2

Principles of Feedback Control

This chapter describes a set of basic principles underlying the conceptual analysis presented in the remainder of the book. We describe the principles here abstractly, with examples and illustrations taken mostly from domains other than personality–social psychology. Our goal is to create a clear sense of the nature of particular processes (for a more detailed account see Clark, 1996) without pressing the argument that human behavior embodies them. We move on to that argument in due course.

CYBERNETICS, FEEDBACK, AND CONTROL

Wiener (1948) defined cybernetics as the science of communication and control. *Cybernetics* is one of several terms intertwined with one another – terms such as *control processes, feedback processes*, and *servomechanisms* (or *servos*). These terms have varying origins, they're used preferentially by different people in different lines of work, and they differ in shades of meaning. For our purposes, though, they refer to roughly the same things. Cybernetics is the science of feedback processes; feedback processes involve the control or regulation of certain values within a system (see also Ashby, 1961; Clark, 1996).

Negative Feedback

A negative feedback loop, the basic unit of cybernetic control, is a system of four elements in a particular kind of organization. The elements are an input function, a reference value, a comparator, and an output function (Figure 2.1). An input function is a sensor. It brings information into the loop. In later discussions we'll treat this input function as equivalent to perception. In the abstract, however, it's simply a process by which information of some sort arrives to be used in a particular way.

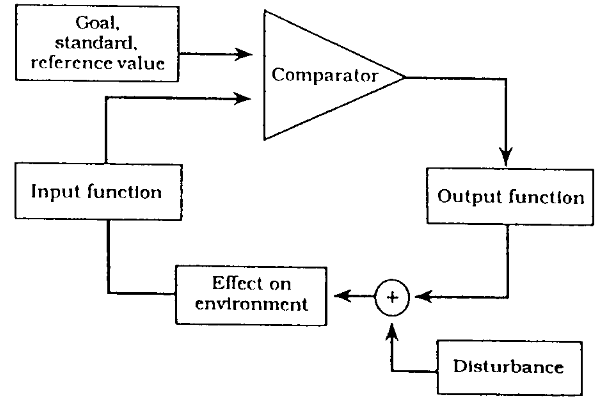

Figure 2.1. Schematic depiction of a feedback loop, the basic unit of cybernetic control. In such a loop a sensed value is compared to a reference value or standard, and adjustments are made in an output function (if necessary) to shift the sensed value in the direction of the standard.

The next step in describing a feedback loop is handled most easily by introducing two elements at the same time: the comparator and the reference value. The comparator is a structure that makes comparisons. Comparison can be made in different ways in different physical systems, and how it takes place is less important than *that* it takes place. The reference value is a source of information other than the information from the input function. As the input enters the system, a comparison occurs between the input and the reference value. This comparison yields one of two outcomes: The values being compared are discriminably different from one another or they're not.

What follows this comparison is an output function (Figure 2.1). In later discussions we'll treat this output function as equivalent to behavior. In the abstract, however, it's more general than that. It's simply an *effector* of some sort – a process that in some way or other has an effect on the system's environment (whatever's external to the system itself). If the comparison fails to find a difference, the output remains whatever it is now (which may be zero or may be some other value). If the comparison finds a discrepancy, the output changes.

These functions in the negative feedback loop were described by Miller, Galanter, and Pribram (1960), in an early statement on feedback control, as what they termed a TOTE unit. TOTE is an acronym for test-operate-test-exit. *Test* is comparison. If comparison reveals a discrepancy, an *operate* (output function) occurs, after which *test* recurs to see

whether the discrepancy still exists. *Exit* indicates that, with no discrepancy, control transfers elsewhere to permit some other activity to occur.

In a negative feedback system, the change in output is aimed at countering any deviation of the input function from the reference value. There are several ways to say this, all of which mean the same thing: The change in output is aimed at reducing the discrepancy between input and reference value, at causing the former to conform to the latter. It's an attempt to create input information that's not discriminable from the standard. It isn't behavior for the sake of behavior, but behavior in the service of creating and maintaining a desired perception.

This last statement has several implications. For one, the value the input function senses depends on more than what happens as output. The output function operates on the system's environment, but so do other forces. Disturbances can change present conditions, either adversely (creating a discrepancy from the reference value) or favorably (closing a discrepancy). In the former case, noting a discrepancy prompts a change in output, as always. In the latter case, however, the result of the disturbance is that *there's no need for an output adjustment*, because the system sees no discrepancy.

