

Figure 2.10. A hypothetical three-level hierarchy of feedback loops. The output function of the superordinate loop consists of the resetting of the reference value of the loop at the next lower level, a pattern that is repeated to the bottom level, where the output function takes another form. All levels are construed as monitoring feedback at the level of abstraction relevant to that level.

CONCLUDING COMMENT

This completes our description of the principles of feedback control. This was the hardest part. In many ways, the rest is all applications. That is, we use virtually all the ideas presented here as descriptive and explanatory principles in examining self-regulation in human behavior. This we do now, beginning with Chapter 3.

3

Discrepancy-Reducing Feedback Processes in Behavior

Our waking lives are filled with behavior. We haul ourselves out of bed each morning and set off for the day. We do much that's mundane and even trivial. Sometimes we do things we think are important at the moment but will later decide didn't matter much after all. Sometimes we do things that seem trivial at the time but will later seem more momentous. Occasionally our actions involve our core sense of self, though this is probably rare for most people. A lot of human behavior, after all, is an extended process of maintenance activity – buying groceries, washing dishes, driving from place to place to get things. Some behavior involves a lot of physical movement, but some of our most important behavior involves moving little more than our mouths and eyes.

FEEDBACK CONTROL IN HUMAN BEHAVIOR

In this chapter we argue that all these behaviors embody the processes of feedback control. We didn't invent this idea. Explicit statements about it go back at least to Miller, Galanter, and Pribram's (1960) book *Plans and the Structure of Behavior* (see also Hunt, 1965; MacKay, 1963, 1966; Powers, 1973a, 1973b), and ideas underlying it go back much further (chap. 3 of Miller et al., 1960, reviews the history of related thought). We're part of a second generation of psychologists who regard it as plausible.¹

¹ We reiterate that *behavior* as discussed here means primarily behavior at the level of personality and social psychology. Keep in mind, though, that theorists examining how the body controls movement find feedback ideas indispensable. This point is relevant for two reasons. First, the nervous system that moves the body is the same one that forms abstract intentions and carries them out. If feedback concepts are useful in the movement part, they're likely to be useful in other parts as well, since it's unlikely that different aspects of the nervous system depend on entirely different principles (cf. Gallistel, 1980). Second, even abstract intentions must be carried out by physical movements. Thus, having built-in conceptual links to theories about how the body carries out actions is a bonus for a theory about self-regulation of behavior at a higher level.

Early Applications of Feedback Principles

The first theorists to argue for the idea that feedback processes are important in macro-level human behavior were Miller, Galanter, and Pribram (1960). Their arguments appeared in a slim, engaging volume that presented a picture of human behavior guided by plans and goals and self-regulated by discrepancy-reducing feedback processes. It's a book of ideas and research possibilities rather than a description of support for the ideas. It articulated a feedback-based vision of behavior, inviting others to consider its usefulness.

Another well-known early user of feedback ideas was Bowlby (1969). He viewed the mind as a hierarchy of evaluating and controlling mechanisms operating on feedback principles, with the goal of creating and maintaining desired conditions. Best known, of course, is his pioneering work on infant attachment, a phenomenon that he viewed in feedback terms. The securely attached child uses its mother as a base. When it feels comfortably secure, it explores. If it's away too long and becomes anxious from the separation (or becomes frightened for some other reason), it returns. If it experiences too much closeness, it pulls away and explores again. The goal is to maintain a desired level of closeness to the attachment figure, neither too little nor (for a securely attached child) too much. Perceptions deviating from the desired level cause change in behavior. Again, although the feedback model provided an organizing principle that was useful, it wasn't a high priority to confirm that this actually is a feedback process.

In later years the functions involved in feedback processes received closer attention. It's now possible to point to studies indicating that feedback processes do occur in at least some kinds of behavioral self-regulation. Some of these studies are our own.

Our Starting Points

The path by which we came to consider the feedback viewpoint began neither with Miller, Galanter, and Pribram nor with Bowlby. We had an interest in phenomena that stem from self-focused attention, an interest prompted by work that others around us were doing at the time (Duval & Wicklund, 1972). Eventually we became convinced that some of these phenomena involved feedback loops.

