer you look at the edge, the more nearly it seems to be an infinitely sharp line separating uniform white from uniform dark. The edge itself is there—but you can't see it as an object. It's just a sense of edginess. Only under special conditions, as in looking at a smooth gradient of illumination going over a relatively short distance from white to black, do you see edge effects like the "Mach band," the only clear subjective evidence of edge enhancement. However those neural signals enhanced at edges are processed, the result is that step changes look like step changes, not outlines as in cartoons. Whatever model we come up with for how the nervous system processes pixel information, it must result in edges that look this way, without borders. If it doesn't, the model is wrong.

The next step is to notice that the edges and corners and broad white areas of the piece of paper add up to—a piece of paper. If you've made this transition properly, it will come as a surprise. Where did that piece of paper, or piece-of-paperiness, come from? It wasn't there in the edge, or the corner, or the whiteness, or the darkness. It comes into being only when all those elements are seen grouped into a thing, a configuration with a familiar shape, orientation, distance, size, and so on. The Gestalt psychologists of old spent a lot of time looking at things like these. They should have kept going. Or perhaps they shouldn't have been cowed by the behaviorists.

You have to go slowly and by the smallest steps you can devise. If you go too fast, you'll miss the smallest steps; if you miss the smallest steps, you'll lose the sense of examining perceptions and start projecting the visual field into an external world again. You'll jump to the more abstract levels and lose the connection from one level to the next. This is, if you like, a form of meditation on experience in which you distance yourself from experience and look at it merely as a display. You're not trying to see anything about the world, but only something about the display. You're trying to see what features the person who constructed it thought of putting into it, just as when you read a program, you think to yourself, "Now he's setting up an array to hold the results," instead of just reading the code, or when you read a novel as a literary critic, you think "Now he's introducing tension," instead of just getting tense. Who the "he" is is immaterial—the point is to see what is before you as a construction that has inner organization, and to try to see how it is put together.

The general principle is that when you have found a level, like sen-

sation, the next level is going to depend on it; also, the current level depends on the one below it. If you analyze a perception to see what it is made of, at first you see just more perceptions of the same level—big configurations are made of little configurations. But when you analyze in just the right way, you suddenly realize that all configurations, of whatever size or kind, are made of sensations, which are not configurations of any kind. And you realize that if it weren't for the presence of those sensations, there couldn't be any configuration to see: a field consisting of a single sensation, such as white, can't lead to any sense of configuration. There's a relationship between these levels of perception. That gives us a hint about building models of perception, a hint about how the brain's perceptual system is constructed.

Sometimes you will identify what seems to be a higher level of perception, some characteristic common to all perceptions, unconnected to lower levels you have previously seen. Then you can use this kind of analysis to try to fill in the gap. What is this new perception made of, besides smaller perceptions of the same kind? When the gap is large, the missing steps are obvious. You can, for example, look at spatial relationships such as "on"—something being "on" something else. You can see the on-ness clearly, it's right in front of you. But what is it made of? If you said "sensations," you would clearly be making too large a jump, because on-ness involves objects, things, configurations. Some kind of object is "on" some other kind of object. If it weren't for the impressions of distinct objects, there couldn't be any sense of the relationship between them. But is that step small enough? I've had to put two levels between relationships and configurations: transitions (which can be zero) and events (which can be as simple as mere duration). Seeing something "on" something else involves more than a brief contact; there must be duration.

Perhaps someone else could find smaller steps still, or would characterize the intervening steps differently. There's still a lot of room for improving the definitions of the phenomena we're hoping ultimately to model.

I'm not talking here about the models themselves. I'm talking about the attitude you take toward your own experiences when you're trying to notice phenomena that need modeling. If you were a physicist, you wouldn't be taking this attitude. You'd treat the world of perception in the normal unanalytical way, as if it lay outside yourself where everyone could see it, and you'd search for laws relating changes of one kind of perception to other kinds of perceptions. You would call these "natural laws" or "behavioral laws" and assume you were discovering truths about an objective universe.