Feedback systems are often called "purposive" because it isn't just the pieces of the system that have roles to play. Rather, the system as a whole serves a purpose: to keep sensed values in conformity to the standard that's in place as a reference value. They're called *self-regulatory* systems because they regulate specific qualities via an internal organization. They're referred to as *closed loop* systems because there's an endless cycle among functions, with output having an impact on subsequent input. They're also referred to as *control systems* for at least two reasons. First, there's an interdependence of processes in the loop, such that the result of each function partially determines (controls) what happens in the next function. Second, the overall purpose of the loop is to determine (control) the quality of the sensed input.

A brief comment on the word *control*: This is a word with several meanings in the vocabulary of psychology. One meaning that's particularly likely to become confused with the meaning intended through most of this book is control as personal causal responsibility for events. We don't mean that here. Here, control refers to the process of maintaining conformity of a sensed input to a reference value, regardless of the source of the influence that does so. That is, it doesn't matter to the feedback loop whether conformity is created by its own output function or by a disturbance from outside. Sometimes the loop's actions are

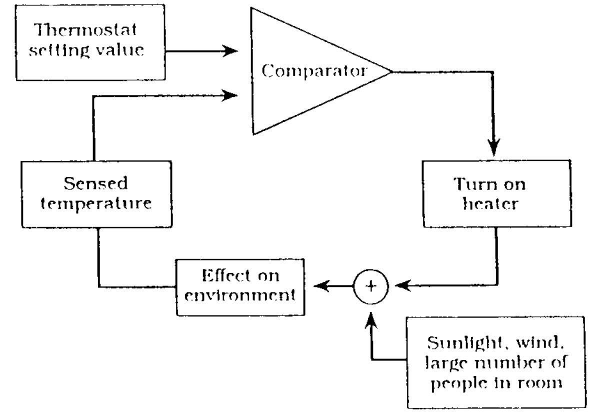

Figure 2.2. The elements of the feedback loop are often illustrated by the functions manifest in a room thermostat and furnace/air conditioner, with sensing of air temperature as the input function, a thermostat setting as the reference value, the turning on and off of the heater as the output function, and a variety of environmental influences as sources of disturbances.

critical, but sometimes a loop creates the desired outcome by *not* acting. Thus, control in this context has more to do with ensuring a result than with having causal primacy.

An Example: The Ubiquitous Thermostat

Let's illustrate some points we've made in abstract terms with a couple of examples from the behavior of physical systems. A commonly used example is the thermostat. The simplest point this example makes is that feedback processes can occur in diverse physical systems. What matters isn't the nature of the physical elements, but the logical relationships among the functions in which the elements engage (see Braitenberg, 1984).

Figure 2.2 shows the schematic displayed earlier in Figure 2.1, relabeled so the functions are identified with the elements of a thermostat and ancillary devices (e.g., heat pump). The system has an input function, continuously sampling current air temperature. This input information goes to the device that compares the sensed value to the thermostat's setting. As long as the two values aren't discernibly different, nothing else happens. If the comparator detects a difference between values, it sends a message that turns on the heater and air transport system, which begins to dump warm air (or cold, depending on the application) into

the room. Eventually enough warm air arrives and mixes with the air in the room that the thermostat can no longer tell the difference between the room temperature and its setting, and its request for activity from the heater ceases.

This example illustrates the elements of the loop, and it illustrates the purposive character of the system as a whole: It functions to keep the room's temperature in reasonably close conformity to the value of its setting. The example also illustrates other points made earlier. This system doesn't act by turning the output function on and off for certain specific periods of time, because it isn't the action of the compressor that's being controlled, but the temperature of the air. If a hot afternoon sun temporarily warms the room on a chilly fall afternoon, or if a crowd of people who are standing in the room radiate excess body heat, the heater doesn't have to heat, and the call to heat doesn't come. The room temperature is controlled by disturbances, rather than by the output function.

This example also illustrates that a not very clearly specified process intervenes between output and input functions, labeled only "effect on environment" in Figures 2.1 and 2.2. The output function has an indirect effect on the input function, rather than a direct one. The output changes reality in some fashion, but the path by which this change influences subsequent input may be quite circuitous. In the case of a heating system, you don't place the heat vent so it points directly at the thermostat, though this would be the fastest way to get the thermostat to register no discrepancy. Rather, you place the vent to dump warm air on the far side of the room, so the entire room is warmed when the thermostat senses no discrepancy.