More specifically, we came to argue that focusing attention on the self is often equivalent to engaging more fully the comparator that's involved in the self-regulation of the activity the person is presently engaged in (Carver, 1979; Carver & Scheier, 1981a). Unbeknownst to us, MacKay

had foreshadowed this argument much earlier, writing in 1963 that "an artifact capable of receiving and acting on information about the state of its own body can begin to parallel many of the modes of activity we associate with self-consciousness" (p. 227).

Before we turn to the work that makes this case, we should say a few words about terminology. The term *self-awareness* has a particular meaning here: focusing of attention on an aspect of the self (Duval & Wicklund, 1972; Wicklund, 1975). It doesn't imply a prolonged or penetrating self-examination or self-absorption, nor does it connote self-knowledge beyond the ordinary. *Attention* is selective processing of particular aspects of the informational field available, such that some information is more salient or more fully processed than others. Self-awareness is self-focused attention, selective processing of information about the self.

The word *self-consciousness* also has a specialized meaning here. In colloquial speech, self-consciousness tends to imply embarrassment. Here it doesn't. Fenigstein, Scheier, and Buss (1975) chose this label to refer to individual differences in the propensity to become self-aware. It provided a verbal shorthand to let authors refer to individual differences by one term (self-consciousness) and situationally created states of self-focus by another term (self-awareness) without constantly reminding readers which was which. Other phrases (self-focus, self-directed attention) are used generically to refer to either the disposition or the manipulated state.

As noted earlier, self-focus appears to have predictable consequences, some of which suggest that self-focus engages the comparator of a feedback system managing behavior (Figure 3.1). The earliest, pioneering work on the effects of self-focus was done by Wicklund and Duval (1971; Duval & Wicklund, 1972), who took a different meta-theoretical view than we're taking. They proposed that self-focus leads people to be more conscious of whatever standards for behavior are salient. As people become more aware of a standard, they tend to compare themselves against it. To the extent there's a discrepancy between their present behavior and the standard, the result is often an attempt to alter behavior so it conforms more closely to the standard. These functions are precisely those of a feedback system.

Self-Directed Attention and Comparison with Standards

If self-focused attention does engage a feedback process, the effects should be manifested in several ways. First, greater self-focus is likely

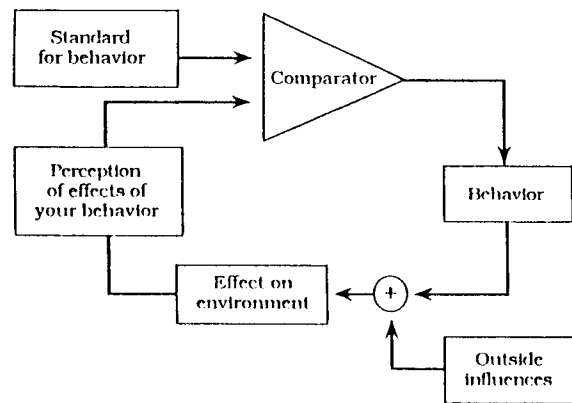


Figure 3.1. Behavior and perception as elements of a feedback loop guiding human action. In this view, self-focused attention enhances the effects of the comparator.

to be associated with the tendency to compare present self or present behavior against the salient standard. This is hard to verify directly, partly because it's an internal process and partly because it's difficult to separate the comparison function from the broader action of the loop. There is, however, indirect evidence for this assertion.

The studies that provide this support (Scheier & Carver, 1983) created situations with an implicit behavioral standard: to perform well at an assigned task. As a reference value, this is relatively abstract. While you're actually engaged in the behavior that's relevant to this goal, it can be hard to evaluate your progress. If your standard for a day's work is to "be productive," how do you decide at day's end whether you've met your goal? You look at the hard evidence of what you've done, and you look at some concrete comparison point (guidelines provided by a work assignment, or perhaps performances of other people engaged in similar activity). In sum, to compare your behavior against an abstract standard, you often seek out concrete information that facilitates the abstract comparison.

This is the principle behind the studies we conducted. In two studies, we asked people to copy a set of complex geometric diagrams onto paper as accurately as they could. Each diagram was projected for five seconds at a time. While the diagram was visible, the subject could view it but not copy. Only when the diagram disappeared could copying begin (or resume). The subject could view the diagram as often as desired, however, simply by pressing a button. Indeed, subjects were told to view each diagram as often as necessary to copy it as exactly as they could.