As a control-theory psychologist, however, you have a different objective: to grasp the natural world as a manifestation of human per-

ception (your own), and to ferret out of it some regularities that tell us about perception rather than about the world perceived. If you stumbled onto this attitude accidentally, without understanding what you were doing, you might well find yourself in a state with a clinical name: dissociation. I don't recommend this attitude as one suitable for ordinary living. It's difficult and uncomfortable, and it tends to strip the meaning from experience (until you get past a certain point, after which you realize that meaning, too, is perception, and let it back in). If you're afraid that understanding your girl friend as a set of intensities, sensations, configurations, transitions, events, relationships, categories, sequences, programs, principles, and system concepts in your brain might strain your feeling toward her (and hers toward you), don't do this with your girl friend. Do it with somebody else's, or a laboratory rat. It doesn't matter who or what you do it to, because you're really talking about your own perceptions. This is a private experience valid only in one person's world. It can become public only to the extent that different people independently arrive at the same analysis. I've always hoped for that, but only a very few people, to my knowledge, have tried this for themselves. Most people just memorize my definitions, which unfortunately are in words. It's easier to push words around than to shut up and examine direct experience.

You'll hear objections to this process, alluding to introspectionism, which failed to get anywhere a long time ago. But introspectionism didn't fail because it looked at the kinds of things I'm talking about here. It failed because it confused the subjective with the objective (and so did its critics). The world that I'm speaking of examining here would be called, by most conventional scientists, the objective world, not the subjective one. I'm not recommending shifting attention off the objective world and plunging into the dim and uncertain world of inner phenomena—or what we imagine to be inner phenomena. I'm recommending a change of attitude toward the world we normally consider to be the objective one, which includes the world outside us and our bodies as we experience them. I'm saying that you will learn something if you look on this world as directly experienced evidence about the nature of your own perceptual system, and only in a conjectural way about the world that is actually outside you.

Instead of treating relationships like on, beside, after, with, and into as properties of the external world, look on them as perceptions constructed on a base of lower-level perceptions. Instead of seeing categories as made of things that are inherently alike, think of categories as ways of perceiving that *make* things appear to be alike—things that are actually, at lower levels of perception, different. Instead of seeing sequential ordering as a fact of nature, see it as a way of putting ordering into an otherwise continuous miscellaneous flow. In short, take noth-

ing about experience for granted, as if some aspects of experience were really outside and others were inner interpretations. Put the whole thing inside, and see what you come up with when you understand that it's all perception. All of it.

In HPCT diagrams, we show signals coming out of perceptual functions and going into higher-level ones (as well as the local comparator, if the signal is under control). I think of these lines as representing single neural signals that vary in only one dimension: how much. This can be confusing, because we don't experience single signals under normal circumstances (when we do, they cease to be meaningful). Instead we experience all the signals within the scope of awareness, at every level in the state we call conscious. To understand what the single-signal concept means, you have to break this world of simultaneous perceptions into its components, the individual and independent dimensions in which the totality of perception can vary. You have truly identified one isolated perception when it can vary only in the degree to which it's present, which we experience as its state. If the perception varies without in the slightest changing its identity, you have probably noticed a single signal.

This can be important when you talk about control. We talk loosely about controlling "a dog," for example. But that way of talking is really lumping many independently variable aspects of the dog together. You don't control its species, or its eye color, or the length of its tail. You don't even control its behavior. If it's behavior you're controlling, you always control *some particular variable aspect of the dog's behavior*. You might control the radius within which it can move by putting it on a chain. You might control its speed of walking by saying "stay" or "follow," and its path by saying "heel." Whatever you control, it must come down to a single variable or small sets of variables independently controlled. If you're controlling in more than one dimension, you must sense more than one variable and have a control system operating independently for each one. That's because independent dimensions can be independently disturbed; you need independent control systems so that a disturbance in one dimension can be corrected without necessarily causing an error in another dimension.