Although the thermostat is a convenient illustration of feedback processes in a device with which most people are familiar, it's certainly not the only example one could use. Indeed, in some ways it's a poor example. For one thing, it tends to imply that feedback systems can produce only conformity to a steady state, which isn't correct. Feedback systems can conform to constantly changing reference values (see Clark, 1996, chap. 10). Another problem is that the purposive quality of the system isn't too striking in the thermostat. It's seen more easily in systems such as the gyroscopic steering system used on a large ship. This system uses feedback principles to keep the ship on course, despite outside influences (changing winds and currents). Such a steering system is a nice metaphor for human action (see also Klein, 1987). Indeed, we'll argue shortly that people have their own internal gyroscopes that keep them on track toward the goals they've put before themselves.

ADDITIONAL ISSUES IN FEEDBACK CONTROL

Many further issues can be raised about feedback principles. We consider three of them, each stemming in some way from the fact that our portrayal thus far has been somewhat idealized. Each suggests at least one way in which the self-regulation displayed by a feedback system can be less than perfect.

Sloppy versus Tight Control

In discussing the elements of a feedback loop, we said that the input information is compared against a reference value and that what happens next depends on whether the comparator detects a discrepancy between them. A question that's rarely raised, but has clear implications for the functioning of the system, is this: How good is the comparator's ability to detect differences? There must be variability in this capability, such that sometimes a comparator can detect minor deviations but sometimes it can detect only substantial ones.

How well a system can detect errors has a straightforward implication for the system's ability to self-regulate (assuming a constant ability to adjust output). A system with precise error detection (high error sensitivity) can maintain close self-regulation. The slightest deviation from the reference value is noted, and compensatory action is taken immediately. A system with poor error detection is more variable. There is a range within which deviations from the reference value won't be noted, and thus won't be countered. Only when the input deviates far enough that it falls outside that range will the discrepancy be noted and an output called for. The functioning of this system will appear sloppier than that of the other system. This is one of several ways in which self-regulation can be less than optimal, while still embodying the elements of feedback control.

Lag Time

Just as we idealized the function of the comparator, we've also idealized the role of time in the self-regulatory process. That is, the influence of the output function isn't instantaneous. As noted earlier, the output has its impact on the input via some "effect on the environment." The time it takes for the effect of an output to influence the input function (an interval sometimes called *lag time*) must be taken into account in the

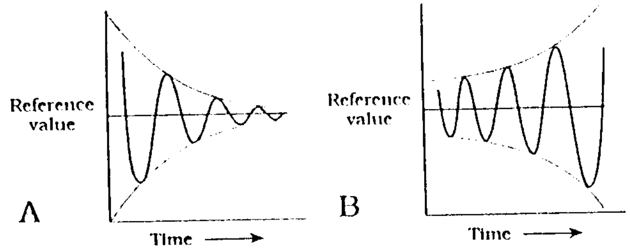

Figure 2.3. Failure to take into account the time required for feedback to register on the system can result in overcompensation, which in turn creates oscillation about the reference value. (A) If the overcompensation is less than the overshoot, the system will converge on the value. (B) If the overcompensation is larger than that, however, the system goes out of control, with deviations becoming more and more extreme.

loop's functioning (Clark, 1996, chap. 9). That is, the system has only one way of knowing whether its output has had the desired effect. Its only source of information on this question is the input function. The desired effect has occurred if the input is no longer discrepant from the reference value. But what if the output of this system takes a long time to exert its influence? What if the influence takes a long time to work its way back through the "effect on the environment" to show up in the input function?

If the system is "well designed," it takes the time lag into account. If the time lag isn't taken into account, the system continues to register a discrepancy long after it has executed the output needed to counter the initially sensed error. Consequently, the output function will continue longer than needed. At some point, the comparator will indicate that the discrepancy no longer exists and stop the output, but by this time a great deal of action has already taken place. Most of the effects of this action haven't yet registered on the input function. When these effects arrive, they'll create a deviation away from the reference value *opposite* to the initial deviation (see Figure 2.3).

This overcompensation (from failure to take into account the system's response time) can potentially lead to repetitive oscillation, as the system overcompensates first one way, then the other. If the overcompensation is smaller than the overshoot, self-regulation is erratic but eventually closes in on the reference value (Figure 2.3A). If it's larger, though, the system spins out of control (Figure 2.3B). It's important, then, to take lag time into account in considering how a feedback system is functioning. Failure to do so is a second way in which self-regulation can be less than optimal, even while embodying the character of feedback control.

