6

A Psychological Perspective on Purpose: Organisms as Perceptual Control Systems

The analysis of behavior in all fields of the life sciences has rested on the concept of a simple linear cause-effect chain with the organism in the middle. Control theory shows both why behavior presents that appearance and why that appearance is an illusion. The conceptual change demanded by control theory is thus fundamental; control theory applies not at the frontiers of behavioral research but at the foundations.

—William T. Powers (1989, p. 127)

Two of the three necessary steps toward a thoroughly materialistic model of purposeful behavior have now been described. The first step was Bernard's and Cannon's discovery of self-regulation in the physiological processes controlling internal body conditions such as temperature and sugar level, acidity, and carbon dioxide concentration of the blood. The second was the cybernetic understanding of circular causality as it recognizes the essential role played by the closed loop of action and feedback in control systems designed by engineers and in self-regulating physiological processes and overt behavior of animals and humans.

But something is still missing: we have yet to come to a clear understanding of how *purpose* operates in such systems, including how it can be represented, where it comes from, and how it manages to bring about controlled consequences by varying actions in the face of unpredictable disturbances. In this regard it is noteworthy that in Cannon's influential book *The Wisdom of the Body* the word "purpose" is not even included in the index. And although it is featured prominently in the title of Rosenblueth, Wiener, and Bigelow's seminal 1943 paper, it again is conspicuously absent from the index of Wiener's *Cybernetics* except for its supporting role in referring to the pathological condition known as purpose tremor.

The Purposeful Behavior of a Cruise Control System

To address these crucial issues concerning purpose, we must go beyond our rather mundane toilet tank example and consider a somewhat more complex feedback-control device that will be familiar to many readers who drive cars. This is the cruise control system commonly found on automobiles that automatically maintains a steady speed with no assistance from the driver.

An automobile cruise control system is engaged by first turning it on and then pushing the "set" button after the car has reached the desired speed. This speed, say 65 miles per hour, somehow becomes the system's goal or purpose (we will soon see how), and the system acts to increase or decrease the amount of fuel it delivers to the motor as necessary to maintain it. So if the car begins to climb a hill or a stiff headwind begins to blow, the system will sense a reduction in speed (being equipped with a speedometer that measures the rate of rotation of the wheels) and will provide more fuel to the engine through a mechanical link to the throttle. This will increase the engine's power output so that speed is maintained despite the hill or wind. As the car begins to descend the other side of the hill or the wind subsides, the cruise control system will sense the increasing speed and close the throttle, reducing the amount of fuel delivered to the engine so that again the desired speed is maintained. Because it responds to too-high speeds by reducing the amount of fuel delivered to the motor and to too-low speeds by increasing the flow of fuel, the system can be easily recognized as a negative-feedback-control system, identical in function to Watt's steam engine regulator.

Now that we have seen that a cruise control system automatically maintains a steady speed in spite of varying road conditions, let's take a closer look at its internal functions to see how it manages to accomplish this. Figure 6.1 is an adaptation of Wiener's control system diagrams from *Cybernetics* (Wiener 1961, pp. 112, 114). The three boxes indicate the three essential components of a feedback-control system: *sensor*, *comparator*, and *effector.* In a cruise control system, the sensor is a speedometer that converts the rate of wheel rotation to an electrical signal. The signal provided by the sensor is compared with another signal, here labeled "input," which represents the desired or goal speed of the car.

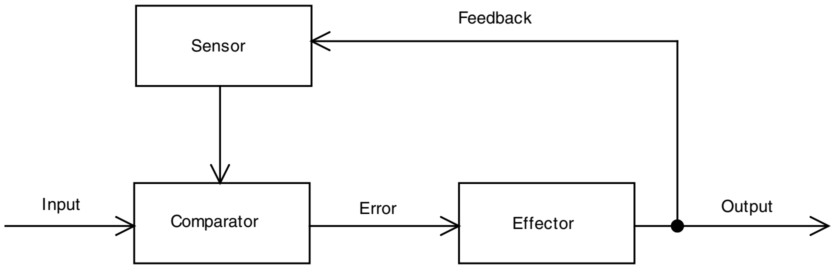

Figure 6.1
Wiener's feedback-control system

The comparator compares the actual speed indicated by the sensor with the desired speed represented by the input signal by subtracting the latter from the former. This comparison results in an error signal, which indicates not only the difference between actual and desired speeds but also the direction in which the actual speed must be changed to match the desired speed. So, for example, if the current speed is 70 miles per hour but the desired speed is 65, subtracting 70 from 65 yields negative 5, indicating that speed has to be reduced by 5 miles per hour. This error signal is then normally amplified and sent to the effector, in this case the throttle that will reduce the amount of fuel provided to the engine until the actual speed matches the desired speed, thereby closing the loop. It should be noted that this is once again a *negative*-feedback system, since the effector increases the amount of fuel sent to the engine if the sensed speed is less than the goal speed, but decreases the delivery of fuel if the measured speed is more than the goal.

It should now be a bit more obvious how the purpose of the system is represented and how it controls the speed of the car. In this diagram, the desired speed, or purpose, is represented by input into the system, which is an electrical signal that indicates the speed of the car when the "set" button is pushed. In this system, as in most engineered feedback-control systems, the desired level of the controlled variable is designed to be manipulable by the human operator. Setting the desired room temperature on a thermostat is another example. In these cases, the goal is provided to the system by a human operator, and is represented in the control system by a signal that is sent to the comparator. The system will then act in a purposeful manner, varying its output as necessary so that the two signals entering the comparator—the signal representing the vehicle's actual speed and the signal representing the desired or goal speed—are the same or very nearly so.

As mentioned earlier, one of the most important insights of the original cyberneticians was the realization that purposefully acting humans and engineered feedback-control systems are alike in certain essential respects. So let us now see how we can use Wiener's diagram to explain the behavior of a human driver controlling the speed of a car the good old-fashioned way, that is, without the assistance of a cruise control system.

We will start again with the sensor. The driver can sense the speed of the car in a number of ways. The speed at which the driver sees road surface approach the car and engine and wind noise (both loudness and pitch) can all be perceived as indicators of speed. But none of these perceptions provides a very precise measure of speed (although I did once know a musician with absolute pitch who claimed she could keep her car at a given speed by keeping the frequency of the engine noise close to a particular musical note!). Fortunately, all cars come equipped with a speedometer that provides the driver with an accurate visual indicator of speed. So the sensor is the driver's eyes and what is sensed is the speed indicated by the speedometer.

But this, of course, is not enough. The driver also has to have a target or goal speed to provide a purpose to his speed-controlling behavior. Let's assume that this goal is the legal speed limit posted as 65 miles per hour. Something within the driver's brain must compare the speedometer reading with the goal speed, subtract the latter from the former, and send the difference (error) to an effector to be acted on. The effector now consists of the muscles of the driver's right leg and foot that act to push down on the accelerator pedal if the perceived speed is less than the goal speed, release the pedal if the perceived speed is above the goal speed, or hold its current position if the perceived and goal speeds match (zero error). Of course, any movement of the accelerator will influence the speed of the car, and this result will be fed back to the speedometer, where the feedback loop from sensor through comparator to effector and back once again to sensor is completed.

So we see that Wiener's basic diagram of a feedback-control system can be readily applied to the purposeful behavior of both machine (cruise control system) and human (driver), even though the physical make-up of the two systems is quite different—electrical wires, sensors, and motors in the former, but living nerves, eyes, and muscles in the latter. However, there is one fundamental difference between machine and driver that seems to have escaped the notice of some early cyberneticians—the origin of what we referred to above as the desired speed or goal speed, but what control systems engineers usually refer to as the *reference level* of the system.

In Wiener's diagram, the reference level is supplied from outside the system and is therefore labeled as an input, since in engineered control systems the reference level can usually be set and manipulated by a human operator. For a cruise control system, the reference level can be changed by pushing the "accelerate" (faster) or "coast" (slower) button until the new desired speed is reached. But there are no "accelerate" or "coast" buttons to be found on the human driver. In fact, the only way to provide input to a human driver is through his senses, as when he sees a speed limit sign or his driving companion asks him to slow down. But there is no guarantee that he will observe such signs or requests. Indeed, our driver may instead decide to speed up when the legal speed limit drops or he is requested to slow down (for example, if traffic decreases or he wishes to annoy his passenger). Or he may slow down when the limit increases or he is requested to speed up (for example, if snow begins to fall or he again wishes to annoy his passenger). So in contrast to the reference level of an engineered control system that is typically provided from the outside by a human user, the reference levels that serve as human goals and purposes seem to originate somewhere inside the brain. If this is the case, it means that the goals of human beings (as well as all other living organisms) are not subject to direct environmental control, as is the case for engineered control systems.

Properties of Engineered and Living Control Systems

We will return shortly to the question of the origin of human reference levels, but only after we first consider some additional ways in which engineered and human control systems are similar. First, although both cruise control systems and human drivers must compensate for many disturbances that would otherwise change the car's speed, they need not perceive the disturbances themselves. The cruise control system has no way of determining whether the road is climbing or descending. Nor can it know if there is a stiff headwind or tailwind, that a heavy trailer was just attached to the car, that a tire is losing air and offering steadily increasing rolling resistance, or that a spark plug has fouled, causing the engine to lose power. All it can sense, and therefore control, is the car's speed. Yet despite its complete ignorance of a multitude of potential and actual disturbing factors, it nonetheless does a good job of maintaining the desired speed. Whereas a human driver may be able to perceive at least some of these disturbances (although wind speed, potentially a very important disturbing factor, is not usually one of these), the performance of the cruise control system suggests that he may not require or use any of this information as long as, like the cruise control system, he pays careful attention to the speedometer reading.