None of this answers your question as to how perceptual signals in a diagram depend on perceptual signals lower in the diagram. The only general answer I can give is that some computation lies between them. The input data consists of lower-level perceptions; the output data, the higher-level perceptual signal, represents the value of the function being computed over and over or continuously. At each level, I presume (judging from the way the context changes every time you consider a higher level), a new type of computation is involved, not simply a repetition of the kind of computation at the lower level. The process

of deriving categories from sets of relationships can't be carried out by the same kind of computation that derives relationships from sets of events or lower perceptions. There is no one kind of computation that could serve at all levels.

But as I say, I am—we all are—a very long way from grasping what these kinds of computations are. Every time people come up with a new computer program for recognizing objects, they try to establish this new computation as the blueprint for the whole perceptual system. This is a waste of time. The blueprint changes with every level. Weighted algebraic summation is simply not going to suffice to model our capacity to recognize and execute a program described in words: a rule. Even though such networks are purported to recognize categories, I think that the category-ness is read into the results by a human observer. I don't think that any category-recognizing back-propagation model will actually create what human beings experience as categories—for example, the category “wife.” Of the 11 levels of perception in my model, I think we know how to model two of them: the first two. All the rest of our modeling presents to us what a human being might recognize as a higher-level perception, but which the circuit or program itself does not recognize or control.

In that I could be wrong, of course, because I speak the truth when I say I don't know how the higher levels of perception work. That means I don't know how they don't work, too. I'm just expressing a hunch.

On Modeling

Bill Powers: Of the hundred-odd people on this net, I don't suppose more than a handful understand what some CSGers mean when they talk about modeling behavior. So I thought I'd explain it a little, at least as the process appears to me. Talking about modeling is a little like talking about control—most people have some concept to go with the word, but not many outside the engineering professions (and not everyone in them) mean what I mean by it.

I'm working now on a model of pointing behavior. On the surface, it's not very impressive. The computer screen shows a little stick man with one arm who reaches out and touches, or continuously tracks, a floating triangle that the user can move around from the keyboard in a perspective drawing of a three dimensional space. It looks like a cartoon of a not very interesting behavior. While movements are a bit more realistic than you find in most cartoons, most people have seen more impressive TV cartoons in which more interesting action occurs. But behind this surface appearance is the model; what's interesting is not so much *what* happens on the screen, but *how* it happens. To explain how it happens, I have to distinguish the kind of modeling I use from other kinds.

The first distinction of importance is that this kind of model is not an animation. That is, the various movements of the arm (and head—the little man always looks at the target) are not simply drawn frame by frame as in the Disney Studios. It's not done the way interactive video games are done, by switching from one animated sequence to another depending on what the user does at the keyboard. Instead, the program is reacting directly to the location and movements of the floating triangle, which are totally unpredictable by the program. I can guarantee that the program makes no attempt to predict the target movements, because I wrote it.

The second distinction of importance is that in this kind of model, there is nothing in the program that computes the actual movements of the arm as we see them. If the arm's fingertip moves in a straight line, this is not because something in the program computes the detailed actions needed to produce a straight line. Likewise for curved movements, or movements that begin fast and slow down as the fingertip nears the target. None of these aspects of movement corresponds to any specific calculation of path or speed in the program.

In some approaches to modeling, such calculations are the heart of the method. One looks at the actions and figures out what commands

would be needed to produce them. If the fingertip is to move along a path and intersect a moving target, such a model would use the target movement information as input and find a path and a speed profile that would bring the finger to the same place as the target some time in the future. Then it would drive the computed arm so as to achieve that path and speed profile, thus bringing about the predicted intersection. Basically, this concept of modeling attempts to reproduce the visible behavior by calculating its details, given all of the physical factors of the situation.

The approach I use is more properly called "simulation." Inside the computer are program modules. Each module computes what some simple element of the real system would do when presented with continually varying inputs. Some of the modules are perceptual modules: they compute what certain nerve signals would do as the aspect of the environment to which a sensor is sensitive changes its effects on the sensor. For example, one module represents a muscle spindle, which emits a signal that depends both on the length of the muscle and on another neural signal, the gamma efferent signal. Another represents the tendon receptors that are affected by the muscle tension.

