Feedback: Beyond Behaviorism

Stimulus-response laws are wholly predictable within a control-system model of behavioral organization.

William T. Powers

The basis of scientific psychology is a cause-effect model in which stimuli act on organisms to produce responses. It hardly seems possible that such a simple and venerable model could be in error, but I believe it is. Feedback theory shows in what way the model fails, and what must be done to correct our concepts of organized behavior.

Responses are dependent on present and past stimuli in a way determined by the current organization of the nervous system; that much is too well documented to deny. But it is equally true that stimuli depend on responses according to the current organization of the environment and the body in which the nervous system resides. That fact has been left out of behavioristic analyses of human and animal behavior, largely because most psychologists (especially the most influential early psychologists) have lacked the tool of feedback theory.

Norbert Wiener and later cyberneticists notwithstanding, the full import of feedback in behavioral organization has yet to be realized. The influence of behaviorism, now some 60 years old, is pervasive and subtle. Shaking ourselves free of that viewpoint requires more than learning the terms associated with feedback theory; it requires seeing and deeply appreciating the vast difference between an open-loop system and a closed-loop system.

Traditional psychology employs the open-loop concept of cause and effect in behavior; the effect (behavior) depends on the cause (stimuli) but not vice versa. The closed-loop concept treats behavior as one of the causes of that same behavior, so that cause and effect can be traced all the way around a closed loop (*1*). When any phenomenon in this closed loop (such as the force generated by a muscle) persists in time, effectively averaging the antecedent causes over some period, the character of the system-environment relationship changes completely—cause and effect lose their distinctness and one must treat the closed loop as a whole rather than sequentially. That is where feedback theory enters the picture. Feedback theory provides the method for obtaining a correct intuitive grasp of this closed-loop situation in the many situations where the old open-loop analysis leads intuition astray.

In this article I intend to show as clearly as I can how a new theoretical approach to behavior can be developed simply by paying attention to feedback effects. There is nothing subtle about these effects; they are hidden only if they are taken for granted. All behavior involves strong feedback effects, whether one is considering spinal reflexes or self-actualization. Feedback is such an all-pervasive and fundamental aspect of behavior that it is as invisible as the air we breathe. Quite literally it is behavior—we know nothing of our own behavior but the feedback effects of our own outputs. To behave is to control perception.

I will not try here to develop all these concepts fully; that is being done elsewhere (*2*). I will provide only some essential groundwork by discussing the development of a hierarchical control-system model of behavioral organization beginning wtih the same sort of elementary observations that led to behaviorism. I hope it will thus become evident that a fully developed feedback model can do what no behavioristic model has been able to do: it can restore purposes and goals to our concept of human behavior, in a way that does not violate direct experience or scientific methods. The human brain is not simply a switchboard by means of which one environmental event is connected to another environmental event. These ideas are not new, but perhaps my synthesis is.

Act versus Result

Behaviorists speak of organisms "emitting" behavior under stimulus control, this control being established by use of reinforcing stimuli. The effectiveness of reinforcers cannot be denied, but behavior itself has not been thoroughly analyzed by behaviorists. Behaviorists have not distinguished between means and ends—acts and results (*3*)—because they have not used the model that is appropriate to behavior.

When a pigeon is trained to walk in a figure-eight pattern, there are at least two levels at which the behavior must be viewed. The first, which is the one to which the behaviorist attends, is that of the pattern which results from the pigeon's walking movements. The other consists of those movements themselves (*4*).

The figure eight is created by the walking movements: the act of walking produces the result of a figure-eight pattern in the observer's perceptions. The observer sees a consistent behavior that remains the same from trial to trial. He generally fails to notice, however, that this constant result is brought about by a constantly changing set of walking movements. Clearly, the figure-eight pattern is not simply "emitted."

As the pigeon traces out the figure eight over and over, its feet are placed differently on each repeat of the same point in the pattern. If the cage is tipped, the movements become still more changed, yet the pattern which results remains the same. Variable acts produce a constant result. In this case the variations may not be striking, but they exist.