Second, a control system *does not control what it does*. Rather, it *controls what it senses*. The word *control* is used here in its technical sense of maintaining some variable at or near a specified fixed value or pattern of values despite disturbances. Both the cruise control system and human driver can control only what they are able to sense or perceive to be the speed of the vehicle, and they do so by changing output (behavior). Technically speaking, behavioral output is not controlled since the only way the car's speed can be kept close to the reference level speed despite disturbances is by varying the output (that is, changing behavior) as necessary. So we see that a feedback-control system, whether artificial or alive, controls its input (what it senses) and not its output (how it behaves). Consequently, maintaining a constant speed using either a cruise control system or an attentive human driver allows one to predict accurately how long it will take to cover a certain distance. But it will not let one predict how much fuel will be used to drive the distance because fuel consumption is not controlled, varying as it must to compensate for unpredictable disturbances. Since a control system controls what it senses, and since an organism's sensing of the environment is generally referred to as perception in behavioral science, application of control theory to the behavior of living organisms is called *perceptual control theory*. Including the word *perceptual* distinguishes this application of control theory to the behavior of living organisms from the control theory applied by engineers and physicists to artificial (that is, nonliving) control systems.

Third, it is important to realize that whereas a control system's behavior is clearly *influenced* by its environment, it is *not determined* solely by its environment. Rather, its behavior is determined by what it senses (or perceives) of the environment *in comparison with its goal or reference level*. It is worth emphasizing again the crucial difference between nonliving control systems designed by engineers and living ones fashioned by biological evolution: an engineered control system is usually designed so that its reference level can be manipulated by the operator, for example, by pushing the "accelerate" button of the cruise control system or by turning up the room thermostat; however, no such direct manipulation of the reference levels of living control systems is usually possible. We can certainly ask a taxi driver to drive more slowly or tell a teenage child to be home by midnight, but we have no way to guarantee, other than by using overwhelming physical force, that either person will comply with our wishes.

Finally, both engineered and living control systems behave in a clearly *purposeful* manner, varying behavior as necessary in the face of unpredictable disturbances to control some perceived variable, in the same way that William James's frog purposefully sought to reach the surface of the water and Romeo sought to reach Juliet's lips (recall chapter 3). This is not achieved by some future state having present effects, but by having a goal state (reference level), comparing it with current conditions (perception), and acting on the difference (error) until it disappears or is made very small.

Note that nothing mystical, psychic, or spiritual is required for this to occur. It is certainly the case that specifying, perceiving, and controlling something like car speed, temperature, or water level in an engineered control system is orders of magnitude simpler than specifying, perceiving, and controlling something like building a house, writing a book, or having a successful career. Nonetheless, the fact that the former can be achieved in a completely mechanistic, materialist way using fairly simple wires, levers, valves, motors, and sensors suggests that the latter can also be achieved just as mechanistically and materialistically using the much more complex neural networks, sensory equipment, muscles, and limbs of the human body.

The cybernetic ideas of Wiener and his associates were greeted with considerable enthusiasm by several leading scientists around the middle of the twentieth century. Between 1946 and 1953 these ideas became the theme of a series of ten meetings sponsored by the Josiah P. Macy Foundation under the title "Feedback Mechanisms and Circular Causal Systems in Biology and the Social Sciences Meeting" that would later incorporate Wiener's new term in the revised title "Cybernetics: Circular Causal and Feedback Mechanisms in Biological and Social Systems." But although many leading figures in the biological, social, and behavioral sciences as well as prominent philosophers, physicists, and mathematicians attended these meetings,[1] the revolution in behavioral science that appeared so ready to occur never did.

One reason was that many participants of the Macy meetings were more interested in applying cybernetics to issues in information theory and communication than to biological, behavioral, and social sciences. Those who were eager to apply these new ideas to the life sciences often lacked basic technical knowledge concerning the design and operation of negative-feedback-control systems. One such individual, who later became president of the American Society for Cybernetics, stated that purposeful behavior could be explained in the same way that Newton's theory of gravity explained the behavior of a drop a water sliding down an inclined plane, totally disregarding the closed-loop character of purposefully acting systems (reported by Powers 1989, p. 261)!

Another factor in cybernetics' lack of lasting impact on the behavioral and cognitive sciences was the emergence of reliable and powerful digital computers in the middle of the century. The digital computer, with its binary zero-one mode of operation, was better suited to symbolic representations and their logical manipulation as practiced in what has become known as the artificial intelligence (AI) approach to investigating brain, cognition, and behavioral processes. Analog computers, with their use of continuously varying electrical currents that is more amenable to a cybernetic approach to modeling nervous systems, were largely replaced by their digital successors.

Many other reasons could be invoked for cybernetics' failure to revolutionize the behavioral and social sciences (see Powers 1989, pp. 129–136). But a major factor that is still operating to impede acceptance of the basic cybernetic insight is the difficulty replacing the well-entrenched one-way cause-effect (stimulus-response, input-output) model of animate behavior with the more complex cybernetic notion of circular causality. And just such a replacement is needed to account for purposeful behavior in which causes are simultaneously effects and effects are simultaneously causes. It wasn't until the 1960s when another combination of two engineers and a medical researcher began to formulate a general feedback-control theory of human behavior.

Understanding Behavior as the Control of Perception

The Contributions of William T. Powers and His Associates

These three individuals were physicist and electrical engineer William T. Powers, physicist Robert D. Clark, and clinical psychologist Robert L. McFarland, who in the 1950s worked together at the Veterans Administration Research Hospital in Chicago. In 1960 they published a two-part article with the title "A General Feedback Theory of Human Behavior." Thirteen years later in 1973 Powers published the first book that focused exclusively on the application of cybernetic and control-system concepts to animal and human behavior. His book finally made good on the cybernetic promissory note issued by Rosenblueth, Wiener, and Bigelow thirty years earlier.

Powers made three important contributions in extending cybernetic concepts to animal and human behavior. The first was to appreciate fully the revolutionary implications that cybernetics had for behavioral science and to share this insight. As indicated by the title of his book, *Behavior: The Control of Perception*, he recognized that organisms, organized as living networks of negative-feedback-control systems, behaved as they do to control their perceptions. This was a blatant reversal of the then- and still-current mainstream view in behavioral science that perception (of environmental stimuli) controls behavior, either directly (as in behaviorist theory) or through intervening brain-based psychological processes (as in cognitive theory). By turning behavioral theory upside-down, Powers achieved what the preceding ninety-four years of psychological research and theory had not: liberation of psychology from the one-way cause-effect view that sees the behavior of living organisms, like that of inanimate objects, as determined by external forces.

Related to this liberation was Powers's realization—mentioned above in anticipation—that unlike engineered control systems such as thermostats, steam pressure regulators, and cruise control systems, reference levels specifying the goals of living control systems originate from *within* the organism and are neither provided nor directly manipulated by the environment. This raises the question as to what within the organism provides these reference levels and how and why they are provided, leading to Powers's second important contribution: a theory and working model of the *hierarchical* organization of control systems operating within the organism.

A Hierarchy of Perception and Control

To understand this hierarchical organization of control systems and its functioning, it will be useful first to take another look at a simple control system. But this time we will use a more complete diagram inspired by Powers's work that is more appropriate to living control systems than Wiener's diagram.

Figure 6.2 differs in several ways from Wiener's original diagram. First, the reader should take note of the purely cosmetic change from Wiener's horizontal orientation to a vertical one.

Second, a dashed horizontal line divides the control system from its environment. This makes it clear that the system is influenced by the environment only through its sensors (for a living organism this could be any sense organ such as eyes, ears, nose, or touch receptors in the skin), and it acts on the environment only through effectors such as those provided by muscles attached to limbs.

Third, input to the system on the left of Wiener's diagram has been replaced by an entity labeled purpose (6) which provides the reference signal (5) to the control system's comparator (4). Whereas in Wiener's diagram it appeared as if the reference signal came from outside the control system, here its source is clearly within the organism itself. We will return shortly to this important component labeled purpose (6) when we consider the hierarchical organization of living control systems.

Finally, three additional components have been added to the bottom environmental side of the diagram. Controlled variable (1) refers to the particular physical aspect of the environment that the organism is controlling. This can be anything that the organism can see, hear, smell, feel, or otherwise sense. In our example of maintaining driving speed, this envi-

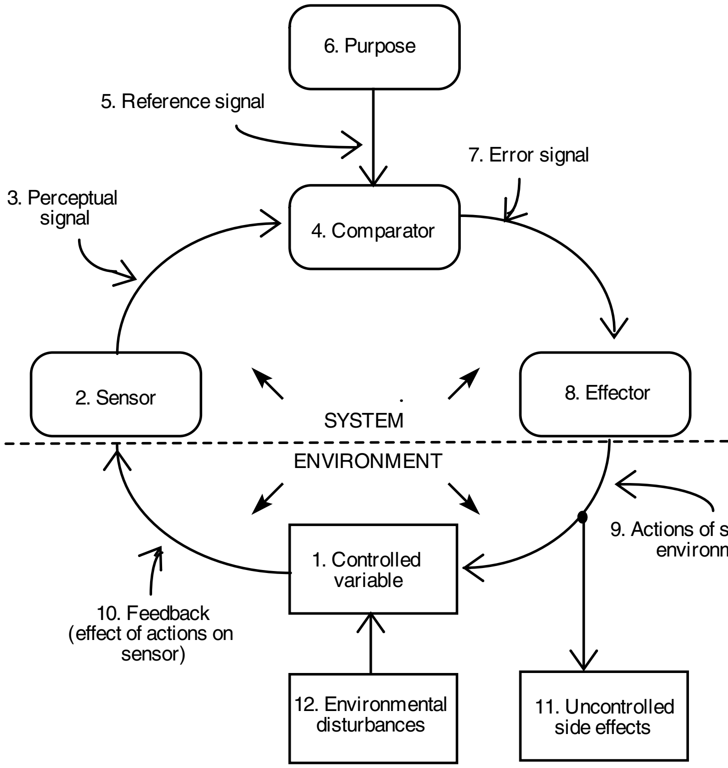

Figure 6.2
Elementary control system

ronmental variable is the position of the needle on the speedometer that the driver must be able to see in order to control the car's speed.