Intermittent Feedback

A third way in which we over-idealized our description of feedback systems was to describe the loop in terms of continuous cycling of its component processes. This might imply that all the component processes function all the time. Sometimes this is so, but sometimes it's not. Schmidt (1988) provides a good description of the difference between these kinds of systems.

A continuous servo is one in which the effects of the output are continuously related to the reference value, and output varies continuously in response to variations in input. An example is the automatic steering mechanism in a ship, in which deviations are sensed continuously and changes in output are also made continuously (although even here it's important to note the need to consider lag time in the system's response).

A discontinuous, or intermittent, system is one in which the input and output functions don't track each other perfectly; rather, one function or the other occurs only intermittently. Sometimes the output is altered intermittently, as in the thermostat and furnace. That is, the room temperature changes continuously and these changes are noted continuously, but the heater is either off (when the temperature is within the range where the comparator doesn't notice the deviation) or on (when a deviation is noted).

There are also cases where the intermittency is on the input side. In such cases an output that's elaborate may be put into action, and its effects are checked only occasionally or only at the conclusion of the output. It's sometimes argued that fast body movements (too rapid for continuous feedback) rely on systems functioning this way (Schmidt, 1988, p. 221). It's also easy to see, however, how intermittency can produce erratic self-regulation. If input is checked only rarely, an output that's actually inappropriate won't be noticed as such until it's been largely or wholly carried out. This idea has obvious application to human concerns. If a person carries out an entire plan of action before checking to see whether it's having the desired effect, a great deal of effort may be spent on a plan that's actually counterproductive.

DISTINCTIONS AND FURTHER CONSTRUCTS

Although the negative feedback loop is the focal control structure in most of what follows, other concepts should also be mentioned. In this section we describe several, and indicate how they differ from those introduced thus far.

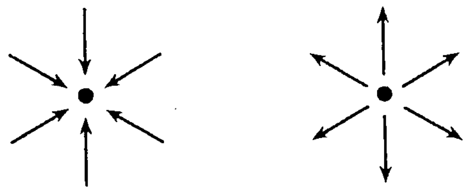

Figure 2.4. Negative feedback loops cause sensed qualities to shift *toward* positively valenced reference points. Positive feedback loops cause sensed qualities to shift *away from* negatively valenced reference points.

Positive Feedback Loops

The systems described thus far are called negative feedback loops, because their function is to negate (remove or diminish) a discrepancy between sensed input and reference value. These systems can be thought of as having reference values that are desired goals. These loops act to create conformity with reference values (see Figure 2.4). No matter the nature or direction of the deviation, the effort is to reduce it.

A positive feedback loop, in contrast, is a discrepancy amplifying system (DeAngelis, Post, & Travis, 1986; Ford, 1987; Maruyama, 1963). These loops create movement *away* from the reference value (Figure 2.4). Think of this as a system with an undesired goal, or an "anti-goal," as its reference value. Discrepancy amplifying loops try to move the currently perceived value away from the reference value. They are believed to be less common in naturally occurring systems than discrepancy reducing systems, because they're unstable. That is, they push away, and awayness goes on without limit. Whereas the conformity caused by a negative loop has a specific goal, the anticonformity caused by a positive loop doesn't. This creates instability, as they go on forever trying to create larger and larger deviations.

Positive feedback processes occur in functional ways in living systems (e.g., McFarland, 1971), but their action is typically constrained in some way by negative loops. Figure 2.5 illustrates symbolically how this can happen. A positive or avoidance loop creates pressure toward deviation from its reference value. Moving away occurs to a point, but the tendency to move away is captured by the influence of a negative loop. This loop then serves as an attractor, working to pull the sensed input into its orbit.

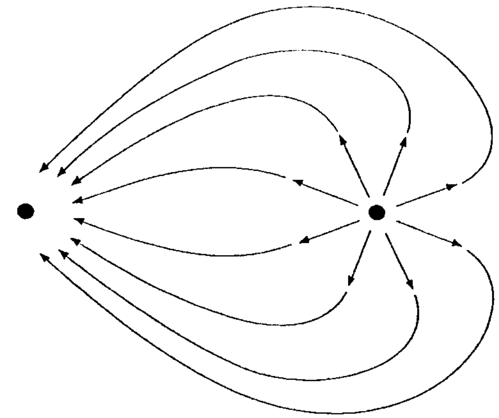

Figure 2.5. The effects of positive (avoidance) feedback systems are often bounded or constrained by negative (approach) feedback systems. A value moves away from an undesired condition in a positive loop, and then comes under the influence of a negative loop, moving toward the latter's desired condition.

