Chapter 12

Models and their worlds

Except for this introduction, this chapter is a reprint of an article by Bourbon and Powers (1993). I include it because it is a paragon of testing a hypothesis straightforwardly, rigorously, quantitatively, and conclusively. It shows the clarity with which a hypothesis in PCT can be confirmed or rejected. If it seems surprisingly simple in a place or two, remember that scientific method must be explicit about every step of procedure, no matter how simple it may seem to some.

If you do not wish, at this point, to delve into the kind of detail contained in Chapters 12 and 13, feel free to skip on. You can return when you feel the urge.

As Richard Marken says, the tracking task is simple in the same way as were the little balls and inclined tracks used by Galileo in his seminal studies of the acceleration of gravity. We do not intend the tracking task to show what particular acts people take when they are driving a car or drinking water or building a house or painting a picture. We do intend to say that the tracking is controlled by the *same sort of neural organization* that is used in those other pursuits and all others, too, that act upon the environment. No matter how simple they are, experiments in PCT are remarkable because (1) both person and model produce quantified results that can confirm the match quantitatively, (2) the model is a material, functioning device that *can* produce quantified results, and (3) the model is tested not against an average over many people but against a single person.

In a posting to the CSGnet on 14 September 1995, here is what Powers had to say about the article below:

> In the physical sciences, the common way to test a theory is to examine it as a logical or quantitative structure, and see where you could vary conditions in a way that the theory would have to predict has some new kind of effect, something that hasn't been observed before.
>
> You'll see this strategy exemplified in the paper "Models and their worlds". . . . The control-system model is matched to behavior under the condition where a target moves in a regular way and the person makes a cursor track the target. Once the model's parameters are set for this condition, we then change the conditions. First, we vary the regular movements of the target so they become irregular. The same control model, with the same parameters, predicts that the behavior will change in a specific way that maintains the tracking, and in fact the real person does change the behavior in just the same way as the model, quantitatively. Then we introduce a smoothed random disturbance added to the cursor position, so now the position of the cursor depends both on the handle position and on an independent arbitrary variable. The control model predicts that tracking will continue, and that the handle movements will now differ from the cursor movements in a specific quantitative way. When the real person does the same task, the predictions are upheld with good accuracy. So now the control-system model has been challenged twice; it could have failed in either of the latter two experiments. All that would have been necessary to make the model fail would be for the person to have moved the handle in some way other than the predicted way. Since there were no constraints on how the person could move the handle, the success of the prediction was highly significant. It was significant because the model's behavior could have failed to match the real person's behavior. . . .
>
> Sooner or later, we would think of a way to change the conditions that results in the model's doing something radically different from the real

> person. Rick Marken and I [Marken and Powers, 1989a] did that when we did an experiment in which the sign of the connection between handle and cursor was reversed in a way that gave no sensory indication of the reversal (i.e., no bumps or joggles at the moment of reversal). The model and the person both showed a very similar exponential runaway after the reversals—for the first 0.4 seconds or so. Then the person did something to regain control, BUT THE MODEL DID NOT. So by thinking up the right change of conditions, we succeeded in making the model fail.
>
> Of course that failure was simply a signal that we had to modify the model, which we did. We added a second level of control that could reverse the sign of the first-level control action when a runaway condition was sensed. That naturally restored the model to working order, and it once again was able to predict behavior correctly. So by finding a way to make the model fail, we learned how we could improve the model so it would no longer fail under that set of conditions, and of course continued to work properly under all the other changes in conditions we had already tried.

The article that follows appeared originally in the now-defunct journal *Closed Loop*, 1993, 3(1), 47–72. Another version of it appeared in the *International Journal of Human-Computer Studies*, 1999, 50, 445–461. *Closed Loop*, 1993, 3(1), along with several other issues has been restored and is available as a PDF-file at www.PCTresources.com

Models and Their Worlds

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

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

ABSTRACT

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

MODELS AND THEIR WORLDS

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

"The world" is the part of the surroundings on which an organism can act, and which, in turn, affects the organism. Every statement about the antecedents or consequences of behavior either includes or implies notions about how the world operates. Every theory of behavior is, in part, a theory about the world in which behavior occurs.