As behaviors become more complex the decoupling of act and result becomes even more marked. A rat trained to press a lever when a stimulus light appears will accomplish that result with a good reliability, yet each onset of the stimulus light produces a different act. If the rat is left of the lever it moves right; if right it moves left. If the paw is beside the lever the paw is lifted; if the paw is on the lever it is pressed down. These different, even opposite, acts follow the same stimulus event.

The author is a consultant to the Center for the Teaching Profession at Northwestern University, Evanston, Illinois 60201.

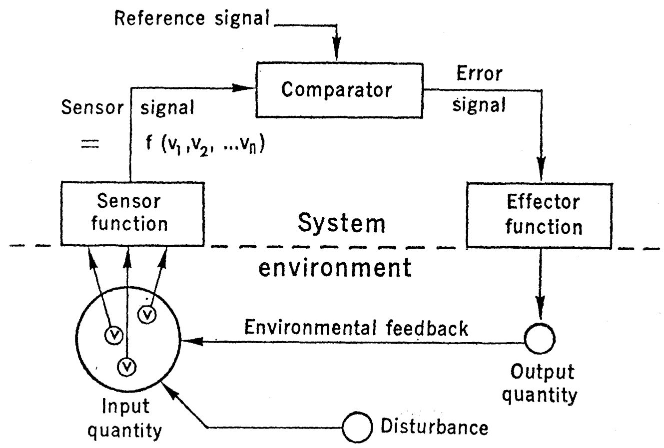

Fig. 1. Basic control-system unit of behavioral organization. The *Sensor function* creates an ongoing relationship between some set of environmental physical variables (v's) and a *Sensor signal* inside the system, an internal analog of some external state of affairs. The sensor signal is compared with (subtracted from, in the simplest case) a *Reference signal* of unspecified origin (see text). The discrepancy in the form of an *Error signal* activates the *Effector function* (for example, a muscle, limb, or subsystem) which in turn produces observable effects in the environment, the *Output quantity*. This quantity is a "response" measure. The environment provides a feedback link from the output quantity to the *Input quantity*, the set of "v's" monitored by the sensor function. The input quantity is also subject, in general, to effects independent of the system's outputs; these are shown as a *Disturbance*, also linked to the input quantity by environmental properties. The disturbance corresponds to "stimulus." The system, above the dashed line, is organized normally so as to maintain the sensor signal at all times nearly equal to the reference signal, even a changing reference signal. In doing so it produces whatever output is required to prevent disturbances from affecting the sensor signal materially. Thus the output quantity becomes primarily a function of the disturbance, while the sensor signal and input quantity become primarily a function of the reference signal originated inside the system. For all systems organized in this way, the "response" to a "stimulus" can be predicted if the stabilized state of the input quantity is known; the stimulus-response law is then a function of environmental properties and scarcely at all of system properties.

The more closely the rat's acts are examined, the more variability is seen. Yet in every case the variations in the acts have a common effect: they lead toward the final result that repeats every time. In fact, if precisely those variations did not occur, the final result would not be the same every time. Somehow the different effects apparently caused by the stimulus light are exactly those required to compensate for differences in initial conditions on each trial. This situation was clearly recognized by the noted philosopher of behaviorism, Egon Brunswik (*5*).

The accepted explanation for this phenomenon of compensation is that the changed initial conditions provide "cues," changes in the general background stimuli, which somehow modify the effect of the main stimulus in the right way. There are three main problems created by this explanation. First, these hypothetical "cues" must act with quantitative accuracy on the nervous system employing muscles which, because they are subject to fatigue, give anything but a quantitative response to nerve impulses. Second, these "cues" are hypothetical. They are never experimentally elucidated in toto, and there are many cases in which one cannot see how any cue but the behavioral result itself could be sensed. Third, the compensation explanation cannot deal with successful accomplishment of the behavioral result in a novel situation, where presumably there has been no opportunity for new "cues" to attain control of responses.