One of the modules is an effector module: it represents the muscle's response to a motor signal from a spinal motorneuron (including the shortening of its contractile part and the consequent stretching of its spring-like component to produce a force). And there are many more modules that represent the way hypothetical sets of neurons respond to neural signals by producing more neural signals. There are sets of modules that are repeated, with the same interconnections, for each muscle in the model.

In this model, by the way, I don't use actual models of individual neurons, although I could. Such a level of detail would not add anything to the performance of the model and would increase the size of the program and slow its operation. What I do instead is use simple calculations similar to what a neural model would do: add signals, subtract one signal from another, amplify signals, and do time integrations and (rarely) differentiations. Nothing more complex.

Each module is meant to represent the way some small part of the real living system works, as nearly as I understand it. Many of the modules represent guesses based on hints from neurology or even from waving my own arms around and paying attention to the details, and they constitute the conjectural parts of the model.

The model is not just a collection of computing modules: it is also a pattern of connections joining one module to one or more others. For example, there are modules representing the static and dynamic parts of the stretch receptors in muscles. The outputs of these modules, conceptualized as neural signals, become inputs to the module represent-

ing the spinal motorneuron. This motorneuron module produces an output that is the sum of several positive inputs from other modules and a negative input from the tendon receptor module. The output of the spinal motorneuron module becomes the input to the module that computes the muscle force output. And so on. Each module is woven into the whole model through its input and output connections from and to other modules.

A more subtle aspect of this process is that the model contains adjustable parameters in the links between modules. The dynamic stretch receptor module, for instance, sends its signal to the spinal motor-neuron module, but there's a parameter that determines how much effect this signal is to have at the spinal motorneuron, and the sign of the effect. If the parameter is set to a high value, the simulated arm behaves sluggishly or, at the extreme, breaks into high-frequency oscillations. If it's set to a low value, the arm begins to wobble around, and even goes into ever-increasing low-frequency oscillations. If the parameter has the wrong sign, the arm will behave more and more wildly, until the whole program blows up.

So it's not enough to model the right kinds of components of the real system, or even to connect them into a network like the real neural network, with the right signals going to the right places. The quantitative parameters can be adjusted to make a model, with any given components and any given pattern of interconnections, do completely different-looking behaviors.

Finally, there's a real-time aspect of this "simulation" kind of modeling. All the computations in all the modules are carried out effectively in parallel. One such parallel computation covering all modules represents one increment of real time, dt . In the arm model, dt represents 0.01 second of physical time (regardless of how long it takes the computer to finish all the computations). The last computation is to recompute all the outputs of the modules, so they have all changed before the next cycle when they will be treated as inputs to other modules. This sometimes requires paying close attention to the way the program is written, so that things supposed to be happening at the same time don't accidentally happen in sequence—one dt too late. In an analog computer, this requirement would be easy to meet, because all of the computing components would be acting at the same time. But in a digital computer, where there is only one busy central processor that has to do everything, achieving the effect of simultaneity isn't always easy.

After each round of calculations, all of the modules have new outputs, which become inputs to other modules (or even the same module) at the start of the next time increment. With a dt of 0.01 second, the result is very close to continuous operation, with all signals (inputs and outputs of modules) varying smoothly and simultaneously. The

test to see whether the incremental approach is sufficiently like a true continuous computation is to decrease the size of dt —let each complete computing cycle represent, say, 0.001 second. If the same behavior results, but in smaller steps of movement, then the larger time increment is short enough. It's nice to use longer intervals, so the movements of the model become fast enough to see between breakfast and lunch. The arm model in its present form runs at about one/fifth of real time (on a 10-MHz IBM-AT-compatible programmed in C).