The box on the lower left, environmental disturbances (12), represents all the factors that influence the controlled variable other than actions of the control system itself. In our driving example these disturbances are factors such as wind speed and its direction, and the slope and condition of the road. These are influences for which the driver must compensate so that the car's speed remains under control.

The last addition is the box on the lower right that is labeled uncontrolled side effects (11). This box shows that the actions of a control system, whether engineered or living, will almost certainly have effects on its environment *other* than the desired effect on the controlled variable. Thus, delivering more fuel to the engine while climbing a hill will have effects beyond that of maintaining the speedometer needle at 65 mph. These effects include greater engine noise and vibration, increased use of fuel, higher engine temperature, and faster flow of emissions from the exhaust pipe. These are all unintended effects of maintaining the car's speed, and we will see later how the distinction between intended (purposeful) and unintended (nonpurposeful) consequences of an organism's behavior is crucial for understanding what a living organism is really doing.

Now that we have a more complete diagram showing what is involved in purposeful behavior, let's take a trip around the closed loop it illustrates to ensure that the functions of all its components, labels, and connections are clear. Staying with the example of a human driver maintaining a constant automobile speed of 65 mph, we will start at the controlled variable (1), which is a reading of 65 mph on the speedometer. But as this is an aspect of the driver's external environment (note that it is in the environment half of the diagram), it must be sensed by the driver to be controlled by him. This is done with his light sensor (2), or eyes. (Obviously, if the controlled variable were a sound, taste, smell, feeling, or some combination of these, other sensory systems would be involved.) The driver's visual system converts the speedometer reading into a perceptual signal (3) that is then provided to the comparator (4) that compares this signal with the reference signal (5) of 65 mph provided by the system's purpose (6). The difference between these two signals (3 and 5) constitutes the error signal (7) that causes the effector (8) to act, which in this case is the driver's foot acting on the accelerator pedal. The action of depressing or releasing the pedal (9) influences the driver's environment in many ways. The intended effect of the behavior is its influence on the car's speed and consequently on the driver's perception of the speedometer reading. This effect of behavior on perception through the system's environment is what is referred to as feedback (10). It is this feedback link from actions through the environment back to sensor that completes the loop from controlled variable (1) to sensor (2) to comparator (4) to effector (8) back to controlled variable (1). The box labeled environmental disturbances (10) represents all of the influences on the car's speed that must be compensated for by the driver. Finally, uncontrolled side effects (11) refer to all the unintended consequences of the driver's manipulation of the accelerator pedal (for example, engine and wind noise, fuel consumption, air pressure on the windshield, and engine and tire temperature).

We are now ready to consider where the all-important reference signal (5) comes from. And important it is, since changing this signal from 65 mph to 55 mph will result in an error if the car had been traveling at the previous goal speed of 65 mph, causing the driver to slow down to and maintain this lower speed. Similarly, increasing the reference signal to, say, 80 mph will cause an error in the opposite direction, leading the driver to accelerate to and maintain the higher speed, perhaps even resulting in a speeding ticket (which is probably one good reason why the reference signal will probably not be increased to 80 mph). Since this reference signal representing the control system's goal does not come from the environment (notice how figure 6.2 shows no connection from the environment to the reference signal), it must be provided as the output from some other component of the nervous system. This other component is a higher-level control system that, instead of sending its output to muscles, sends it to the comparator of a lower-level control system.

Powers hypothesized that the nervous systems of animals and humans are made up of many networks of control systems with the basic hierarchical arrangement shown in figure 6.3 whereby higher-level systems send their outputs as reference signals (and thereby constitute higher-level goals) to the comparators of lower-level ones (note that to save space in figure 6.3 comparators are indicated by the letter C, sensors by I for input, and effectors by O for output). For humans, Powers proposed eleven levels of perception. And since each higher-level control system must be able to sense what is happening in the control systems below it, the human control-system hierarchy also requires eleven levels of perception, with higher-level perceptions being made up of weighted combinations of lower-level ones.

Although combining many basic control systems in this hierarchical fashion adds much complexity (and capability) to the overall network, it should be kept in mind that each elementary control system compares its perceptual signal with its reference signal and acts on any difference to reduce it to close to zero. But instead of sending its output to a muscle or group of muscles to act on some aspect of the environment, a higher-level system sends its output to one or more lower-level control systems where it acts as a reference signal for the lower-level systems.

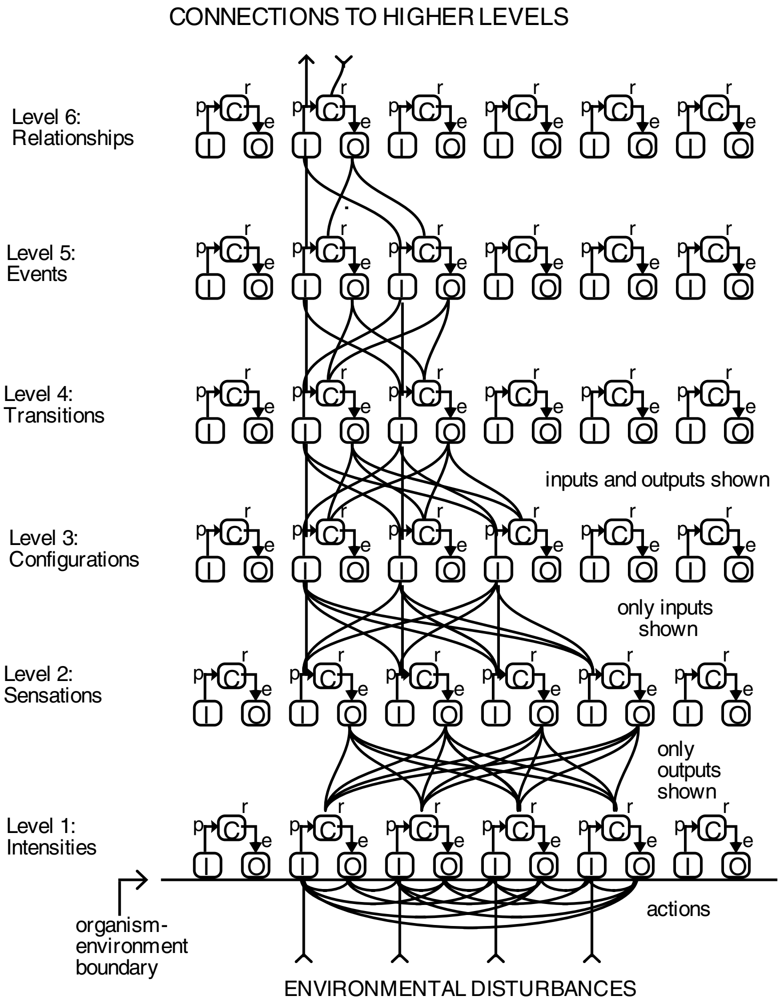

Figure 6.3
A hierarchical network of control systems

This model of the nervous system makes certain predictions about behavior, some of which can be easily demonstrated. But we will save this for a bit later in this chapter where several interesting demonstrations of perceptual control will be described. Instead, let us now consider how Powers's proposed organization provides a new perspective on the physiological control of an organism's inner environment as studied by Bernard and Cannon.

It will be recalled that Bernard wrote of the "constancy of the internal environment" and Cannon introduced the term homeostasis to describe the process by which the body maintains constant internal conditions in spite of the disturbances to which it is continually subjected. But it turns out that at least some of these internal conditions are not so constant after all, and vary in functional ways.

Human body temperature is a particularly interesting example. It is normally maintained close to 98.6° F regardless of ambient air temperature. But we have all experienced fevers during which body temperature increases to 100° or even 102° or 103° F. It used to be thought that these higher temperatures were the harmful effects of bacterial or viral infections. Research has shown, however, that the elevated body temperature characteristic of fever is actually an adaptation in that it helps the immune system eliminate harmful microorganisms. This is accomplished by setting a higher reference level (often called a set point by physiologists) that, like 98.6° F, is also defended against disturbances.

Consider the stages of a typical fever. First, your body temperature begins to rise. But even though it may already be higher than normal, you feel cold and may shiver and put on additional clothing or blankets. This is an indication that the reference level for body temperature has been reset to a higher setting by a higher-level control system. Until your body reaches this new temperature goal you feel cold despite the fact that your body may already be warmer than normal. When your temperature attains the new reference level, you are more comfortable but you feel very warm to anyone who touches you. Finally, your fever "breaks," which means that your reference level for body temperature has been reset once again to its normal temperature of close to 98.6° F. But since it takes a while for your body to cool down to the reference level of the new target temperature, you feel very warm during this time and may perspire profusely until your body temperature once again matches its normal reference level.

While it is not yet completely clear how the reference level for body temperature is manipulated, it is clear that homeostasis is not the best word to describe a control process that involves a changing reference level. Consequently, physiologist Nicholas Mrosovsky (1990) used the term rheostasis[2] to describe such changing reference levels, and he described many such varying reference levels, including those involved in body weight, calcium stores, blood acidity, blood gases, and blood pressure.