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

Three Models

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

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

We expect most behaviorists to say that our S-R model is "reflexological"—a version of behavioristic theory many behaviorists disavowed years ago—and to echo the comment: "There may not be a reflexologist alive" (Shimp, 1989, p. 163). Protests aside, at the core of every behavioristic theory is a claim that the environment controls behavior. From the beginning, behaviorists have asserted, like Donahoe and Palmer, "Although the organism is the locus of environmental action, it is the environment, and not the organism, that is the initiator and shaper of behavior" (1989, p. 410). When Hayes and Brownstein (1986) discussed prediction and control as criteria for evaluating behavioristic analyses of behavior, they said, "One could ask, for example, how do we know that *this* is the relevant stimulus for *this* behavior? The answer is of the general form that when we change *this* stimulus (and not *that* stimulus), we get a change in *this* behavior (and not *that* behavior)" (p. 178, emphases in the original). And Skinner claimed, "The ways in which behavior is brought under control of stimuli can be analyzed without too much trouble. . ." (1989, p. 14).

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

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

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

> . . . three stages may be viewed. . . as three components of a single dynamical system mechanistically tied to the organism's nervous system. The encoding of information would. . . correspond to initial inputs, computational rules correspond to transient dynamics, and representations would correspond to the equilibrium configurations resulting from the transient dynamics. The animal reaches a representation of the environment through the operation of specific computational rules applied to a particular pattern of incoming sensory information (p. 980).

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

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

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

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

The perceptual control theory model. The PCT model, which we discuss later at some length, is the least familiar of the three models. In brief, it proposes that there is a simultaneous two-way interaction between organism and environment (see Hershberger, 1989; Marken, 1990; and Powers, 1973, 1989, 1992). In PCT, the basic unit of behavior is not the linear input-output chain, but the negative-feedback loop, which has properties different from the units of the other two models and implies interesting consequences about the way an organism's actions alter the outside world.

"Models"

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

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

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

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

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

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

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

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

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

To avoid problems of this sort, we give cognitive models the same benefit of the doubt that we give S-R models. Given proper knowledge of the history and properties of the environment, and the correct internal computations, the ideal cognitive model should calculate exactly the motor outputs required to produce a preselected result. Of course, even a perfect cognitive model would require experience with an environment to build up knowledge of its properties: if the environment changed, the model would need new interactions with the altered form before it could again compute the correct action.

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

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

Perceptual Control Theory as a Model

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

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

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

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

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

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

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

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

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

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

THE EXPERIMENT

The Task

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

We can easily modify the experiment to include perceptual variables other than spatial position. For example, the handle can be set to alter the size or shape of a geometric figure, change the magnitude of a number displayed on the screen, or alter the pitch of a sound. And tracking can occur across stimulus attributes and sensory modalities, as when a person uses the handle to make the pitch of a sound match the magnitude of a number or the vertical position of a target. All relationships observed during a simple tracking experiment are found in these other tasks; any of them can be used to make the points we make here.

The Conditions: Three Phases

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

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

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

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

The experimental variables. During each 60-second experiment, each variable is sampled every 1 /30 second, for a total of 1800 values per variable. In the figures illustrating the results, every third value is plotted. There are three measured variables: the positions of the target (T), handle (H), and cursor (C).