The central fact that needs explanation is the mysterious fashion in which actions vary in just the way needed to keep the behavioral result constant. The "cue" hypothesis comes after the fact and overlooks too many practical difficulties to be accepted with any comfort. Yet what is the alternative? It is to conclude that acts vary in order to create a constant behavioral result. That implies purpose: the purpose of acts is to produce the result that is in fact observed. This is the alternative which I recommend accepting.

Feedback Control

Behaviorists have rejected purposes or goals in behavior because it has seemed that goals are neither observable nor essential. I will show that they are both. There can be no rational explanation of behavior that overlooks the overriding influence of an organism's present structure of goals (whatever its origins), and there can be no nontrivial description of responses to stimuli that leaves out purposes. When purposes are properly understood in terms of feedback phenomena, acts and results are seen to be lawfully related in a simple and direct way. We will see this relationship using a simple canonical model of a feedback control system.

Engineers use negative feedback control systems to hold some physical quantity in a predetermined state, in an environment containing sources of disturbance that tend to change the quantity when it is uncontrolled. Every control system of this kind must have certain major features. It must sense the controlled quantity in each dimension in which the quantity is to be controlled (*Sensor function* in Fig. 1); this implies the presence of an inner representation of the quantity in the form of a signal or set of signals. It must contain or be given something equivalent to a reference signal (or multiple reference signals) which specifies the "desired" state of the controlled quantity. The sensor signal and the reference signal must be compared, and the resulting error signal must actuate the system's output effectors or outputs. And finally, the system's outputs must be able to affect the controlled quantity in each dimension that is to be controlled. There are other arrangements equivalent to this, but this one makes the action the clearest.

This physical arrangement of components is further constrained by the requirement that the system always oppose disturbances tending to create a nonzero error signal; this is tantamount to saying that the system must be organized for negative (not positive) feedback, and that it must be dynamically stable—it must not itself create errors that keep it "hunting" about the final steady-state condition. There is no point in concern with unstable systems, because the (normal) behavior we wish to explain does not show the symptoms of dynamic instability—and we do not have to design the system.

This system is modeled after Wiener's original concept (*6*). In the system I describe, however, there are certain changes in geometry, particularly the placement of the system boundary and the identification of the sensor (not reference) signal as the immediate consequence of a stimulus input. This is a continuous-variable (analog) model, without provision for learning.

A system that meets these requirements behaves in a basically simple way, despite the complexities of design that may be required in order to achieve stable operation. It produces whatever output is required in order to cancel the effects of disturbances on the signal generated by the sensor. If the properties of the sensor remain constant, as we may usually assume, the result is to protect the controlled quantity against the effects of unpredictable disturbances of almost any origin.

Goal-Directed Behavior

The reference signal constitutes an explanation of how a goal can be determined by physical means. The reference signal is a model inside the behaving system against which the sensor signal is compared; behavior is always such as to keep the sensor signal close to the setting of this reference signal.

With this model we gain a new insight into so-called "goal-seeking" behavior. The usual concept of a goal [for example, William Ashby's treatment (*7*)] is something toward which behavior tends over some protracted period of time. We can see that idea now as describing the behavior of a sluggish control system, or a control system immediately after an overwhelming disturbance. Many complex control systems are sluggish, but only because any faster action would lead to dynamic instability. The appearance of "working toward" a goal may result from nothing more than our viewing the system on an inappropriately fast time scale.

It is useful to separate *what* a control system does from *how* it does what it does. Given two control systems controlling the same quantity with respect to the same reference signal, one system might be able to resist disturbances lasting only 0.1 second while the other could not oppose a disturbance lasting less than 1 second. After a disturbance, one system might restore its error signal nearly to zero in one swift move, while the other makes that correction slowly and after several over- and undershoots of the final steady-state condition. These are dynamic differences, and have to do with the details of system design. Both systems, however, do the same thing when viewed on a slow enough time scale: they control a given quantity, opposing disturbances tending to affect that quantity. On a time scale where we can see one system "working toward" the goal state, we might see the other as never allowing significant error to occur—as reacting simultaneously with the disturbance to cancel its effects.