Control of body temperature holds further interest since for humans and many other animals it may involve voluntary overt behavior in addition to involuntary internal physiological processes. Shivering and constriction of blood vessels close to the skin are two automatic physiological responses designed to raise body temperature by generating and retaining heat. But a human may also act on the external environment to raise body temperature, as when a person puts on a sweater, adds another log to the fire, prepares and consumes a hot drink, turns up the room thermostat, or adds insulation to the attic. Behavioral means of regulating body temperature are particularly noticeable in cold-blooded animals (technically *ectotherms*) that have no internal physiological means of controlling body temperature. Lizards climb up the sides of rocks and walls in the early morning to catch the first warming rays of the sun. The desert iguana will move closer to a source of heat (such as an electric lamp in laboratory conditions) when infected with a pathogen, thus producing a reptilian version of fever that facilitates elimination of disease-causing microorganisms (see Mrosovsky 1990, p. 77).

Humans can control many variables that are much more complex than body temperature. Imagine for example that Mary, living in San Francisco, learns that her son has become ill in New York City and is being cared for in a hospital there. It is very likely that this situation would be disturbing to Mary in many ways and she would likely feel compelled to make arrangements to visit her son during his illness. We could invoke all kinds of reasons for why Mary is going off to New York, such as love for her son, concern for his well-being, or even that she was rewarded in some way for previous visits she made either to her son or to other individuals in similar situations. But another way of understanding Mary's actions is that she sees herself as a good and loving mother, and not visiting her son during his illness would constitute a serious disturbance to her self-perception.

This is surely a much more complex variable than body temperature or driving speed, but the basic principles of perceptual control are still applicable. For Mary to control her perception of herself as a good mother, she will have to manipulate many lower-level reference levels and control the many perceptions they specify. This is just another way of saying that she will have to accomplish many subgoals to accomplish her higher-level goal of visiting her son. To go from San Francisco to New York, she will have to obtain an airline ticket. To obtain her ticket, she must telephone an airline or travel agent. This involves pushing buttons on her telephone, accomplished by manipulating the tension of her arm muscles in a certain pattern. Only if all these (and many other) lower-level perceptual-control systems are successful in achieving their goals (each subject to unpredictable disturbances) will Mary be able to visit her son and thereby control her perception of herself as a good mother. Doing so, however, will likely cause disturbances to other goals she has, such as those related to her family and work in San Francisco. Thus goals can be related to each other within the same hierarchy as lower-level and higher-level, but can also be situated in different hierarchies, creating the possibility of someone being "of two minds" with accompanying stress and conflict.

The What, Why, and How of Behavior

Powers's model of a hierarchy of perceptual control systems provides a new way of understanding the what, how, and why of animal and human behavior and how this understanding is very different from views provided by other psychological theories.

We will first consider the what of behavior. When behaviorism came to dominate American psychology at the beginning of the twentieth century, one of its major goals was to make psychology a "real" science like physics, and *objective* measurement of behavior became an essential part of its methodology. The number of seconds taken by a rat to run through a maze, the rate at which a pigeon pecked at a key, and the number of times a child disrupted his class during a day at school are examples of behaviorists' objective measurement of behavior. But whereas many aspects of an organism's behavior can be measured by such apparently objective means, such an approach ultimately fails to be either objective or useful. This is because every behavior has very many consequences, and all that a behavioral scientist can ever do is describe one or more *subjectively selected consequences*.

Take the example of Mr. Smith walking down the street. By mentioning walking, I already described one of the consequences of his behavior, namely, that his legs are moving in such a way as to propel him over the ground. I could conceivably obtain more quantitative data about his behavior, such as the frequency of his gait, the speed of his travel, or the force with which he pushes his feet against the ground. But he is also doing many other things that I might have described. He may be out for exercise, calming himself after an argument with his wife, breaking in a new pair of shoes, or going to buy a newspaper at the corner store. He is probably also breathing, perspiring, and even slowly wearing out the soles of his shoes. These are all possible consequences of his behavior, but it is not at all obvious from simply observing Mr. Smith walk down the street which of these descriptions, if any, provides the best answer to the question, what is he doing?

So how does one provide an objective account of behavior when there are so many possible behavioral consequences from which to choose? Figure 6.2 provides a clue. Note that when a control system acts on its environment it has two major types of behavioral consequences. One is that some aspect of the environment, what we called the controlled variable (1), is affected. But many "uncontrolled side effects" (11) are also brought about. Objective observation and measurement do not themselves tell us which of the many effects that one's behavior has on the environment is being controlled—that is, which is the one for which there is a reference level and therefore matters to the individual.

A perceptual control system analysis informs us that one or more of these behavioral consequences matter to the behaving system, and others do not. But how do we find which consequences are being controlled by the individual's behavior and which are unintended side effects? Fortunately, the nature of perceptual control is such that it may be quite easy to find out which is the controlled variable because disturbances to this variable will be resisted whereas disturbances to uncontrolled aspects of the environment will not be resisted. This method of finding out what a particular behavior is intended to accomplish is called the test of the controlled variable by Powers, or more simply, the test.

Let us consider how we might apply the test to Mr. Smith. If we guess that he is out for exercise we might offer him a ride to wherever he is going. His refusal to accept would be consistent with the hypothesis, since a car ride would disturb his goal of getting exercise; but if he accepted, the hypothesis would not look good. If we suspected that he is out to buy a newspaper, we might tell him that the corner store is out of newspapers but the vending machine in the other direction still has some and then observe his actions. A change of heading toward the vending machine would be consistent with the newspaper hypothesis and no change of direction would be evidence against it.

In the case of human behavior, we might save ourselves considerable trouble by simply asking what someone is doing, or more accurately, what he or she is attempting to achieve by his or her actions. But although we may obtain useful information in this way, we have no guarantee that it will be accurate, particularly if the individual has some reason to conceal the real motives for his or her actions or is not conscious of them. And asking is not an option when dealing with very young children or animals.

So we see that perceptual control theory provides a new approach to understanding the what of behavior. Because an action on the environment is initiated when there is a difference (error) between a goal (as represented by a reference signal) and one's current perception, a useful answer to what one is doing is the intended consequence of the behaving organism. Jack may knock over a glass of wine into the lap of his dining companion while reaching for the salt, but a wine-stained skirt was not the intended consequence of his behavior, only the rather unfortunate unintended side effect of the combination of a reference signal for more salt on his steak and the location of the salt shaker behind his glass of wine. The goal-based analysis of behavior provided by perceptual control theory not only provides a new approach but in so doing provides, by the test, a scientific method for distinguishing between the intended (purposeful) and unintended (accidental) consequences of behavior, a distinction that is not even considered meaningful in the objective behaviorist approach. Indeed, the key to understanding behavior as the purposeful control of perception is to attempt to perceive the world from the perspective of the behaving organism. In this important sense, behavior is best understood from a subjective viewpoint, not an objective one.

From a control theory perspective, the answer to the question concerning the why of behavior partially overlaps with the answer to the what. To return to our example of Mr. Smith's walk, knowing what he is doing in terms of his goals (say buying a newspaper) is also to answer why he is walking down the street. But as every parent of a young, inquisitive child knows, one can always continue the why game to the next level and ask why he is getting a newspaper.

To answer this question we must make use of the hierarchy of control systems as shown in figure 6.3. As can be seen in this diagram, comparators receive their reference levels (goals) from the output of higher-level control systems. So obtaining the newspaper is a subgoal on the path to satisfying some higher-level goal, one specified in the reference signal to a higher-level perceptual control system. This higher-level goal could be to check the closing stock market prices. And why is Mr. Smith interested in the closing stock quotes? This brings us up one more notch to a yet higher-level perceptual control system that has as its goal the accumulation of wealth. Why accumulate wealth? Perhaps to be able to retire comfortably at age sixty. If, like the perpetually inquisitive child, we keep on asking why, we will eventually run out of reasonable higher-level goals and be tempted to answer with a simple unadorned "because." But the important point for the present discussion is not to provide an accurate list of higher-level goals for this particular example but rather to show that such why questions can in principle be answered by discovering what the next higher-level control system is controlling, and understanding all goals (except perhaps the one or ones at the very top of the hierarchy—more on that later) as being in the service of still higher-level goals.

The final question about behavior concerns how, and once again the hierarchy of control systems suggests an approach. Just as the why question can be answered by finding the reference level of the next-higher control system, the how question can be addressed by considering the reference levels of lower-level control systems. This is because higher-level goals typically require the control of many lower-level perceptual variables, and higher-level systems control their perceptions by manipulating reference levels they send as outputs to lower-level systems. If Mr. Smith discovers in the newspaper that he is not accumulating wealth according to his plan, he will have to modify certain lower-level goals so that, say, he will change his portfolio from 60 percent bonds and 40 percent stocks to 60 percent stocks and 40 percent bonds. Or, more drastically, he may have to modify his plans, postponing retirement from age sixty to sixty-five to ensure that he will have sufficient funds to retire in comfort.

We can now appreciate that answering a what question about behavior is actually more complicated than first suggested whenever we are dealing with a hierarchy of control systems. This is because the control of a variable such as buying a newspaper involves simultaneous control of many lower-level perceptions (such as reaching the store, taking the newspaper off the shelf, and putting a certain quantity of money on the counter). Yet buying a newspaper is itself a lower-level goal from the perspective of the higher-level goal that has set it, such as checking one's investments or preparing for retirement.