Phase 1

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

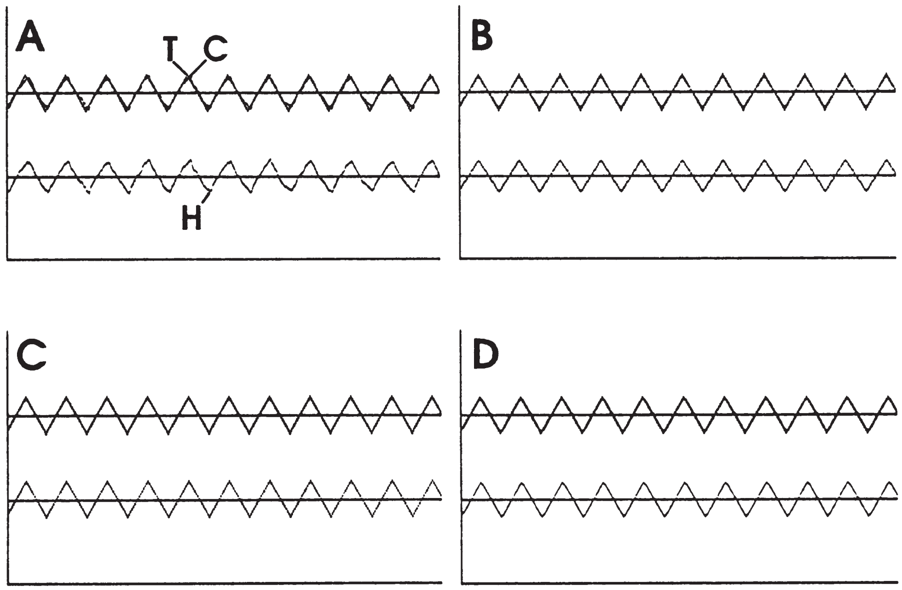

Figure 1. Results of pursuit tracking, Phase 1: data from the person (A); reconstructions of the data by the stimulus-response model (B); by the cognitive model (C); and by the control-system model (D). In A, H = handle, T = target, and C = cursor. For target and cursor, "up" in the figure is toward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

position plot; we use this scaling in all figures).

The mean vertical distance between the cursor and target was -0.8 units of screen resolution (S.D. –1.8; total vertical range on the screen = 200 units). The following Pearson correlation coefficients describe the relationships among variables in Fig. 1A: between positions of the cursor and target, .977; handle and target, .977; and handle and cursor, 1.0. In the regression of handle on target, the slope was 0.89 (the person moved the handle the equivalent of 0.89 screen units for every movement of one unit by the target), and the intercept was -0.8, identical to the average difference between positions of the cursor and target.

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

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

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

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

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

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

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

$h = mt + b,$

and the position of the cursor is modeled as

$c = h.$

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

The positions of handle and cursor created by the model resemble those from the person: the correlation between modeled and actual handle positions is .977; between modeled and actual cursor positions, also .977. Our simple reflexological model accounts for 96 per cent of the variance (r-squared) in the behavioral data from Fig. 1A; the regression equation is highly meaningful.

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

A detailed version of this model would use a program loop simulating a "higher cognitive process" to compute handle positions independently of target movements. It would generate commands for the amplitude, frequency, and shape of the movements. But severe phase errors (mismatches in timing between the positions of the target and the model's handle) would develop unless we gave the model exact information about the frequency of the target and started it at exactly the right moment with exactly the right initial conditions. To assure that there were no errors, we would tell the model exactly how to move the handle to re-create the results of Phase 1. To achieve the same result, without the complex computations, we simply assume that, however the cognitive model works, it works perfectly: it computes handle movements to match the average pattern of previous target movements. For the last minute of practice, it uses information accumulated earlier to command movements that reproduce the movements of the target (of course the model we use here does not actually need any practice).

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

$h = t.$

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

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

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

$c = h + d.$

In Phase 1, the disturbance magnitude is zero.

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

$h = k[\mathrm{int}(p^* - p)].$

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

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

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

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

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

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

Results of running the PCT model. In the person's run during Phase 1, p* was estimated as –1 unit on the screen (–0.8 rounded), which means that, on average, the person kept the cursor slightly below the target. Following the procedure described above, the estimated best value of the integration constant k was 8.64 in units of resolution per second.

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

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

Phase 2

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

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

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

Prediction of the cognitive model. The results for the cognitive model, shown in Fig. 2C, reveal the first obvious failure of a model. The positions of handle and cursor created by the model are not like those from the person: the correlation between modeled and actual handle positions is .230; between modeled and actual cursor positions, also .230.