The proper time scale for observing what a control system does is that on which the response to an impulse-disturbance is apparently zero. That automatically restricts our observations of disturbances in the same way: all disturbances appear to be slow. On such a slow time scale, it is apparent that a control system keeps its sensor signal nearly matching its reference signal by producing outputs equal and opposite to disturbances, in terms of effects on the controlled quantity.

The normal behavior of a good control system, viewed on the appropriate time scale, is therefore not goal-seeking behavior but goal-maintaining behavior. The sensor signal is maintained in a particular goal state as long as the system is operating within its normal range, in the environment to which its organization is matched. If the properties of the sensor do not change, this control action results in the external controlled quantity being maintained in a state we may term its reference level.

Much of what we interpret as a long process of goal-seeking (and perhaps all) can be shown to result from higher-order goal maintenance that involves a program of shifting lower-order reference levels, but that anticipates what has yet to be developed here.

Controlled Quantities

The key concept in this model, as far as observable behavior is concerned, is that of the controlled quantity. If it were possible to identify a controlled quantity and its apparent reference level, the model just given would provide an adequate physical explanation for existence of this quantity and its goal state, just as the telephone-switchboard model of the brain has heretofore been taken as an adequate physical explanation for stimulus-response phenomena. To be sure, the source of the reference signal that sets the system's goal remains unspecified, but that is of no consequence in a part-model of a specific behavior pattern. We are concerned here with immediate causation, not ultimate causes.

If a quantity is under feedback control by some control system, that fact can be discovered by a simple (in principle) procedure, based on the fact that the system will oppose disturbances of the controlled quantity.

Suppose we can observe the immediate environment of a control system in terms of detailed physical variables $(v_1, v_2, \ldots v_n)$. We postulate a controlled quantity $q_c = f(v_1, v_2, \ldots v_n)$, where f is a function of the variables. According to the definition and known physical principles, we can then devise a small disturbance d affecting some v's such that (in the absence of behavioral effects) $\Delta q_c = g(d)$, where g is the function describing the environmental connection between the disturbance and the controlled quantity. Applying the disturbance we predict a change in q_c, and compare it with the observed change, $\Delta^* q_c$. If we have hit upon a definition of q_c that is accurate, and if a reasonably good control system is acting, we will find $\Delta^* q_c / \Delta q_c \ll 1$.

By progressively changing the definition of q_c [that is, the form of f in $f(v_1, v_2, \ldots v_n)$], we can find a minimum in the ratio $\Delta^* q_c / \Delta q_c$; that is, we can find a definition of the controlled quantity such that the observed effect of a disturbance is far less than the effect predicted according to physical principles, omitting the behavior of the system.

The reason for the "failure" of the prediction is of course the fact that the control system actively opposes effects of d on q_c. Let h be the function describing the environmental connection between the output o of the system and the controlled quantity. If the output o affects q_c additively according to the relationship $\Delta q_c = h(o)$, then the total effect on q_c is the sum of the effects of the disturbance and the system's active output: $\Delta q_c \cong g(d) + h(o)$. When control is good, this sum will be nearly zero.

Defining the zero points of the controlled quantity and the system's output as their undisturbed values, we can see that the controlled quantity will remain nearly at its zero point $(\Delta q_c \approx 0)$, while the disturbance and the system's output will be related by the approximation, $g(d) \approx -h(o)$.