So it turns out that there is usually no one simple answer to a what question concerning behavior (e.g., what is he doing?) but rather the answer must be a description of a set of interrelated goals, some of which may be consciously accessible to the individual (if human) but others not necessarily so. Mr. Smith may be consciously aware of his goal to buy a newspaper, but he is certainly not consciously aware of the complex pattern of perceptual control that is involved in walking down the street (so complex, in fact, that no robot has mastered the bipedal gait). He may not even be conscious at the moment of his goal to retire at age sixty. The test, however, can still in principle be applied to any of these controlled variables, and answers to why questions of behavior can be answered only by moving up the hierarchy, whereas answers to how questions can be addressed only by moving down.

Demonstrations of Perceptual Control

We now come to Powers's third and final (as least as presented here) major contribution. Many behavioral scientists have produced block diagrams of their theories of behavior and perception of the types shown in figures 6.2 and 6.3, as well as verbal arguments to go along with them. But Powers took an important step beyond diagrams and words in producing several convincing demonstrations of the phenomenon of perceptual control and simulations of control-theory models of behavior. These models and demonstrations also inspired several other researchers to develop additional working demonstrations. Since they provide a useful hands-on approach to understanding perceptual control, we will explore several of them and see how they exemplify the concepts introduced above.

The Classic Rubber-Band Demonstration

Our first demonstration, developed by Powers (1973, pp. 242–244), only requires for equipment two rubber bands, a coin, a table, and a willing participant. The two rubber bands are knotted together as shown in figure 6.4 and the coin is placed on the table.[3] Seated across from you, your participant puts a finger through one of the two rubber-band loops and you do the same with the other loop. You then ask your participant to keep the knot that joins the two rubber bands centered over the coin while you gradually and repeatedly move your end of the rubber band toward and away from the coin, keeping it taut, but not so taut that it might break.

If your participant understood your request, you will see that the hand he is using to hold his end of the rubber bands mirrors the actions of your own hand. As you pull your end of the rubber bands away from the coin, he pulls in the opposite direction to keep the knot over the coin. And as you move your hand toward the knot, he does the same.

Since the movements of your participant's hand mirror those of yours, a third person observing this demonstration might well conclude that the participant was simply copying your actions with the position of your hand as the stimulus and moving his hand in response. But it is easy to show that this stimulus-response appearance is really just a seductive illusion (referred to by Powers as the behavioral illusion) and not at all what is really happening. This can be shown by blocking your participant's view of your hand by putting a large book (or magazine or newspaper) between your hand and the knot while taking care not to interfere with your participant's view of the knot and coin. You will then see that even with your hand hidden from your participant's view, he will have no difficulty keeping the knot over the coin in spite of your hand's movements. So contrary to what may appear to be happening, your participant is not responding directly to your hand's movements.

We can get a better idea of what is going on here by referring back to figure 6.2. In this demonstration, the participant is the control system above the horizontal system-environment boundary and you are acting as

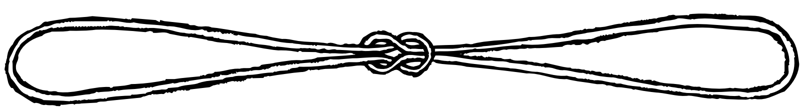

Figure 6.4
Knotted rubber bands

a source of environmental disturbances (12). The participant is able to keep the knot above the coin and achieves this by observing the controlled variable (1) with his eyes serving as sensors (2) that provide a perceptual signal (3) to the comparator (4) that compares the perceived position of the knot with the reference signal (5) provided by his purpose (6). The error signal (7), indicating the discrepancy between the intended perception and actual perception, is sent to the effector (8) that causes muscle contractions to increase or decrease tension on the participant's end of the rubber bands. So whereas your disturbances (12) do result in the participant counteracting them, the diagram makes it clear (as did blocking the participant's view of your hand) that he is responding to disturbances to the position of the knot only because of their effect on the controlled variable (1).

So is it not your movements in themselves but rather their effect on the position of the knot relative to the coin that causes the participant to move his hand. But then isn't it also the case that the participant's actions influence the position of the knot? So what is causing what? Is the position of the knot causing the participant to move his hand, or are his hand movements causing the position of the knot to change? The correct answer, which I hope is obvious by now, is that *both* are happening at the same time: changes in the position of the knot lead to movements of the participant's hand that simultaneously lead to changes in the position of the knot. Here we once again find circular causality operating in a closed loop from perception to action back to perception that defies a one-way, cause-effect analysis.

Computer-Based Demonstrations of Perceptual Control

Although the rubber-band example is a simple and useful demonstration of the phenomenon of perceptual control (and countless variations of it demonstrate other aspects), it does not permit a quantitative analysis of the relationships among disturbance, controlled variable, and action. For this reason, Powers developed a computer demonstration, called Demo 1, that runs on any IBM-compatible computer running DOS (or a DOS window) and that can be obtained on the Internet at *www.uiuc.edu/ph/www/g-cziko/twd*.

Demo 1, the phenomenon of control, provides a computer version of the rubber-band demonstration called a tracking task. The participant's task is to keep a short horizontal line, the cursor, between two target lines (see figure 6.5) by manipulating a computer mouse or trackball, referred to generically as the handle. Instead of pulling on the end of a rubber band, the participant moves a mouse or trackball up and down. Instead of keeping the knot centered over the coin, the participant keeps the cursor horizontally aligned between the two target lines. And instead of you as demonstrator providing disturbances by pulling on your end of the rubber band, disturbances are generated automatically by the computer program.

But now the similarities with the rubber-band demonstration end as the computer demonstration is able to store, display, and analyze relevant data. Figure 6.6 shows a typical run of step F of Demo 1 called compensatory tracking. Time is represented along the horizontal axis, which also serves as an indication of target lines. The positions of the handle, cursor (C. Var), and disturbance are represented by the three lines as they change over time during the course of the 30 or so seconds of the run.

The most striking pattern of this graph is the symmetrical relationship between the disturbance and handle, the latter forming a mirror image of the former. This corresponds to the symmetrical movement of the participant's and your hands in the rubber-band demonstration. This result is even more striking using the computer since we know that the participant never saw the disturbance but only its effect on the cursor while the cursor's position was simultaneously being influenced by the participant's movement of the handle. Yet the disturbance and handle movements are

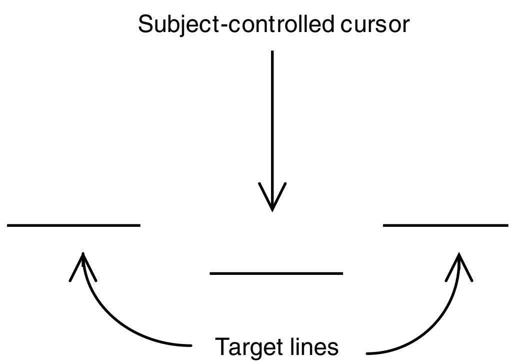

Figure 6.5
Cursor display for Demo 1, compensatory tracking task

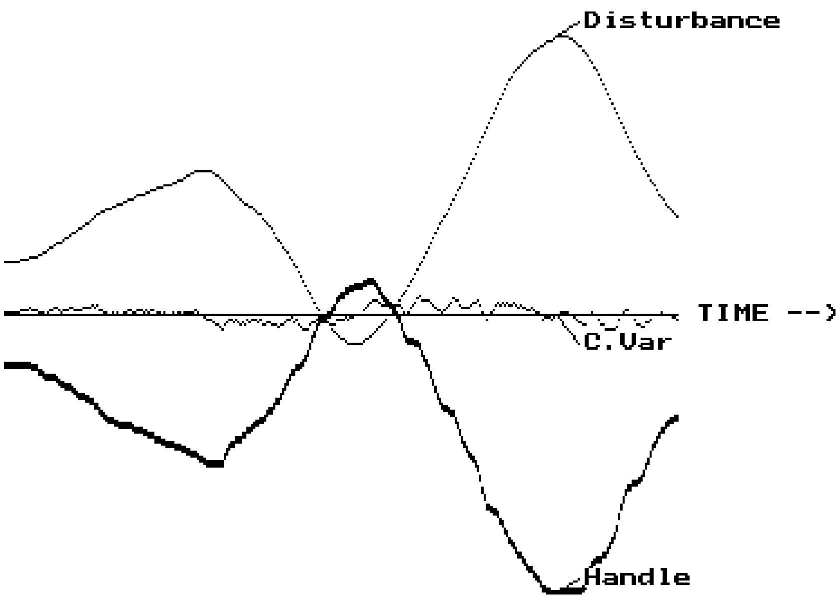

Figure 6.6
Results of Demo 1, compensatory tracking task

very highly corrrelated, with the program indicating for this particular run a correlation coefficient of negative 0.996 between the variables (see box 6.1 for an explanation of correlation coefficients).

Box 6.1
The Correlation Coefficient and Causality

To measure the direction and strength of the relationship between two continuous variables, behavioral scientists use an index called the *correlation coefficient* (usually denoted by the letter *r*), which was developed by Karl Pearson (1857–1936), a British applied mathematician and philosopher of science.

The value of the correlation coefficient varies from -1.00 to 1.00. Its sign (negative or positive) indicates the *direction* of the relationship between two variables, let's call them *x* and *y*. A positive sign indicates a direct relationship, so that as *x* increases *y* also increases and as *x* decreases so does *y*. A negative sign indicates an inverse relationship, so that as *x* increases *y* decreases, and vice versa. As examples, a positive correlation would most likely be found between the height and weight of a group of individuals (since taller people tend to be heavier than shorter people). A negative correlation would likely be found between weight and the number of pull-ups a person can do (since heavier people tend to be able to do fewer pull-ups than lighter individuals).