The diagram itself was a concrete standard of comparison for subjects' drawings. Looking at the diagram and comparing it to your own drawing is conceptually equivalent to engaging in the comparison assumed in the feedback model in Figure 3.1, at a concrete level. We assumed this process would occur in the service of a more abstract comparison between present behavior and the goal of "accurately duplicating the diagrams," or "performing well at the task." The dependent variable was how often subjects viewed the drawings. Those who did the task in the presence of a manipulation previously shown to increase self-awareness viewed the diagrams more frequently than those who did the task without it. In a second study, subjects whose dispositional tendencies toward self-consciousness were high viewed the diagrams more frequently than those whose self-consciousness was lower.

Two additional studies examined the comparison process with a different paradigm. These studies followed from the idea that one way to judge how well you're doing at a task is to compare your performance against performances of others on the same task. These studies were portrayed as part of a process of establishing norms for new tests of verbal ability. Subjects in one study completed a series of anagrams, calling out the solution for each one to the experimenter, who in turn told the subject how much time had elapsed, which the subject wrote down.

Before starting the task, the experimenter had drawn the subject's attention to booklets on the far end of the table which contained ostensible testing norms from previous semesters. After the anagrams were finished, subjects could consult the norms and compare their solution times against them. The norms thus constituted concrete standards that would facilitate the mental comparison between the abstract standard (performing well) and the performance. As predicted, subjects whose self-consciousness scores were higher consulted more pages of norms than did those whose self-consciousness was lower.

A final study was similar in logic to this one, but subjects had the chance to seek normative information before beginning the task. Participants were told they'd be working on 10 items from a test of abstract reasoning that was being evaluated. They were to choose their 10 items from among 4 categories, 2 of which concerned spatial reasoning, 2 semantic relationships. The item sets also varied (independently) in another way: For 2 sets, normative information was available; for the others, no norms were available. Subjects were to indicate how many items they wanted to work on from each of the 4 sets.

As expected, subjects who completed the form in front of a self-awareness-inducing stimulus chose more items for which norms were available than did control subjects. Presumably, this seeking of normative information occurred in the service of what would later be an attempt to compare abstract performance against a standard.

Self-Directed Attention and Conformity to Standards

If the effect of directing attention to the self in such situations is to engage more fully the comparator of a feedback loop, another thing should also happen, concerning the functioning of the feedback loop as a whole. If the comparison process is more fully engaged, the loop should do a better job of doing what it does: maintaining conformity between its sensed input and the standard. Thus, self-focus should promote closer self-regulation to the person's reference value.

This was, in fact, one of the first effects of self-awareness to be studied. In the first test of this hypothesis, Wicklund and Duval (1971) created a situation in which subjects were given the task of copying prose (printed in a foreign language). They were to copy quickly, and were left alone to work, either in the presence of a self-awareness-inducing stimulus (a mirror) or not. The experimenter surreptitiously timed the subject for a specific period, then interrupted. The question was how much the person copied in that time period. The answer was that subjects copied more when self-awareness was higher than when it was lower.

This is the first illustration of a phenomenon that apparently is sufficiently counterintuitive that it is sometimes misunderstood or ignored (e.g., Ingram, 1990): Increases in *self-focus* can promote increases in *task focus*. In this study, a standard for behavior was made salient: Copy quickly. Subjects with higher self-focus apparently applied themselves to the task more than those with lower self-focus.

This situation, of course, was an extremely simplified one. Further, it made use of one particular kind of standard: an instruction from the experimenter. Is closer conformity to reference values under high self-focus a general phenomenon? It would help to see the effect emerge under a variety of conditions. By this we mean seeing the effect in a variety of behavioral acts (copying prose isn't really the best example of an activity within the broad category "social behavior," though it has the advantage of being easy to measure). We also mean seeing the effect with regard to a variety of types of behavioral standards.