Here is a very simple example. Suppose we observe a soldier at attention, and guess that one controlled quantity involved in his behavior is the vertical orientation of one of his arms, seemingly being held in a straight-down position (the zero point). If this quantity were not under active control, we could predict that a sideways force of 1 kilogram would raise the arm to about a 30-degree angle from the vertical. Applying the force, we observe that in fact the arm moves only 1 degree, or 1/30 of the predicted amount. The effective force-output of the soldier is thus just a trifle under 1 kilogram in a direction opposite to our 1-kilogram disturbance, the trifle being the restoring force due to the slightly deflected mass of the arm, and gravity. This is a reasonable verification of the initial guess, and we may claim to have found a control system in the soldier by identifying its controlled quantity.

The reference level of a controlled quantity can better be defined as its value when the system's output is totally unopposed (even by friction or gravity). Because that state normally implies no error-correcting output, the reference level of the controlled quantity can also be defined as that level (state, for multidimensional quantities) which results in zero error-correcting output.

A controlled quantity need not have a reference level of zero. The soldier, for example, might be persuaded to raise his arm to the horizontal position, so that in the same coordinate system used before, the apparent reference position is now 90 degrees. The weight of the arm now constitutes a natural disturbance, and we would guess that the system's output is now equivalent to an upward force equal to the weight of the arm. If that force were 10 kilograms, we would also predict that an upward force disturbance of 10 kilograms would cause the arm muscles to relax completely, or at least that the net force-output would drop to zero (arm muscles can oppose one another). Our pushing upward with a force of 11 kilograms should result in an output of ½ kilogram downward.

Hierarchies of Controlled Quantities

Suppose that the soldier is now ordered to point at a passing helicopter. He will raise his arm and do so. We can verify that arm position is still a controlled quantity by applying force-disturbances, but now the picture is complicated. The test still works for relatively brief (but not too brief) disturbances, but over a period of some seconds we find that arm position does not remain constant. Instead, it moves slowly and uniformly upward and sideways, as the helicopter approaches.

This suggests that a second controlled quantity has entered the picture. If the helicopter stops and hovers, this new controlled quantity is invisible—the force-test cannot distinguish it, for the arm simply remains almost still as before. But if we radio the helicopter pilot to move his craft in various ways, we can test the hypothesis that the soldier is controlling the angular deviation of his pointing direction from his actual line of sight to the helicopter. If that were not a controlled quantity, the pilot's moving the helicopter would create a predictable deviation. In fact, movement of the helicopter results in no observable deviation at all (barring slight tremors). We are reasonably assured that the pointing direction relative to the direction of the helicopter (and nothing else) is a two-dimensional controlled quantity, with a reference level of zero deviation.

Now we have a slight dilemma. We established, and could reestablish at any time, arm position as a controlled quantity. (The position-control system will react to disturbances within the lag time of the pointing-control system.) Yet control of the new controlled quantity requires a change in arm position, which would constitute a disturbance of the first system. Why does the first control system not resist this change?

The answer is obvious. The second control system opposes disturbances not by direct activation of force outputs, but by altering the reference level, by means of changing the reference signal for the arm-position control system.

Now two controlled quantities (and implied control systems) exist in a relationship that is clearly hierarchical. One controlled quantity is controlled by means of changing the reference level with respect to which a second quantity is controlled.

This immediately suggests a partial answer to the question raised by Fig. 1: Where does the reference signal come from? It is clearly the output of a higher-order control system, a system that senses a different kind of quantity and controls it with respect to an appropriate reference signal by using the whole lower-order system as its means of error prevention (the appropriate time scale for the higher-order system will be slower than that for the lower).

We now have a plausible physical model for a two-level structure of goals. The goal of pointing is achieved by setting—and altering—a goal for arm position. In fact the higher-order system must adjust reference levels for two lower-order control systems, one governing horizontal arm position and one governing vertical arm position: both can be shown to be under feedback control. Of course we do not know yet the actual nature of the lower-order systems—any two noncollinear directions of control would give the same observed results. But we have achieved a first approximation.

The source of the lower-level reference signals has been identified but the question of the ultimate source of reference signals has simply been pushed up a level. The range of explanation for immediate causes, however, has been considerably extended.