The *strength* of the relationship between *x* and *y* is indicated by the absolute value of the correlation coefficient, that is, its distance from zero and closeness to either negative one or positive one. Correlation near zero would likely be found between weight and intelligence since heavier people would not be expected to be more or less intelligent than slimmer people. A correlation around 0.7 would likely be found between the height and weight of a group of people, indicating a fairly strong but less than perfect relationship between the variables (it is not perfect since some people will be shorter but heavier than some other people). Perfect (or close to perfect) correlations are not usually found in the behavioral sciences, but can be found in Newtonian physics, such as for the relationship between the mass of an object and the force necessary to accelerate it at a given rate.

It is generally well understood among behavioral scientists that a strong correlation between variables *x* and *y* does not mean *x* is the cause of *y*. First, it may be that *y* is really the cause of *x*. For example, a strong positive correlation may be found for a sample of people between wealth and level of education. Although it may be that wealth leads people to pursue education, it could also be the other way around so that one's education level determines wealth (more highly educated people may earn more money than less-educated individuals). Second, it may be that another variable (or variables) may cause both *x* and *y*, so that wealthy people receive both wealth and educational opportunities from their wealthy parents.

But although a strong correlation does not imply causation, we nonetheless should expect to see a strong correlation between two variables if one of them *is* the cause of the other. For example, if smoking really does cause lung cancer, we should find a strong positive correlation between smoking behavior and incidence of this disease, and we do. This is why in Powers's Demo 1 it is of such interest to find a near-zero correlation between what the participant sees and what he does, since this is strong evidence that what the participant does (response) is *not* directly caused by what he sees (stimulus). Instead, what the participant does *controls* what he sees.

Less striking, at least initially, is the relationship between the cursor (which is what the participant saw) and his handle movements (what he did). The small movements of the cursor above and below the horizontal axis of the graph indicate that the participant was successful in keeping the cursor close to the target position but did not achieve perfect control. And the correlation between the cursor and handle in this run was only 0.179, which is quite close to zero as far as its strength is concerned. But it is this near-zero relationship that is remarkable since we might naively expect what the participant saw to influence what he did. Once we realize, however, that what he *did* also influenced what he *saw* (he was, after all, using his behavior to control his perception of the cursor), the lack of relationship makes more sense. The lesson being, once again, that the circular causality characteristic of perceptual control does not work according to rules of one-way cause-effect phenomena characteristic of the behavior of nonliving objects. In Demo 1 the indication that the participant is actually controlling his perception of the cursor is that there is virtually no measurable one-way relationship between what the participant saw and what he did.

This rather curious characteristic of perceptual control is demonstrated more clearly in step I of Demo 1, intentional vs. accidental effects. In this demonstration, there are now three cursors between the target lines (see figure 6.7). All three are influenced by the participant's movement of the handle, but each is affected by a different disturbance. This would correspond to a task in which three knotted pairs of rubber bands were looped around a participant's finger with three separate disturbers on the other ends. Although the participant's actions move all three cursors, having three disturbance patterns means that only one of the three cursors can be kept between the target lines. The participant's task is to pick one of the three cursors to control, and it is the computer's task to figure out which one it is.

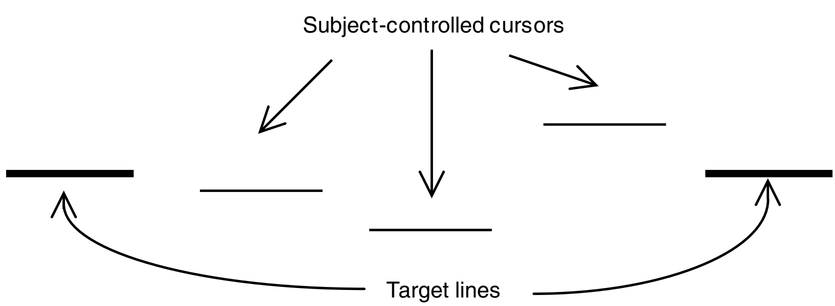

Figure 6.7
Screen display for Demo 1, Step I tracking task

Someone watching the participant do this task would have no difficulty deciding which cursor was being controlled since it is the one that remains close to the target position while the others wander up and down the computer screen. But this is not how the computer makes its decision. It does so by computing correlations between handle movements and all three cursors and picking the cursor that has the *weakest* (closest to zero) correlation with the handle. This counterintuitive approach works very well. In a typical run, correlations between 0.70 and 0.90 are obtained between the handle and the two uncontrolled cursors, while a virtually zero correlation (such as negative or positive 0.10) is obtained between the actually controlled cursor and the participant's handle movements.

An interesting variation of this method of distinguishing the intentional effects of actions from their unintended side effects was developed by psychologist Richard Marken. In his Mind Reading demonstration (developed for Macintosh computers and for Java-enabled Web browsers such as current versions of Netscape Navigator and Internet Explorer) on any computer platform, several numbers (boxes in the Java version) roam the computer screen, each continuously pushed around in two dimensions by its own disturbance. What is seen is not unlike a few scattered fallen leaves being blown around on the ground by its own gusts of wind. But the participant's computer mouse, along with the disturbances, also influences the movements of each number, pushing them all in the same way. By focusing on one number, the participant can control its position on the screen. The participant can decide to keep the chosen number stationary (counteracting its disturbances) while the other numbers continue to be buffeted by their disturbances. In this case it would easy for an observer to find the number being controlled, as it would be the only nearly stationary number on the screen.

But the participant could also decide to move his chosen number in any desired pattern, as in tracing out a circle, square, or figure eight, or even writing his name across the screen with the number. In these cases, since all the numbers will be moving around the screen in irregular patterns, an observer would be hard pressed to tell which one was being controlled by the participant. But the computer only has to find the weakest correlation between the movements of each number and the movements of the participant's mouse to determine which number the participant is intentionally moving. When found, the program indicates the controlled number by highlighting it in boldface. This mind reading of the participant's intentions works no matter what type of pattern the participant imposes on his number, as long as he has an intention concerning where he wants the number to be and varies his behavior to bring about the desired perceptions.

Another program developed by Marken called Find Mind allows the subject to do some mind reading of her own. Now we have numbers (boxes again in the Java version) roaming around the screen as before, but one of them is different from all the others, although this is not at first apparent from watching them move. All the numbers but one have been programmed to move around the screen not "caring" where they roam. If one of these numbers had been programmed to move one inch to the left while a disturbance pushed them all an inch upward, the number would simply move about an inch and a half toward the upper left corner by combining its own movement with that of the disturbance. But one of the numbers represents the actions of a control system with a varying reference signal specifying where it should be at any given instant and the means to counteract disturbances to achieve its goals. As in the previous demonstration all the numbers are influenced by the computer operator's mouse movements, but the one acting as a control system will go where it intends to go (the intention, of course, having been provided in the computer program by the programmer) and will resist disturbances to its movements. By trying successively to keep each number contained in a box at the center of the screen, the user will soon find the one number that has a mind of its own in not "wanting" to be in the box. This one number actually "feels" quite alive in its resistance to the user's mouse-induced disturbances.

These demonstrations were designed to give the user a better understanding of the phenomenon of perceptual control and to show some of its rather surprising characteristics, such as near-zero correlation between perception and action when one's actions are used to control one's perceptions. But Powers and his associates did not stop there. They wanted to show not only that perceptual control is a real phenomenon but that control systems can provide useful working models for animal and human behavior.

Powers's Demo 2, modeling compensatory tracking, leads the user step by step to the construction of a working control system whose behavior in a tracking task is compared with that of the user. In step F, closing the loop, the user sees how a working control system keeps the cursor centered on a target location and how changing the system's reference signal influences the consequences of its behavior. In step J, matching the model to real behavior, the user can compare his behavior to that of the model control system and make adjustments to the model until its behavior closely matches his own. In figure 6.8, the top diagram portrays the computer model's behavior (with plots of cursor, handle, and disturbance provided) and the bottom diagram is that of the human operator. The smaller

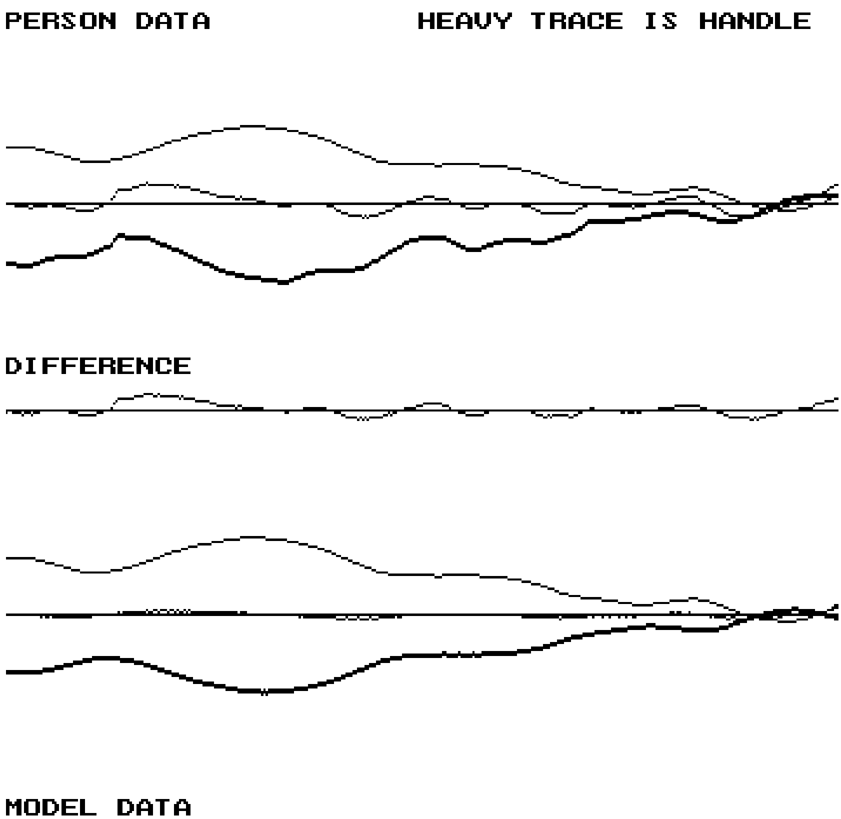

Figure 6.8
Matching person and model data in Demo 2

graphic separating the two shows the difference in their behavior. In the particular case shown, the difference was very small, with the correlation between the control system's behavior and the human's (in this case yours truly) equal to a very strong correlation of 0.986. This near-perfect correlation indicates that the control-system model fits the human's behavior extremely well (it should be noted that correlations stronger than 0.70 are quite rare in the behavioral and social sciences). Thus, Powers's Demo 2 goes well beyond the typical diagram of a psychological theory in that it can be turned into a working model that does what it was designed to do, that is, control some aspect of its environment as a purposeful human performs this same task.