Many types of standards exist. Instructions, social comparison information, the norms of a society or group, a person's attitudes – all these are standards. Consider attitudes. A person's belief that something is desirable or undesirable can serve as a point of reference for that person, a value around which to regulate behavior. The idea that self-focus causes closer self-regulation suggests that when the attitude is salient as a reference value, the person will act in closer conformity to it when self-focus is high than when it's low. (Note that this hypothesis doesn't indicate when attitudes are naturally salient as standards, an issue we'll defer.)

This line of thought provided the basis for another study (Carver, 1975), which used a restricted pool of participants. First they completed a questionnaire expressing their opinion about punishment as a teaching technique. Some favored it (they thought it was effective and they'd be willing to use it); others opposed it (they thought it was ineffective and they'd be unwilling to use it). Weeks later, each participated in an experiment portrayed as a study of concept formation, in which they would teach another person to recognize an embedded visual pattern by a conditioning procedure. As part of that procedure, subjects would reward the learner after correct responses and punish him after incorrect responses. No talking would take place, and even nonverbal communication would be prevented by a partial wall.

Punishment was to be delivered as electric shock to the learner's finger. The punishment could vary in intensity from 1 to 10. Subjects received "sample" shocks to provide a sense of the different levels. They then were told they had freedom of choice over what intensity to administer for any given punishment. The experimenter said that psychologists are unsure about the effects of punishment on learning, and that participants were being asked to use their own opinions as guides. Thus each subject's attitude was made salient as a standard of comparison. For half the subjects in the study a mirror had been placed on the apparatus at eye level; for the other half it was removed.

The learner in this procedure was actually part of the research team. When the experiment began, he opened a door in the apparatus and turned off the shock. He then responded to the problems with a predetermined series of correct and incorrect responses. The punishment levels the subject chose for incorrect trials were recorded.

The results (Figure 3.2) were as predicted. Subjects relied on their attitudes to guide their actions when self-focus was high. Those who favored punishment used intense punishment, those who opposed punishment used very weak punishment. Subjects whose self-focus was low during




Figure 3.2. Punishment intensities delivered by subjects who either favored or opposed the use of punishment to produce learning, under conditions of either high or low self-awareness. (Data from Carver, 1975.)

the task apparently didn't use their opinions. There are many possible reasons for this; the one we regard as most plausible is that the task itself demanded a great deal of attention. It would have been easy for subjects to lose track of the fact that they were supposed to be using their attitudes to direct their behavior, unless their attention was periodically directed back to the comparison between what they were doing and what they were intending to be doing.

The studies just described make the case that self-focus causes closer conformity to salient standards, but they're far from the only studies to make this case. Induction of self-focus has caused effects as wide ranging as these: men conformed more to an implicit social standard of "chivalry" when giving punishment to a woman (Scheier, Fenigstein, & Buss, 1974); students allocating group earnings responded more to equity and equality norms when each was salient (Greenberg, 1980; Kernis & Reis, 1984); students rated erotica more consistently with their own standards (Gibbons, 1978); children in the midst of trick-or-treating conformed more to the instruction to take one (implicitly *only* one) piece of candy from a bowl (Beaman, Klentz, Diener, & Svanum, 1979); people opposed to stereotyping restrained themselves from doing so, whereas those who condoned it stereotyped even more (Macrae, Bodenhausen, & Milne, 1998). Individual differences in self-consciousness also related to improvement in sales performance among salesmen who saw that goal as important (Hollenbeck & Williams, 1987).

We want to emphasize two points about these studies: First, in all cases, self-focused attention caused subjects to conform more closely to the standard that was salient for them in that situation. Second, this effect of self-focus is an influence on a *process*, not a direct effect on the *content* of behavior. That is, being self-aware can make you *less* punitive if the salient standard is nonaggression (Scheier et al., 1974), but it can also make you *more* punitive if the standard calls for it (Carver, 1974, 1975). The content of behavior when self-focus is high depends on the reference value. People often can easily plug in one standard or another, and the effect of self-focus on overt action changes correspondingly.

Brain Functioning, Self-Awareness, and Self-Regulation

Though we have no intention of delving deeply into the nature of the brain's organization and functioning, we do want to address briefly a body of brain research that bears on some of these arguments. In particular, Stuss (1991) has argued that the prefrontal cortex is responsible for phenomena captured by the term *self-awareness* (see also Stuss & Benson, 1986).