This hierarchical analysis of behavior can now be continued indefinitely, the only restriction on the number of levels being that imposed by experimental findings. The model of the brain's organization (for that is what it is) can be extended accordingly. Each time a new level of control is found, the range of explanations of immediate causes of behavior is extended to cover more kinds of behavior and to span longer periods of time. Each such extension redefines the question of ultimate causes, for each new level of reference signals represents goals of greater generality.

Our going up a level in this analysis is equivalent to our asking what purpose is served by achievement of a given set of lower-order goals: *why* is the man doing that? Why does the soldier raise his arm? In order to point at the helicopter. Why does he point at the helicopter? Perhaps—we would have to verify this guess by test—perhaps to comply with an order. And why comply with an order . . .?

Going down a level is equivalent to asking *how*. How must the man behave in order to point? He must control his arm position. How must he behave in order to control arm position? He must control net muscle-generated forces. And the chain extends further down, to the control systems in the spine which control the effort in whole muscles, as sensed kinesthetically. Each level must be verified by finding a way of disturbing the controlled quantity without affecting lower-order quantities.

Oddly enough, behaviorists may have already found the answer to the ultimate *why* at the top level of this model. Why are the highest-order behavioral goals set where they are set? In order to control certain biologically important variables, which Ashby called critical variables and which I term intrinsic quantities. These are the quantities affected by deprivation and subsequent reinforcements that erase, or at least diminish, the errors caused by deprivation. This makes the highest order of reference levels into those extremely generalized ones that are inherited as the basic conditions for survival. But that takes us to the verge of learning theory, which is beyond the intent of this article. Briefly, I view the process of reorganization itself as the error-driven "output" of a basic inherited control system which is ultimately responsible for the particular structure of an adult's behavioral control system (*8*). For a human being, the "intrinsic reference levels" probably specify far more than mere food or water intake. We cannot arbitrarily rule out any goal at this level—not even goals such as "self-actualization."

Implications for Behaviorism

The most important implication of this analysis for the traditional view of cause and effect in behavior lies in the fact that control systems control what they *sense*, not really what they *do*. In the total absence of disturbances, a control system hardly needs to do anything in order to keep a controlled quantity at a reference level, even a changing reference level. By far the largest portion of output effort is reserved for opposing disturbances.

This is expressed in the approximate relationship, $g(d) \cong -h(o)$. Because of the way negative feedback control systems are organized, the system's output is caused to vary in almost exact opposition to the effects of disturbances—the chief determinant of output is thus the disturbance. If we read "stimulus" for disturbance and "response" for some measure of output, stimulus-response phenomena fall into place within the feedback model.

Stimuli do cause responses. If one knew the controlled quantity associated with a given stimulus-response pair, one would see more regularity in the relationship, not less. In fact one would see an exact quantitative relationship, for the effects of the response on the controlled quantity must come close to canceling the effects of the stimulus on that same quantity, and both these effects are mediated through the environment, where the detailed physical relationships can be seen. That implies, of course, that given knowledge of the controlled quantity one can deduce the form of stimulus-response relationships from physical, not behavioral laws (*9*).

Knowledge of the controlled quantity makes the stimulus-response relationship even clearer by pointing out the right response measure and the right measure of the disturbance, or stimulus. An organism's muscle efforts produce many consequent effects, no one of which can be chosen on the basis of behavioristic principles as being a "better" measure than any other. A stimulus event impinges on an organism and its surroundings in many ways and via many paths, again undistinguishable under the philosophy of behaviorism. Knowledge of the controlled quantity eliminates irrelevant measures of stimulus and response.

Let us consider a rat in a Skinner box. The rat responds to a light by pressing a lever for food. Whatever the immediate controlled quantity may be, it is clearly not affected by the current that flows to the apparatus when the lever is depressed: opening the circuit will not in any way alter the rat's next press of the lever. But holding the rat back with a drag-harness as it moves toward the lever would create immediate forward-pushing efforts, so we would know that the rat's "motion" is close to a controlled quantity. We would of course try to do better than that.