The use of the word *orbit* in the preceding sentence suggests a metaphor that may be useful for readers to whom these concepts don't feel terribly intuitive. The metaphor is feedback processes as gravity and antigravity. The negative loop exerts a kind of gravitational pull on the sensed input it's controlling, pulling the input closer to it. The positive loop has a kind of antigravitational push, moving sensed values ever farther away. A similar metaphor is feedback as magnetism: negative feedback as similar to magnetic attraction, and positive feedback as similar to the repulsion that occurs when two poles of the same sign face each other.

Open Loop Systems

Most of our emphasis in this book will be on the closed loop systems described in the preceding sections, systems in which the effects of actions have an impact on the future perceptions of the systems, and ultimately on subsequent outcomes. We should also note, however, the difference between these systems and what have been termed open loop systems.

In an open loop system (Figure 2.6), the output of the system is triggered by the recognition of a particular state of affairs, and output

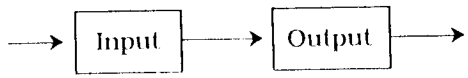

Figure 2.6. What's called an open loop system exists when some input cue triggers an output that is executed without subsequently checking on its impact.

simply executes. There's no checking on its consequences. There is a presumption that the consequence will be as intended. In an open loop system the output is always preprogrammed, with no provision for modifying it. Such systems are invoked when the output is very rapid, where there appears to be insufficient time for feedback information to be used (for discussion see chap. 7 of Schmidt, 1988). To some extent, the open loop system looks much like the stimulus-response depiction of behavior: A stimulus prompts an action, which simply occurs.

When is it necessary to assume the existence of closed loop systems, and when will an open loop system do? This is an important issue in discussions of movement control, and there remain substantial disagreements as to the answer (for a critique of the idea that open loop systems are used in movement control, see Marken, 1986, p. 275). Schmidt (1988, pp. 145–146) also pointed out that complex systems can be hybrids of closed and open loop subsystems. As an illustration, car engines have elements that follow closed loop principles – regulation of engine temperature through an arrangement of thermostat and control of coolant flow. They also have elements that are open loop – order and timing of the spark plugs.

The issue of combining open and closed loop systems is one that's important in discussions of movement control. In general, at the level of abstraction we're interested in, we'll argue that open loop systems are dysfunctional. In general, human functioning at this level involves closed loops.

Feedforward

Another concept to address here is feedforward. The principle of feedforward is used in somewhat different ways in different contexts, but the uses have some commonalities. Whereas feedback is information about the consequences of an output (reflected in a subsequent input), feedforward is anticipatory. This principle can be conceptualized in two ways.

First, feedforward can be considered as *anticipatory output*. In the cases of most interest to us, feedforward represents a first approximation of output, which occurs before any input is taken into account. Because

it would seem to precede the noting of any discrepancy between present state and reference value, the feedforward theoretically comes first. It doesn't take into account the present state, because nothing has come in yet to take into account.

The feedforward signal can be viewed as a "best estimate" of the output that's going to be needed (e.g., Marteniuk, 1992), an output that's "in the ballpark" (Greene, 1972). If the first approximation is good enough, the outcome will occur exactly as desired, with little or no subsequent correction needed. This result – behavior as intended with no need for correction – is often what leads people to invoke the concept of feedforward. (Indeed, this also can cause people to infer open loop control when that's not the case; even if the first approximation is exactly on the mark, it doesn't mean the system isn't checking to be sure.)

This use of the feedforward concept raises questions that are hard to answer. It's convenient in some respects to assume that the feedforward somehow bypasses the comparator (Figure 2.7A), but it's hard to be sure this is right. Maybe the feedforward instead occurs through the ordinary mechanism of the loop, *but is simply the first cycle of output to occur* (Figure 2.7B). That is, given the sudden presence of a reference value but no input yet available, the comparator by default should register a discrepancy. Thus the feedforward may be no different in principle from any other output.

In either case, this application of the feedforward concept raises questions about how the first-approximation output is created. In an electronic system the answer is often quite easy: The reference signal is sent directly on as the output, and the feedback loop acts simply to eliminate subsequent unwanted disturbances from sources outside the system. In a behavioral application of these ideas, all this becomes a little trickier.