Stuss (1991) holds that the frontal cortex has three levels of function. The first involves the ability to organize and maintain information in meaningful sequences. The second is an executive function involved in moving toward goals in novel or nonroutine situations. This control function is divisible into processes such as goal selection, means-end analysis, reflective evaluation of behavioral outcomes, and performance maintenance in light of those evaluations. The third level of frontal function is consciousness itself, the ability to be aware of oneself and one's relationship to the environment.

Much of this argument rests on studies of persons with frontal lobe damage. Stuss (1991) reviewed portions of this literature and (though noting inconsistencies) concluded that several themes can be extracted from it. First, many patients display a dissociation between knowledge and its use – for example, although they can detect errors, they don't use this knowledge appropriately. Second, there's a disturbance of subjective time. Patients with frontal lobe damage often show a loss of the sense of temporal order, a sense that's obviously necessary to implement plans or to experience the continuity of the self (see also Ingvar, 1985; Tulving, 1989). A third theme, more diffuse, is that there are differing levels of awareness of self, with higher levels being more likely to be disturbed by frontal lobe damage.

A very different kind of evidence bearing on a related idea has recently been reported by Gehring, Goss, Coles, Meyer, and Donchin (1993). Subjects in this study performed a long series of simple choices while electroencephalogram (EEG) data were recorded to assess aspects of their brain activity. Of particular interest was what occurred on trials where subjects made errors. The errors were associated with a particular pattern in the EEG, which indicated that a brain mechanism was noting the error even as it was being made. Further, this pattern was also associated with several measures indicating attempts at error correction. Taken together, the data suggest the existence of a brain system that detects errors and attempts to compensate for them.

Although these links between neuropsychology and social-personality psychology are more than just a little tenuous, they are also exciting and encouraging. They suggest that in work such as this there may emerge a better understanding of the physiological mechanisms within which the phenomena described in the preceding sections take place.

How Does Attention Shift to the Self in Ordinary Life?

A question that naturally arises concerning the effects of self-focused attention is how the effects translate to reality outside the lab. What causes attention to shift to the self in the natural course of experience? There are several answers to this question.

One answer is that attention fluctuates naturally back and forth, toward and away from the self, throughout moment-to-moment experience. For example, there's evidence that attention is demanded at decision points in behavior (cf. Norman & Shallice, 1986). Although cognitive psychologists rarely construe decision points in terms of self-focus, the need to make a decision relevant to a present intention would seem to imply that attention is being diverted to the self – more particularly, the aspect of self that's managing the execution of the intention. Thus, natural brief fluctuations in attention to the self probably occur throughout the course of carrying out an intention whose execution is not wholly automatic.

Attention to the behavior-managing aspect of the self also increases if something about your behavior produces an unexpected result (Norman, 1981). For example, if you go to where you parked your car and it isn't there, you check your behavior more closely to see where you walked, to assess whether you went where you intended to go. This sort of event-induced self-focus seems relatively specific to a review of the

most recent behavioral output as it relates to the intention, though if this doesn't reveal the error the search quickly broadens.

There are also social stimuli that serve to increase self-focused attention, and others that decrease it. Having an audience looking at you promotes self-focus (Carver & Scheier, 1978), as does spending time writing a story about yourself (Fenigstein & Levine, 1984). Attention to the self diminishes when you're immersed in a crowd (Diener, 1979; Prentice-Dunn & Rogers, 1980) or when you become involved in the story line of a movie.

Aspects of your inner experience can also draw attention inward. The need to make a decision might be considered an inner experience, but more obvious inner experiences can draw attention as well. The experience of physical arousal draws attention to the self (Wegner & Giuliano, 1980). So does the experience of emotion (Carr, Teasdale, & Broadbent, 1991; Salovey, 1992; Wood, Saltzberg, & Goldsamt, 1990), though positive emotion appears to do so less reliably than negative emotion. Although these studies make it clear that inner experiences draw attention, it isn't clear whether the self-focus they prompt shifts to comparison between self and salient standards, or whether it centers only on the state that attracted it.