The reason for this failure is obvious. The cognitive model assesses properties of the environment and computes an action that will have a required result. But now the environment, in the form of target movements, is subject to unpredictable variation. The cognitive model gets no information about the next target speed before it is experienced. Thus, the best that a cognitive "central-process" model can do is command its output to match the best estimate of average target speed; in the present case, that average is the speed that occurred throughout Phase 1, when

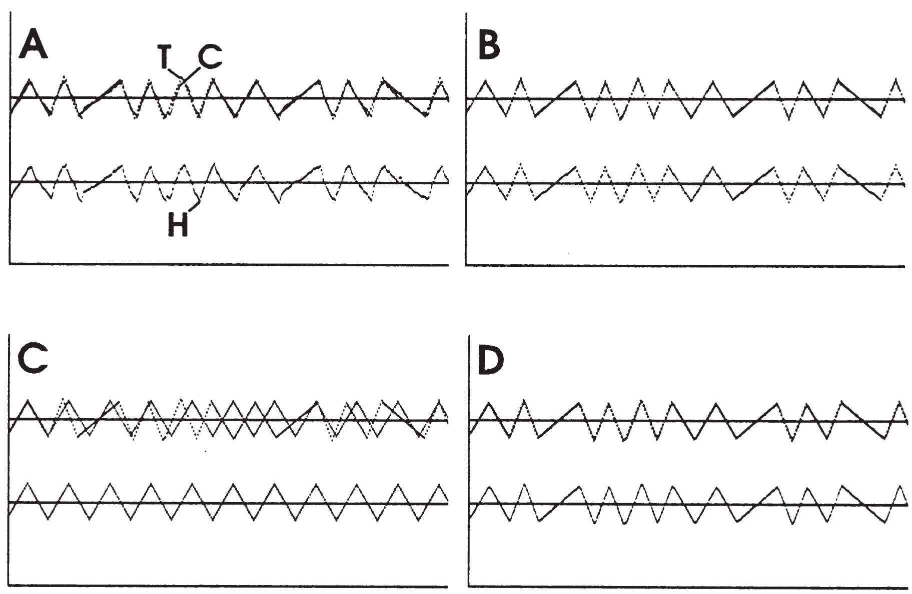

Figure 2. Results of pursuit tracking in Phase 2: data from the person (A); predictions of the data by the stimulus-response model (B); by the cognitive model (C); and by the control-system model (D). In A, H = handle, T = target, and C = cursor. For target and cursor, "up" in the figure is toward the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

the motor plan was established. The cognitive model continued to produce a triangular wave of handle and cursor movement that conformed to the average waveform of target movement—a form not like the waveform of the target in Phase 2.

One might think of modifying the cognitive model so that the central processor re-assesses the environment's properties on an instant-by-instant basis. That would solve the problem, but only at the expense of converting the cognitive model into a control-system model intent on making its output match its input: the new model would be a control-system model acting like a stimulus-response model. The core concept of a cognitive motor plan would be abandoned.

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

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

Phase 3

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

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

The person's data. As we show in Fig. 3A, the person still made the cursor track the target (mean distance between cursor and target = –1.0 screen units, S.D. = 3.0), despite the unpredictable variations in target speed and the unpredictable interference of a disturbance. Had the person not moved the handle, the correlation between positions of the cursor and momentary values of the disturbance would have been + 1.0; that between positions of cursor and target, near 0.0. Instead, the correlation between the disturbance and cursor was only .101, while that between cursor and target was .940.

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

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

Successful simulation can no longer be attained by moving the handle in synchrony with target movements. That is why the person moved the handle in a pattern that deviated radically from the pattern of target movements; the deviations were exactly the ones needed to counteract the effects of the new disturbance. But the S-R model responded to the target stimulus just as before, and moved the handle proportionately to any movement of the target. The simulated cursor, now subject to an independent disturbance, did not follow the target.

To salvage the S-R model, one might propose that the cursor, too, be included in the definition of the stimulus. However, the person's data in Fig. 3A show that the cursor moved in nearly the same pattern as the target, but neither pattern resembled what the