A second application of the feedforward concept treats feedforward as creating a change in reference value. This application occurs in discussions of functioning of visual centers in the nervous system. When your eye makes a shift in position, you rarely notice it. If the room shifted the same amount, however, you'd notice it and be disoriented. You can tell whether you've made a movement in a stable environment or whether the environment itself has suddenly shifted, but how?

The usual answer (Evarts, 1973; Gallistel, 1980; Sperry, 1950) is that when the instruction to shift is sent to the muscles that move the eye (an output function within its own loop), a copy of that instruction (termed *efference copy*) is sent to the part of the brain that interprets the visual input. This copy causes the movement to be taken into account in interpreting

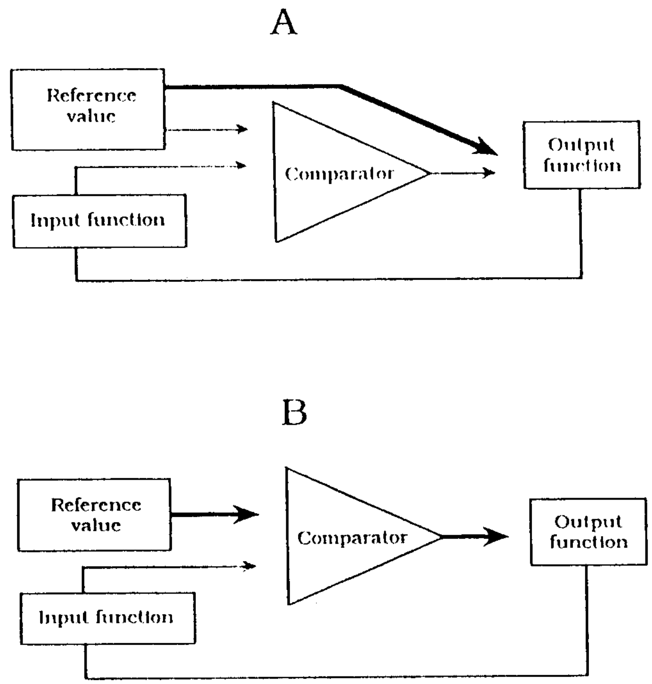

Figure 2.7. Two construals of feedforward as a first approximation of the output needed to create a desired result (i.e., to be identical to the reference value). Paths that are associated with feedback are gray, to indicate that they've not yet conveyed information when the feedforward occurs. (A) Here the feedforward signal is shown as bypassing the comparator as first-cycle output. (B) Here the feedforward signal is portrayed as occurring within the structure of the loop. This portrayal assumes that the comparison in the absence of input information registers a discrepancy, and the initial output is the feedforward signal. In this latter view, there's some question as to whether the feedforward signal really differs from any other output.

the visual input. The signal that a movement was ordered causes the visual system to adjust its reference value, to compensate for the fact that the movement had been ordered by some other part of the nervous system. The signal to make this adjustment is termed a *feedforward* signal.

INTERRELATIONS AMONG FEEDBACK PROCESSES

Thus far in the chapter we've described the functioning of feedback processes in terms of single cases, one loop at a time, although we've made occasional reference to information coming to the loop from elsewhere in a broader organization. We turn now to a consideration of several ways in which organizations can exist among feedback loops.

Interdependency

One way in which feedback loops can be related constitutes a kind of interdependency between separate feedback systems. In this arrangement, two systems that have their own goals and purposes are both responsive to the outputs of the other system.

An example often used in discussions of feedback processes in ecology concerns population sizes. Imagine a closed ecosystem (perhaps a small island) where there's a population of rabbits and a population of foxes. The rabbits constitute the major food source for the foxes, who are the only predators to threaten the rabbits. The size of each population is influenced by changes in the other population.

Given plentiful vegetation, the rabbits multiply rapidly and populate the island. As the rabbit population grows, the population of foxes also grows, because there's now more food (rabbits) for them. As the population of foxes grows larger, the population of rabbits begins to shrink, because so many of them are being eaten. This reduction in the available food supply has an adverse effect on the population of foxes, which also begins to diminish. With a depletion in the ranks of the predators, the rabbit population begins to reemerge.