Demonstrating a Hierarchy of Perceptual Control

Powers and his associates also developed a number of demonstrations of the hierarchical organization of human control systems that was described earlier and illustrated in figure 6.3. Recall that in a hierarchy of control systems, higher-level systems send their outputs as reference signals to lower-level systems. In this way the higher-level control systems do not tell the lower-level ones what to do but rather what to perceive as the consequence of their actions. This proposed hierarchical organization has at least two implications. First, it makes some interesting predictions about the performance of certain tasks. Second, it should prove useful in modeling certain types of animal and human behaviors.

Our first demonstration requires a human participant and you as experimenter. First, ask your participant to extend her arm fully toward the front so that her hand is at the same level as her shoulder, and to maintain it in this position. Now you apply disturbances to her extended arm by pushing her hand gently up and down and from side to side. If the participant indeed has the goal of maintaining her arm in this fixed position (as you have asked her to do), she will resist your disturbances, pushing back on your hand with the force required to keep her arm more or less stationary. This is a rather simple feedback-control system of the type shown in figure 6.2, with you acting as the environmental disturbance. You will notice that your participant's control of her arm is not perfect, but she should be able to keep her arm fairly close to her intended position as long as you don't apply too great a force to her hand or make very rapid changes in the force you apply.

Now as your participant maintains her extended arm position, place your own hand above and lightly touching hers and tell her that when given a certain signal she should bring her arm quickly down along her side. The signal will not be a verbal one, however. You will give it by pushing down on her extended hand (remember your hand is already touching hers) when you want her to change the position of her arm. When you provide the signal as described, you will notice a curious reaction from your participant. Instead of quickly bringing her arm down to her side as soon as you push down on it, she will at first resist your push for a fraction of a second. You can do this again and again, and each time this momentary resistance and hesitation will occur. This resistance seems at first rather odd since you are pushing her hand in the direction that she intends to move it. So why does she initially resist your push?

The hierarchical organization of control systems makes it clear why this must happen. By asking your participant to move her hand down when you push on it, you are actually asking her to change her reference signal (goal) for the position of her hand from straight out to down. But the only way this reference signal for hand position can be changed is by the output of the control system above it, the one that supplies the reference signals to the lower control system and is concerned with your participant's higher-level goal of complying with your request. It would be a disturbance to this higher-level system if your participant were to keep her arm and hand extended after you have pushed down on it, and so to correct for this error the higher-level system changes the reference level for the arm-position control system below. But before the higher-level system can perceive the push on her hand, the lower system has already sensed it (since it is lower in the hierarchy) and taken appropriate action to maintain the original position before the reference level can be changed to the new position by the higher-level system. So this momentary resistance and hesitation in bringing her arm down when pushed are exactly what a hierarchical control-system model of behavior predicts.

Many other manual demonstrations of the hierarchical organization could be described (see, for example, Robertson & Powers 1990, p. 21). But we will now move on to another interesting computer program developed by Powers known as Arm 1.

This computer demonstration (which again can be run on any IBM-compatible computer running DOS or in a DOS window) shows both how a hierarchy of control systems could be used to model human pointing behavior and how such a model could be used to create a robot arm. The task for the computer-simulated arm involves bringing its fingertip in contact with the center of a suspended triangular target and maintaining contact while the target is moved anywhere within reach in the three-dimensional space in front of the arm. This may seem to be a rather simple task for a robot arm to accomplish, but it turns out to be quite complicated, as least when pursued in the typical manner of using what the robot sees to compute what it should do. For this one-way cause-effect approach to work, the robot first has to see the target, determine its position in space, convert this position to the angles required at the shoulder and elbow joints for its fingertip to touch the target (this calculation is known in robotics as *reverse kinematics*), and finally calculate the forces required to bring the arm to this position without undershooting or overshooting the target using what is known as *reverse dynamics* (see Bizzi, Mussa-Invaldi, & Giszter 1991 for evidence of the extreme complexity of this feed-forward approach to pointing to a target).

But this pointing behavior is actually quite easy to accomplish using seven simple control systems, with six of them organized into a two-level hierarchy. At the higher level are three visual control systems, each of which sees both the target and the robot arm's fingertip and also has a reference level of zero for the perceived distance between fingertip and target. One of these visual control systems controls horizontal distance between fingertip and target by sending its output as a reference signal to the comparator of a lower-level kinesthetic control system that controls the side-to-side angle of the shoulder joint. The second of the upper-level visual control systems controls the vertical distance between fingertip and target by sending a reference signal to another lower-level system that controls the up-and-down angle of the shoulder joint. And the third upper-level visual control system makes sure that the fingertip is not behind or in front of the target by controlling for zero perceived difference in the distance of the target and fingertip from the eyes by manipulating the reference level sent to the elbow joint. These six simple control systems, plus a separate seventh one that keeps the robot facing the target, are sufficient to keep the simulated robot pointing to the target as the user manipulates the position of the target in simulated three-dimensional space using the keyboard or a mouse (see figure 6.9). Powers's Arm 2 program does the same, but is more realistic (although slower) in that it includes the effects of gravity on the arm, real arm dynamics (related to the physical characteristics of human arms and muscles), and the possibility for the robot to learn to point more effectively over time (Powers 1999).[4]

Demonstrating Social Systems

Social systems composed of interacting purposeful individuals also were modeled using perceptual control theory. Powers, together with sociologists Clark McPhail and Charles Tucker (1992), developed a program called Gather[5] that models the movements of temporary gatherings of individuals (persons or animals). In these simulations, individuals are programmed as control systems that begin their existence at a certain point on the screen and move to satisfy the reference levels they are given for their locations. Each individual's location goal is either a fixed point on the computer screen or a certain proximity to another individual who also has a goal of either a fixed location or distance to another individual. Each individual also has reference levels for avoiding too-close proximity to other individuals and the fixed obstacles that are scattered across the screen. The user can manipulate the number of individuals, their goals, the number of obstacles present, and various other parameters

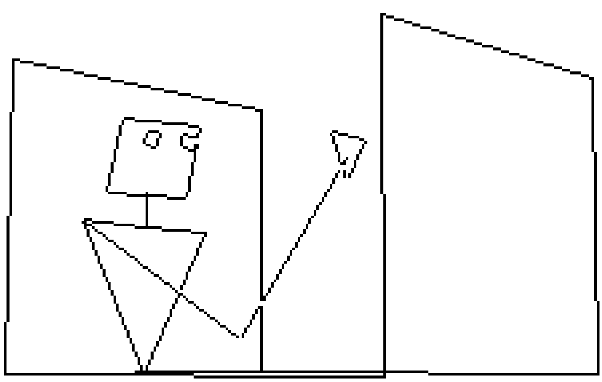

Figure 6.9
Pointing arm simulation

of the individuals' control systems, and see the effects on their collective behavior.

Figure 6.10 is the result of one run of Gather in which one individual, labeled M, moves from the left of the screen to its destination goal in the circle on the right side of the screen. The goals of the four individuals labeled G are not fixed locations but rather the intention to remain close to M without being too close to each other or run into any of the obstacles indicated by small circles.

With the choice of the right control-system parameters, all of the individuals are successful in achieving their goals (as the traces on figure 6.10 indicate) regardless of the distribution of the obstacles they must avoid. Their collective behavior is similar to that of a human mother being followed by her four children across a shopping mall while avoiding other individuals and objects, or a mother goose followed by her four goslings as they waddle from meadow to lake avoiding rocks and trees along the way. It is also of interest to note that the arc formed by the four Gs does not exist as a goal for any of the individuals but rather emerges as an uncontrolled (but reliable) side effect of the outcomes that the Gs are controlling, namely, maintaining a certain distance between themselves and M.[6]

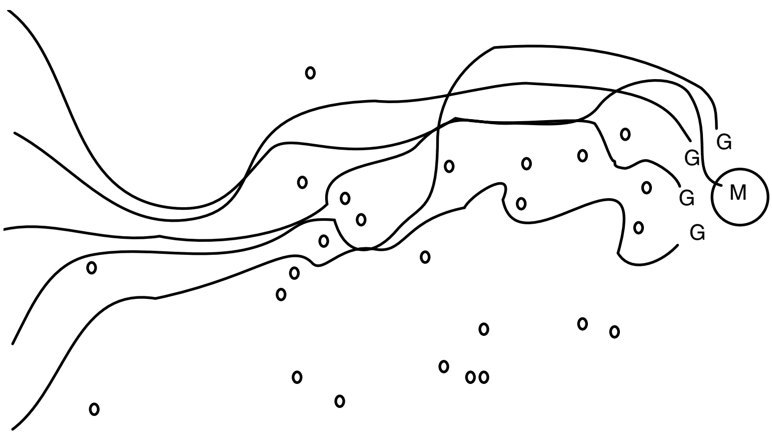

Figure 6.10
Gather simulation of four individuals (G) following another (M) (after McPhail, Powers, & Tucker 1992)

Three Final Demonstrations: Controlling Another Person, “Ballistic” Movements, and the Coin Game

Three final demonstrations, none requiring a computer, are worth describing since each shows another interesting characteristic of perceptual control. The first requires the same knotted rubber bands (see figure 6.4) used in the first demonstration, a table, two coins placed about 10 inches apart on the table, and, of course, our indispensable willing human participant. As in the first demonstration, you and your participant each grasp an end of the two knotted rubber bands, and you ask your participant to keep the knot over the coin that is farther from him. But this time as you watch the position of his hand, you move your hand so that he places his hand over the other coin.