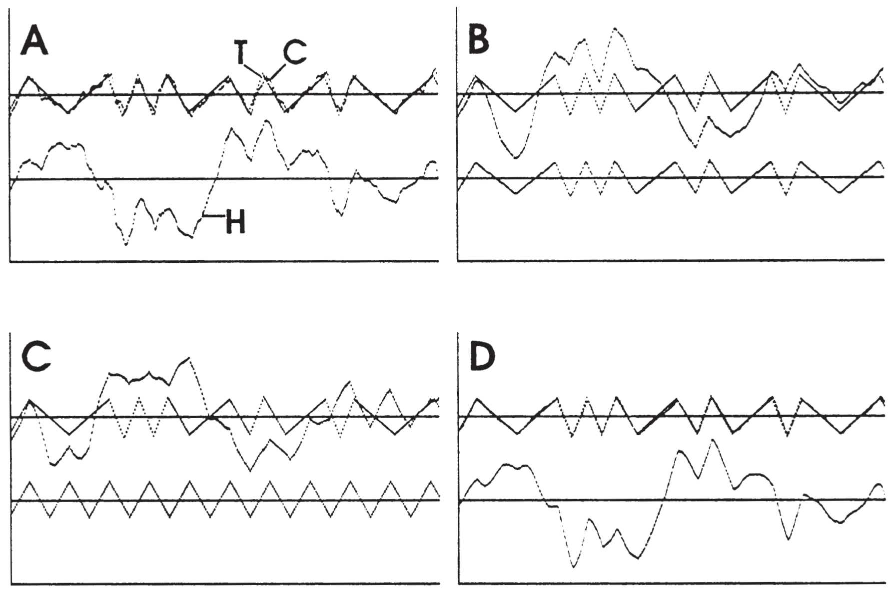

Figure 3. Results of pursuit tracking in Phase 3: data from the person (A); predictions of the data by the stimulus-response model (B); by the cognitive model (C); and by the control-system model (D). In A, H = handle, T = target, and C = cursor. For target and cursor, "up" in the figure is toward the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

handle did. To include the cursor in the definition of the stimulus, we might conclude that the difference between the target and cursor positions is the stimulus. On further examination, we would find that this difference does not match the handle movements, either, but its time-integral does: perhaps the time-integral is the stimulus. That change is acceptable, but if we adopt it, we are left with the fact that cursor position depends, simultaneously, on handle position and the independent random disturbance: now there is no true independent variable in the causal chain, and the core premise of any model of stimulus control over behavior is abandoned. Neither the cursor nor any relationship between the cursor and any other variable can be described as a pure independent variable, because it is also, at every moment, a dependent variable.

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

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

Prediction of the control-system model. As we show in Fig. 3D, the control-system model produced precisely the outputs required to maintain a pre

-selected target-cursor separation, despite two kinds of random variation that called for pronounced changes in the output pattern. The PCT model faithfully predicted the person's behavior. The correlation between actual and predicted handle positions is .996; between actual and predicted cursor positions, .969. Correlations as high as those here, between tracking behavior and predictions by PCT, are commonplace, even when the interval between predictions and behavior is as long as one year as is reported by Bourbon, Copeland, Dyer, Harman, and Mosley (1990).

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

DISCUSSION

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

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

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

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

Similarly, Dar (1987) wrote:

> In physics. . . theories are tighter and lead to precise predictions. As a consequence, (a) if the numerical result is as predicted (that is, close enough to the predicted point value or curve), it will be very difficult, in contrast to the situation in psychology, to

> offer a reasonable alternative theory for that. This is because it is difficult to imagine alternative states of nature that will lead to the exact same curve or numerical result. (b) If the experimental result is not as predicted, some serious revision of the theory would be required. This is because a tight theory simply does not allow for significant (I do not mean "statistically significant") discrepancies from predicted outcome (p. 148).

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

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

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

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

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

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

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

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

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

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

ACKNOWLEDGMENTS

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

REFERENCES

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Powers, W.T. (1973). *Behavior: The control of perception.* Chicago: Aldine. Second edition (2005), revised and expanded, New Canaan CT: Benchmark Publications.

Powers, W.T. (1978). Quantitative analysis of purposive behavior: Some spadework at the foundations of experimental psychology. *Psychological Review, 85,* 417–438.
Reprinted on pp. 129–165 of the following reference.

Powers, W.T. (1989). *Living control systems; Selected papers of William T. Powers.* Gravel Switch, KY: Control Systems Group, distributed by Benchmark Publ., New Canaan CT.

Powers, W.T. (1992). *Living control systems II; Selected papers of William T. Powers.* Gravel Switch, KY: Control Systems Group, , distributed by Benchmark Publ., New Canaan CT.

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

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

Runkel, P. (1990). *Casting nets and testing specimens: Two grand methods of psychology.* New York: Praeger. Second edition (2007), revised and updated, Hayward CA: Living Control Systems Publishing

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

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

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