We should also note that in thinking about self-focus in feedback loops (as opposed to other effects of self-focus²), the issue isn't exactly what makes attention be self-directed. The issue is what makes attention go to *the comparison between existing condition and reference value*. This comparison can happen if you attend to your "self," but it can also happen if you attend to an intention, task instructions, or some personally held value. That is, thinking about any one of those as a standard tends to make you think about your present behavior in relation to it. Many settings incorporate cues that remind people of their intentions, instructions, and values. These settings thus engage self-focused

² The self-awareness literature is one of several bodies of work that can be used to discuss how feedback concepts provide a useful model of human behavior. We should note, however, that not all effects of self-focus fit the self-regulatory model we're discussing here. Sometimes self-awareness has simpler effects. Attending to the self can make people more aware of existing affect (Scheier, 1976; Scheier & Carver, 1977) and the *absence* of expected internal states (Gibbons, Carver, Scheier, & Hormuth, 1979; Scheier, Carver, & Gibbons, 1979). Self-focus can cause people to make internal attributions (Arkin & Duval, 1975; Duval & Wicklund, 1973; Ross & Sicoly, 1979), and over time it can cause people to develop more elaborated and firmly anchored self-concepts (Hjelle & Bernard, 1994; Nashy, 1985, 1989a, 1989b). Thus, some of the consequences of self-focus are outside the realm of this discussion.

self-regulation by inducing the comparison process via the standard, rather than via the self more generally.

BROADENING THE APPLICATION OF FEEDBACK PRINCIPLES

In describing feedback processes in behavior in the preceding sections, we've simplified our description of the nature of the feedback that's available to people. Reality is more complex than we implied, and we need to address that complexity.

Sources and Nature of Feedback of the Effects of One's Behavior

The feedback that allows you to tell whether a discrepancy exists between the act intended and the act you're doing comes from several sources, and the relative importance of a given source varies from one situation to another. Often, you use several kinds of information at once. In some cases, proprioceptive and efference information is important. These cases tend to be represented more in the domain of motor control than in that of social behavior, but not always. If you're trying to impress someone with your gracefulness or athletic ability, proprioceptive information will be part of your input for judging whether you're doing what you intended to do.

Other sources of information are sight and sound. For example, as an untrained typist types a line of text, he decides whether he's executed each bit of behavior intended by checking to see if the right character or word appears on the computer screen. He also uses the clicking sound of the keys as feedback about whether or not the key has been fully depressed.

With even this many possibilities for feedback, a complex issue begins to emerge. For a feedback system to function properly, the information used as feedback must be relevant to the reference value (Harver, 1994; Langewiesche, 1993). As an absurd example, you could try to regulate typing by attending only to the sounds made by the keys, or (even more absurd) by whether the light coming in the window gets brighter or dimmer as you type. But since these sources of information won't be informative about the consequences of your actions for the intention you're trying to fulfill, they're useless as feedback.

The feedback channels that tell you what your movements have been and whether they were the movements you intended are important. But they aren't the only sources of feedback you use to track your behaviors. Indeed, it could easily be argued that they're the least important sources you use. Most social behavior has consequences that go far beyond the question of whether your act took the form you intended. These more remote consequences are typically the ones you need to monitor, rather than the question of whether the behavior took the form intended.

When you behave in a particular way calculated to make a good impression on someone, part of what you monitor is whether you're acting in the way you intend, but part of what you monitor is *whether you're making the kind of impression you want to make*. The input for this is quite different from the input for monitoring the form of the act. Discerning whether you're making the intended impression involves assessing other people's reactions to you. Your "effect on the environment" in this case means an effect on the other person. The effect is exerted through the filter of that person's preconceptions, interpretations, preoccupations, and distractions (all of which are disturbances). The impact eventually shows up in behavioral cues that he or she emits back to you, implying that he or she has some sort of impression of you.

The additional complexity involved in this kind of feedback raises a number of issues. For example, sometimes the cues you get are clear, but often they have to be interpreted. This means that to have the perceptual input you need in order to tell whether your behavior is appropriately regulated, you need to decode the cues. The process of interpreting the cues isn't easy. Although most people learn to do it fairly well, there are clear differences in how effective they are (i.e., how well their interpretations match normative ones). For example, people high in social anxiety perceive greater rejection from an interaction partner than do people lower in social anxiety, even when the facial cues they see are identical (Pozo, Carver, Wellens, & Scheier, 1991).