What you have done is controlled the behavior of your participant by “making” him put his hand over the other coin. This control was achieved by knowing what perceptual consequence he was controlling and providing the disturbances that would lead him to put his hand where you wanted it to be. But this control of your participant’s behavior works only as long as he maintains his goal of keeping the knot over the more distant coin and *does not care* (that is, has no higher-level goal or reference signal for) where his hand is located over the table. If either of these conditions no longer holds (your participant either no longer wishes to comply with your request to keep the knot over the one coin, or does not want to keep his hand over the other coin) you will no longer be able to control his behavior without recourse to overwhelming physical force. This indicates a general principle of the control of one person’s behavior by another: Other than using irresistible physical force, an individual can control another individual’s behavior (or more accurately, the *outcome* of his behavior) only by causing disturbances to goals that will elicit the desired behavior, and only if the desired behavior does not disturb the goals of higher-level control systems.

Another example shows the fallacy of the common belief that certain so-called ballistic behaviors take place too quickly for continuous sensory feedback to be involved in their execution. Two such behaviors are hammering and throwing a ball or stone. Neurobiologist William Calvin (1990, p. 239) made just such an argument and proposed it as a factor contributing to the evolution of the human brain:

> . . . ballistic movements [are] quite unlike the ones where an intention and feedback corrections suffice to get the job done: Brief movements have to be carefully planned in advance. Any trial and error has to be done while planning, checking a proposed movement against memory as you "get set," and discarding the plans that don't jibe.

To see if feedback can actually be used in these actions, one has to figure out how to apply a disturbance to the behavior while it is occurring and see if it is resisted to any extent. This can be easily done by attaching an elastic band to your participant's wrist (I use a large loop of rubber about an inch wide cut from an old bicycle inner tube) and have him throw or hammer while you apply a disturbance by pulling on the elastic band after his action has begun (still better would be to use two elastic bands with two disturbers pulling on one, or the other, or both, or neither so that the participant could not anticipate what the disturbance would be).

For throwing I have my participant throw a tennis ball underhand from a distance of about 15 feet against a chalkboard on which a target consisting of a circle of about 1.5 feet in diameter has been drawn. For hammering, I place a coin on a table and let the participant hammer on it with his closed fist (it's technically pounding, not hammering, but much easier on the table if not on the fist). While disturbances applied by the elastic band will likely have some effect on the accuracy of throwing or hammering, the effects are quite small compared with the magnitude of the disturbance. This is something you should also experience as the thrower or pounder with your participant attempting to disturb your actions, since you will experience how you automatically adjust your actions "on the fly" to compensate for the disturbances.

The fact that these disturbances can be corrected after the throwing or hammering action has begun indicates that negative-feedback control *is* involved in these supposedly ballistic behaviors. If they were the result of preplanned motor commands (as Calvin and many others believe), no real-time corrections would be possible at all. The results of these demonstrations are instead consistent with the operation of a hierarchy of control systems in which upper-level systems do not tell lower-level systems what to *do* (that is, provide motor commands) but specify what lower-level systems should *perceive*. The controlled perception is that of a certain sequence of joint angles (known as *proprioception*) that has been associated with the perception of previously successful throwing or pounding and that will itself be adjusted by still higher-level systems depending on the perceived outcome of each trial. It is important to note that a form of associative learning is occurring here. But it is not that of associating a stimulus with a behavior. Rather, it is associating higher-level controlled perceptions with lower-level ones.

The final demonstration is the coin game devised by Powers (1973, pp. 235–236). It shows how difficult it can be to figure out what perception another person is controlling, even when you have the opportunity to make repeated disturbances and guesses.

To play the game you need four coins, a table, and your human participant. Have your participant first arrange the coins in any configuration she wishes (for example, rectangle or square, or even something like two coins closer to each than the other two coins are to them or to each other) and ask her to write down in words on a piece of paper the configuration or condition that she has adopted as her goal. You as experimenter attempt to guess what your participant is controlling by disturbing the coins any way you wish and having the participant say "no error" or correct the error (by moving a coin or coins) that you have created. Once you are fairly certain that you know what the participant is controlling, test your hypothesis by making three moves, each of which you believe will be corrected by the participant, followed by three moves you believe will cause no error. If successful, you then describe what you believe to be your participant's controlled variable (such as, any three coins in a straight line) and compare it with what the participant wrote down.

Playing the coin game will reveal how difficult it can be to determine what the participant is "doing" (actually, what perception she is controlling) even though her actions are completely visible to you, and you can repeatedly disturb the configuration of coins and observe her reaction. Of course, the game is none other than a form of the test for the controlled variable mentioned earlier and provides an illustrative example of how the test can be used to understand another person's behavior.

I hope that I have provided useful descriptions of these demonstrations and what they reveal about the process of perceptual control. Verbal descriptions alone, however, cannot come close to providing the understanding and insights that hands-on experiences with these demonstrations can provide. For this reason, I strongly urge that readers take the time to do at least the rubber-band demonstrations, and that those with access to a personal computer and the Internet obtain and try out the computer demonstrations. Only in this way can one realize that seeing behavior as the control of perception is not just another cute slogan or cliché, and that cybernetic models of perceptual control are more than just boxes and lines on pieces of paper. Rather, perceptual control is a real and easily demonstrated phenomenon that cannot be understood from the traditional one-way cause-effect view of animal and human behavior, and networks of negative-feedback perceptual control systems can be fashioned into working models that behave remarkably like the purposefully behaving animals and humans that they were meant to simulate. Most important, however, is understanding that we now have a basic theory (and model) of animal and human behavior that can explain its purposeful nature in purely materialist and mechanistic terms, but which requires a rejection of the one-way cause-effect view of living behavior.

The Puzzle of the Ultimate Why Question

We have now seen how considering animate behavior as an organism's means to control aspects of its environment provides a new way of addressing questions concerning the what, how, and why of behavior. From this perspective, what questions are addressed by considering the perceptual variable that an organism is controlling, keeping in mind that any given action may have many uncontrolled side effects that are of no concern to the behaving organism, and that the behavioral consequences specified in reference levels need not be static but instead can be continually changing.

How questions are answered by considering the subgoals, or lower-level reference levels, that must be controlled for a higher-level perceptual variable to be controlled. From this perspective, a professional golfer is able to drive her ball onto the green not because her nervous system is able to send a certain fixed sequence of motor commands to her muscles, but because she has learned to control a sequence of lower-level perceptions involving the positions and velocities of her limbs, head, and trunk, as well as the relationship of these kinesthetic and proprioceptive perceptions to the visual perception of the green she is aiming at.

In contrast to behavioral how questions that focus our attention on lower-level control systems and their reference levels, why questions about behavior are addressed by moving up the hierarchy of control systems to find higher-level reference levels (or goals) that determine lower ones. Someone observing my behavior at this moment would notice that I am currently tapping keys on my computer. Why? To make certain letters and words appear on my computer screen (not to make the tapping sound that accompanies each keypress, although objectively my typing is creating noise as well as words). Why make these words appear? Because I want to write and publish a book. Why write and publish a book? Maybe to became famous and make lots of money from royalties (not very likely). Or perhaps so that I can make a lasting contribution to human knowledge (somewhat more likely?). But why bother contributing to human knowledge (I could be outside enjoying this beautiful late spring day rather than sitting in my office in front of a computer)? Good question. As we have noted earlier, as we continue to ask why questions about behavior we usually come to a point at which we no longer can easily imagine what words to put after "because."

But the hierarchy of goals posited by perceptual control theory provides at least a framework for considering answers to why questions. And the answers we find are very different from the ones proposed by one-way cause-effect theories that look for answers not within the organism but rather in the effects that the environment has on the organism. Because we attempt to answer these questions by searching for the next higher-level control system and its reference level, these can be considered the proximate causes of behavior.

But for any theory of behavior to be complete, ultimate causes of behavior must also be considered. We observe a robin pecking in the soft earth during a rainstorm and understand its behavior as a way of getting food in the form of earthworms into its stomach. But why earthworms and not the seeds that the sparrows and finches consume? A male robin pursues a female until she allows him to mount her. From where did this urge to copulate come? We later see the female robin regurgitating her food into the gaping mouths of her newly hatched chicks. But why should she share her hard-earned food with this chorus of seemingly insatiable little beaks?

Similar questions concerning the ultimate reasons for behavior could easily be posed for humans, but answers cannot be found by staying within an individual organism's hierarchical network of perceptual control systems. Instead, we have to consider the process responsible for life itself and its continued evolution.