One of the modules is a physical model of the arm. The inputs to this module are three torques being applied by the muscle modules to the three joints during one time increment. Using kinematic equations, calculating Coriolis forces and all that, these torques are transformed into angular accelerations around the three joints (taking the moments of inertia and masses of the arm segments into account). Those accelerations are integrated to produce angular velocity, which is integrated to produce angular position. The three angular positions are inputs to the behavioral model, determining the new joint angles and angular velocities, and the new muscle lengths and rates of change of muscle length for the start of the next dt .

There are two inputs to each muscle control system: an alpha efferent and a gamma efferent. When these signals are varied (for testing purposes), the arm will go through certain motions on the screen. I use a standard test signal which simply switches from a positive value to zero and back again, with a half-second interval between transitions.

What the arm segment being tested *should* do is move quickly from one angle to another, stay there for a half second, and move and dwell for another half second, over and over. What *does* happen, of course, is initially something very different. There are five parameters to adjust, representing five meaningful aspects of the control system: three sensor sensitivities, one sensitivity of muscle contraction to driving signals, and the spring constant of the muscle. Only the muscle spring constant can be estimated from observations and data in the literature. The other four have to be guessed at. Finding the right combinations of values can be done in part through computations, but there are so many interactions and nonlinearities in the model that exact predictions are impossible (certainly for me). So what one ends up doing is changing the parameters experimentally until the arm begins behaving properly, or as nearly properly as possible, without adjusting the parameters of the other control systems, too.

Estimates of parameter values and, especially, of the behavioral effects of varying parameters can be made, but only for small segments of the model such as a single control system for a single joint. Such estimates get you in the right ballpark for each control system's parameters. But it's impossible to write the equation for the whole model and

solve it for the best values of parameters. The equations are all nonlinear differential equations (made more nonlinear when the visual part of the model comes into play), and the interactions among parts of the model are large (extending the arm at the elbow joint affects both arm segments, for example, through inertial interactions). This brings us to the heart of simulations.

The reason we do simulations is precisely that we can't analyze or even understand the whole model at one time. The postulates of the model are in the definitions of the modules. These modules are each very simple and are closely related to simple properties of the nervous system and muscles. So we can easily understand what each module does or is postulated to do.

What we can't easily understand is what will happen when we conjoined the modules together in some specific way, with specific interconnection parameters. Our postulates about the modules completely determine the behavior that is implied; the only problem is that we can't deduce our way from the postulates to their actual implications.

A simulation shows us the implications directly. It says to us, "I don't know what you thought you were modeling, but here's what you *did* model." It's just like a computer program, which does what you told it to do, instead of what you wanted done. A simulation cuts through all the fuzz of verbal explanations and imprecise reasoning about what a particular model *ought* to do. A simulation is a way of finding out the implications of propositions that are linked together in such a complex way that human reasoning is inadequate to reach a conclusion.

Human reasoning becomes inadequate for most real systems with more than three or four components. Even mathematical analysis is usually impossible in the real world, which doesn't fit the idealized forms that we know how to handle analytically. One result of this fact is that people regularly try to fit the real world to those mathematical methods they *do* know how to handle. Every new discovery of some tractable mathematical phenomenon is followed by a hoard of people trying to make nature behave that way. Hence, chaos theory and its application to literally every unsolved problem, particularly in the nervous system. There are phenomena to which chaos theory applies; in fact, chaos was discovered through observing a working simulation of the weather. But in other contexts it's a solution looking for a problem.

An alternative to analysis is simulation. You hook up a model of the system in which the simple components are represented or plausibly conjectured, turn it on, and gaze at what it does. The model then becomes an experimental object. You can play with it, altering its components, their interconnections, and the connection parameters, and learn the effects of each kind of change. Each variation leads you to under-

stand something about the real system. You find out why a given connection is positive instead of negative. You find out why certain connections are present in the real system and others are not. The “why” in every case is simply that the model doesn’t act like the real system in some relatively dramatic way. And you can *see* why it doesn’t.