Even though the current to the experimental apparatus does affect the appearance of food, which is quite likely to be a controlled quantity (q_c), the current is still not a controlled quantity, for we could leave the circuit open and actuate the food dispenser in a different way, and the rat would still do nothing in opposition, nothing to restore the current. There is no need to assume what is controlled except as a starting hypothesis, and this method can disprove wrong hypotheses.

The irrelevance of some stimulus measures is common knowledge; rats, for instance, have been found to respond quite well to a burned-out stimulus light, provided that the actuating relay still clicked loudly enough.

Systematic experimental definition of controlled quantities will eliminate irrelevant side effects of stimuli and responses from consideration. But it will also negate the significance of most stimulus-response laws, for once a controlled quantity has been identified reasonably well, a whole family of stimulus-response laws becomes trivially predictable. Once it is known why a given response follows a given stimulus, further examples become redundant. Knowing why means knowing what is being controlled, and knowing the reference level.

When a controlled quantity is found, variability of behavior is drastically lowered, simply because one no longer considers irrelevant details. The remaining variability is lowered even further as one explores the hierarchy of controlled quantities. If all we observed about the soldier in the example were his force outputs, we would have to fall back on statistics to predict them. If we then understood that the soldier was using these outputs to control arm position we could find many cases in which there would be scarcely any variability; applying the correct stimuli (forces) would result in quantitatively predictable force outputs. There would still be many unpredicted changes, but a good fraction of those would become precisely predictable if we understood that the soldier was using arm position in order to point at a specific moving object. Of course as we push toward higher and higher orders of control organization we will find more complex systems employing many lower-order systems at once so that prediction depends on our determining which of several apparently equivalent subsystems will be employed. In principle, however, we can become as thoroughly acquainted with one individual's structure of controlled quantities as we please, if cooperation continues to satisfy his higher-order goals.

Control systems, or organisms, control what they sense. The application of a disturbing stimulus does not affect for long what matters to the organism at the same level as the disturbance, because the organism will alter its lower-order goals in such a way as to cancel the effects of the disturbance. If a position disturbance is applied, the organism will alter its force goals and prevent disturbance of position. If a relative position disturbance (movement of the helicopter) is applied, the organism will alter its absolute position goals and prevent disturbance of relative position.

In this way the system continues to oppose disturbances, making adjustments at every level in the hierarchy of control. The organism will not let *you* (the experimenter) alter what it senses (if it can prevent it), but it will without hesitation alter the very same quantity itself in order to prevent the experimenter's disturbing a higher-order controlled quantity. Hence the well-known perversity of experimental subjects!

It is this hierarchical character of control systems that makes it seem that organisms value self-determinism. And that is not only appearance: organisms are self-determined in terms of inner control of what they sense, at every level of organization except the highest level.

Only overwhelming force or insuperable obstacles can cause an organism to give up control of what it senses, and that is true at every level. In order to achieve ultimate control over behavior, one must obtain the power to deprive the organism of something its genes tell it it must have, and make restoration contingent on the organism's setting particular goals in the hierarchy of learned systems, or even on acquiring new control systems. But one attempts that at risk. Human beings are more prone to learn how to circumvent arbitrary deprivation than they are to knuckle under and do what someone else demands in order to correct intrinsic error. In the sequence deprive, reward, deprive, reward . . ., one person may see the reward as terminating deprivation, but that is only a matter of perceptual grouping. Another person may learn that reward leads to deprivation, and take appropriate action against the cause of deprivation. Pigeons in Skinner boxes, of course, do not have that option.

Summary

Consistent behavior patterns are created by variable acts, and generally repeat only because detailed acts change. The accepted explanation of this paradox, that "cues" cause the changes, is irrelevant; it is unsupported by evidence, and incapable of dealing with novel situations.