In a well-balanced ecosystem, the populations of both species will find levels at which there is mutual stability. However, it's not impossible for the system to be driven outside this stability. For example, if the rabbit population is reduced to zero, the foxes have no food source, and the fox population also dies out. Here's a case where both population sizes are causes and both are also effects. Each changes over time in response to changes in the other, and each helps determine how the other changes.

Another example of interrelated feedback processes embeds one within the other, such that the output of each one acts as a disturbance on the other. Consider (from Schmidt, 1988) the case of two temperature-controlling feedback processes that are active simultaneously in many homes: a refrigerator and a furnace. The refrigerator keeps the temperature inside it at a set level, through a feedback loop with an internal thermostat. The furnace does the same thing, but with respect to the temperature inside the house.

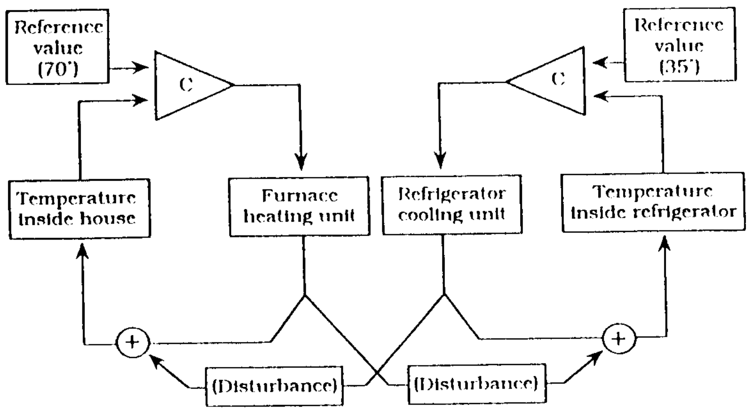

Figure 2.8. Interdependent feedback loops: One loop controls the temperature inside the refrigerator; the other controls the temperature inside the room. The output function of each loop acts as a disturbance on the other loop, influencing the perception registered in its input function. (Adapted from Schmidt, 1988, p. 145.)

Each of these systems also has an impact on the other (Figure 2.8). The output function of the refrigerator cools the air inside the refrigerator by transferring heat from inside to outside. The input function of the furnace (indicating the temperature of the room) responds not only to the actions of the furnace, but also to disturbances. One disturbance is the behavior of the refrigerator, which periodically dumps some heat from inside it into the rest of the room. When the refrigerator's compressor is running, it's not just cooling the refrigerator but heating the room.

In the same way, the furnace is acting as a disturbance on the refrigerator. Why does the refrigerator have to run its compressor periodically in the first place? Because heat gets into the refrigerator, raising its temperature. Where does the heat come from? From the air surrounding the refrigerator, the very air that the furnace is trying to keep warm. Thus the furnace's action makes the refrigerator work harder to keep its temperature down.

Reference Value and Input Function: How Do They Differ?

Let's return to the beginning of this chapter's story for a moment, think again about the elements that make up a feedback loop, and consider a

question. The loop has two sources of information, the input function and the reference value. It should be apparent that both sources are important – indeed, both are necessary for a feedback loop to exist and to function. The two aren't entirely equivalent, however. There's a sense in which the reference value is more important, in that it's more "demanding" than the input function.

We said earlier in this chapter that detecting a discrepancy between input and reference value leads to an attempt to reduce it. In principle, there are two ways by which the discrepancy might be reduced (leaving aside fortuitous environmental disturbances). The path we've focused on involves changing the output function in order to change the input so that it comes into closer conformity with the reference value. Why not change the reference value so that it more closely matches the input?

If these two values didn't differ in their stability, self-regulation would be truly haphazard. Instead of holding fast to some reference standard, the system would be all over the place, shifting the standard and shifting the input erratically, as a function of which moved first. Although there do appear to be cases in which standards change, such changes indicate something more complex than a single feedback loop. In a single loop, the reference value is rigidly in place (compared to the input function), and adjustments try to move input toward the standard.

A shift in standard appears to imply the operation of a second feedback path. The second path shown in Figure 2.9 uses the same information about the relation between input and reference value as the original loop uses. Its action is aimed not at changing the input, but at changing the reference value. We're not going to explore how it actually does so; we're only going to note that this path typically has a more gradual effect than the other path. As a result, this path has a discernible effect on the reference value only if a great deal of attempted output fails to reduce the discrepancy.