A simulation is like an X-ray into the real system, showing you aspects of its functioning that can’t be observed directly. Like an X-ray, the simulation can be ambiguous; the observed behavior can be accomplished by more than one plausible model. As with X-ray interpretations, however, we don’t have to rely on ambiguous indications; we can think up alternative diagnostic tests that will rule out some possible models, and with increases in technical skill, we can even open up the system and see some of the connections, even monitor some of the circuit activities. Every added piece of observational evidence narrows the field of models that would behave correctly *and* work by the right means.

There’s another side to the subject of observational evidence. Often the observational evidence is available, but isn’t understood. To say it isn’t understood is to say that there’s no model that needs that evidence. The combined stretch and tendon reflexes are a case in point. These reflexes have been known for close to a century, but nobody has understood what they are for. There have been vague qualitative conjectures, of course. But the arm model I’m working on shows quantitatively what these reflexes do. The tendon reflex controls applied force. The dynamic stretch reflex controls the integral of applied force, or angular velocity. The static stretch reflex controls the integral of velocity, or angular position at a joint. The model shows that with certain values of the parameters, this combination of control systems makes the arm extraordinarily stable, quick to respond to driving signals, and consistent in response over a wide range of external conditions and internal conditions of the muscles. While I haven’t demonstrated this yet, it’s clear now that this combination of reflexes easily compensates for the extreme nonlinearity of the muscle’s tension-extension curve. In fact, when I realized finally how this system works, I was amazed at its cleverness and simplicity.

But those who traced the circuits and measured their details couldn’t have seen that cleverness and simplicity, because not having modeled the system, they didn’t see all the problems that it solves with such economy. These reflexes can be seen as a remarkable design only after you have looked into the problem of controlling a jointed arm in some detail. I couldn’t have designed that system. I simply designed the model to be as much like what I knew about the stretch and tendon reflexes as possible, turned it on, played with the parameters, and discovered beauty.

The whole arm model is built up this way. It behaves as it does because of the interactions among its modules. It reaches out and touches the target, and follows the target around when it moves, and looks at the target, and resists gravity, and moves at various speeds and along various paths in the process, because there is nothing else it can do. We are seeing in this kind of behavior the necessary consequence of organizing a system the way the model is organized. Maybe another organization would also have to behave this way. But this one behaves like a human being, at least at these levels of organization, and to the extent possible, its modules are similar in function to known modules in human systems. The external physics and optics in the model conform to what is known about physics and optics, near enough. Some parts of the model are in one-to-one correspondence with direct observations. Some parts are conjectured. But the X-ray seems to be showing a convincing shadow of the real system, at least as it is seen from this angle.

73



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

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

Inside back cover

Now available:

LIVING CONTROL SYSTEMS II

Selected Papers of William T. Powers

ISBN 0-9624154-2-1, 1989, 277 pages, illustrated, softcover, 5.5" x 8", \$22.00 postpaid (KY residents please add sales tax).

There are 22 previously unpublished papers in this book, including two chapters from the preliminary version of *Behavior: The Control of Perception* ("Emotion" and "An Experiment with Levels") and the complete version of "CT Psychology and Social Organizations." W. Thomas Bourbon wrote the Foreword, and there is also a list of the published works by William T. Powers on living control systems, 1989-1991, which updates the list in the first volume of *Living Control Systems*.

Some Implications of Feedback Theory Concerning Behavior

The Wiener Feedback Model—A Strategic Error

Emotion

An Experiment with Levels

The Illusion of Control

Control Theory for Sociology

CT Psychology and Social Organizations

A Bucket of Beans

Deriving Closed-Loop Transfer Functions for a Behavioral Model, & Vice Versa

Learning and Evolution

An Agenda for the Control Theory Group

Control Theory, Constructivism, and Autopoiesis

Control Theory: The Road to Utopia

On Violence and Aggression

The Good, the True, and the Real

A Manifesto for Control Theorists

After Galileo

The Epistemology of the Control System Model of Behavior

The Fallacy of Behaviorism

Human Nature from the Standpoint of Control Theory

Standing at the Crossroads

The X-Phenomenon

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

Back cover