The apparent purposefulness of variations of behavioral acts can be accepted as fact in the framework of a control-system model of behavior. A control system, properly organized for its environment, will produce whatever output is required in order to achieve a constant sensed result, even in the presence of unpredictable disturbances. A control-system model of the brain provides a physical explanation for the existence of goals or purposes, and shows that behavior is the control of input, not output.

A systematic investigation of controlled quantities can reveal an organism's structure of control systems. The structure is hierarchical, in that some quantities are controlled as the means for controlling higher-order quantities. The output of a higher-order system is not a muscle force, but a reference level (variable) for a lower-order controlled quantity. The highest-order reference levels are inherited and are associated with the meta-behavior termed reorganization.

When controlled quantities are discovered, the related stimulus-response laws become trivially predictable. Variability of behavior all but disappears once controlled quantities are known. Behavior itself is seen in terms of this model to be self-determined in a specific and highly significant sense that calls into serious doubt the ultimate feasibility of operant conditioning of human beings by other human beings.

References and Notes

1. L. von Bertalanffy, in *Toward Unification in Psychology*, J. R. Royce, Ed. (Univ. of Toronto Press, Toronto, 1970), p. 40.
2. W. T. Powers, *Behavior: The Control of Perception* (Aldine, Chicago, in press).
3. This distinction is akin to the older distinction beween movement and action, the more recent distinctions between molecular and molar, or proximal and distal aspects of behavior. What I term an *act* is a behavior that is arbitrarily left unanalyzed, while a *result* is defined as an understandable physical consequence of an act. Act and result are relative terms, whereas those they replace are absolute. In some circumstances it may be appropriate to consider a movement as a result, in which case the *acts* would be the tensing of muscles. What is proximal or molecular at one level of analysis may be distal or molar at another level. "Distal achievement," in this feedback theory, becomes *perceptual* achievement, and is multiordinate.
4. W. W. Rozeboom, in *Toward Unification in Psychology*, J. R. Royce, Ed. (Univ. of Toronto Press, Toronto, 1970), p. 141.
5. E. Brunswik, *The Conceptual Framework of Psychology* (Univ. of Chicago Press, Chicago, 1952).
6. N. Wiener, *Cybernetics: Control and Communication in the Animal and the Machine* (Wiley, New York, 1948).
7. W. R. Ashby, *Design for a Brain* (Wiley, New York, 1952).
8. W. Powers, R. K. Clark, R. L. MacFarland, "A General Feedback Theory of Human Behavior," in *General Systems—Yearbook of the Society for General Systems Research 1960*, L. von Bertalanffy, Ed. (Society for General Systems Research, Ann Arbor, Mich., 1960), pp. 63–83.
9. W. T. Powers, *Behav. Sci.* **16**, 588 (1971).
10. This article is adapted from a series of lectures given at a faculty seminar on "Foundations of Science," held at Northwestern University, 1971.

NEWS AND COMMENT

NIH Training Grants: Going, Going, Gone?

In reality, there is no such entity as the NIH training program. . . .—From a 1972 analysis prepared by the office of the director of the National Institutes of Health.

Every January, close to the end of the month, the President sends his budget for the next fiscal year to Congress. And every January, during the weeks before that happens, Washington is caught up with dark rumors about programs whose death warrant will be issued in the budget message. Around town, bootleg copies of pages of the budget pass surreptitiously from hand to hand, becoming a special currency whose value lasts a fortnight or so. It will be worthless by 29 January, when the budget is revealed.

From such documents and from the people who have had a hand in shaping them, or who have tried to, came word a couple of weeks ago that the training and fellowship programs of the National Institutes of Health (NIH) are about to meet their end (*Science*, 19 January). This rumor—and it must be considered that until the budget is finally cast in type—has aroused considerable unhappiness and brought no small measure of confusion to the nation's medical schools and research institutions. No more training grants? Is it true? How can that be? Will we survive? Deans, department chairmen, and young investigators seem to have been repeating these questions to themselves, to Washington officials, and to national journalists as the rumor spread.

Several inquiries by *Science* indicate