Such a secondary process probably is not inevitably in effect, but it is there sometimes. For example, most people sleep roughly eight hours per night (as a standard). If you get too little sleep over a night or two, your body typically compensates by catching up on a later night. Anecdotal evidence (from one of our colleagues) suggests that if you get smaller amounts of sleep over a long enough time, the standard appears to undergo a shift, so your body asks for less sleep than it used to. Because this change in reference value seems to occur through a secondary output

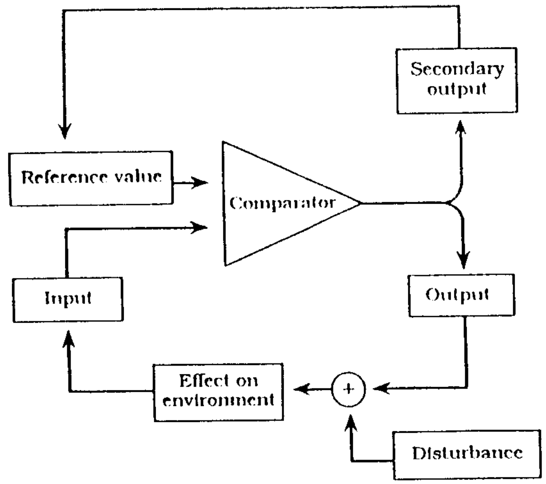

Figure 2.9. Feedback loops act to create changes in the input function, to shift its value toward the reference value. *Sometimes* another process is also in place (indicated by gray lines), which works to adjust the reference value toward the input. If it's in place, this additional process is typically very gradual, which keeps the reference value comparatively stable.

of the very comparison process that's using the reference value, we'll call this a "self-referential" change in standard, to distinguish it from the idea we're just about to introduce.

Hierarchies

Thus far we've talked about individual feedback loops and loops that are mutually interdependent. Even in the latter cases, the loops implicitly exist at the same level of abstraction. It's also been argued, though, that feedback loops can be organized in a hierarchical fashion, such that there are superordinate systems and subordinate systems (Powers, 1973a, 1973b). In this view, superordinate systems don't act directly on the outside environment. They act on subordinate systems. More specifically, they act by resetting the reference values of the subordinate systems.

Consider the thermostat one last time. As is true of many control structures, it's a flexible device that can operate with respect to a wide range of reference values. It gets its reference value from a superordinate system, the woman on whose wall the thermostat is wired. This woman also has a reference value (be comfortably warm). Rather than operate directly on the environment to produce heat – for example, by building a

fire – she operates by providing a new reference value to the subordinate system – resetting the thermostat from 60 degrees to 75 degrees. Given this change in reference value, the thermostat calls on the furnace, and the room temperature rises.

This example illustrates several points. First, the action of the superordinate system involves changing a reference value. Unlike the change discussed in the preceding section, however, this change process isn't self-referential. Rather, it occurs by the action of one system on the other system. A second point is that as the two systems act to create their desired conditions, each is monitoring its own input, which exists at its own level of abstraction. The thermostat assesses air temperature; the woman assesses her comfort level. Third, since the subordinate system is operating in the service of the superordinate system, progress toward discrepancy reduction occurs simultaneously for both systems as the air temperature rises.

A final point is that it's possible for a superordinate system to exist but not to act in a superordinate fashion. The woman in our example has actually been sitting in her living room for a long time, inattentive to the fact that she isn't comfortably warm, because she's been thinking about a problem she's facing at work. Only when someone came into the room and asked her if she was warm enough did she realize she wasn't. At that point she walked across to the thermostat and adjusted it. Until then, the superordinate system wasn't acting in a superordinate way. It was as though it had been temporarily disconnected. The fact that the higher-level system wasn't being used doesn't mean that it and its reference value weren't there. But while the superordinate system wasn't engaged, this two-level hierarchy was operating with the thermostat functionally superordinate.

In principle, it's possible for there to be many levels in a hierarchy of control (Figure 2.10). Among theorists interested in the control of movement of the human body, it isn't unusual for models to be proposed in which several layers of control are assumed (see, e.g., Greene, 1972; Rosenbaum, 1987, 1991b), though how to conceptualize their arrangement is debated among the theorists (Schmidt, 1988). William Powers (1973a, 1973b), whose background is in engineering rather than motor control, argued that the sensorimotor aspects of the nervous system are organized in such a way that as many as nine layers of control may be involved in voluntary movement. Because Powers tried to take higher-order intentions into account, as well as motor control, his model has some interesting implications, which we consider later.