The process of interpreting the meaning of other people's behavior has been studied extensively in children by Dodge and his colleagues (Dodge, 1986; Dodge & Crick, 1990). An important difference between inappropriately aggressive children and those who aren't is that the aggressive children don't do a very good job of interpreting other children's actions toward them. As a result, they're more likely to infer hostile intent and regulate their own actions accordingly. This research makes an important point about self-regulation: Whether you use cues

emitted by other people to decide how to act in the first place, or whether you use them to decide how your actions have been received, these cues are critical information.

This sort of interpersonal feedback can also raise another issue, concerning time. Sometimes the cues that tell you the effects of your behavior come immediately, but sometimes not. Often they take some time to accumulate to the point where they can be clearly interpreted. With delays in receiving an interpretable input signal, the problem of lag time arises. As discussed in Chapter 2, effectively functioning feedback processes take into account the fact that feedback isn't always immediate. In such a case it's counterproductive to keep creating output until enough time has passed to determine the effects caused by the initial output (cf. Levine, 1992). Sometimes the lag in providing return input can be speeded up (for example, you can ask other people to tell you what they're thinking), but sometimes it can't.

In sum, the process of creating an input function from the various sources of information available to you is complex. The more abstract the reference value is, the more complex is the problem. As an illustration of the breadth of the problem, consider the use of social feedback as a way of verifying not simply the result of your behavior, but of your very self.

Use of Feedback for Self-Verification

People use social feedback to tell them not only how they're doing in their intended actions, but also to tell them what they're like as people. The information you get from others is sometimes used to fill in missing pieces in your picture of who you think you are (e.g., Cooley, 1902; Mead, 1934), but it also has other uses. It can also be used to verify your self-image (for overviews see Swann, 1990, 1996).

In some cases this use of social feedback works in counterintuitive ways. For example, common sense suggests that people want to be told good things about themselves. It turns out, however, that that's not always true. In particular, people want to be told they're who they think they are, even if they don't think they're very good. That is, someone with a negative self-image apparently wants feedback from others that confirms the negative image, rather than responses that contradict it. These are the cases that seem counterintuitive.

These efforts at confirmation of negative self-images have been shown in a variety of ways. For example, people prefer interaction partners

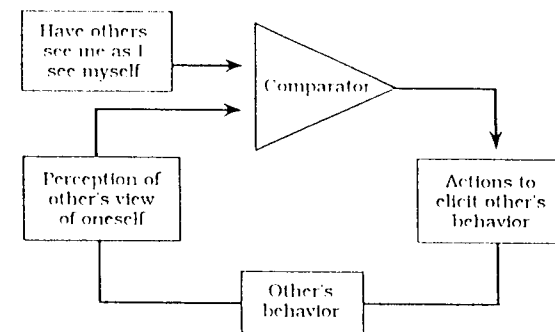


Figure 3.3. In self-verification, a person acts in ways that elicit responses from others which confirm that they think the actor is the person he thinks he is, even if the characteristics being confirmed are unpleasant ones. (From Swann, 1990).

who give confirming social feedback over those who don't (e.g., Joiner & Metalsky, 1995; Swann, Pelham, & Krull, 1989; Swann, Wenzlaff, Krull, & Pelham, 1992). Of particular interest at present is the fact that people act in ways that appear calculated to *elicit* the preferred feedback. This tendency is especially strong among people who think their interaction partners already have a misimpression of what they're like (Curtis & Miller, 1986; Swann & Ely, 1984; Swann & Read, 1981) and among those who are confident about what they're like (Pelham & Swann, 1994).

The process of eliciting self-confirming feedback is a feedback process (Figure 3.3). The person acts so as to promote a particular impression in the mind of another person. This behavior also elicits behavior, in return, from the other that confirms the actor's impression of himself. If this doesn't happen (if the cues from the other indicate the desired impression hasn't developed), the actor will modify behavior to try to promote the correct impression in the mind of the other.

It's not that people eliciting negative reactions from others necessarily want to *retain* negative self-views. Rather, they have two kinds of feedback processes occurring at once. The more pressing process is an attempt to make others' impressions conform to the current self (Figure 3.4). The more hidden process is an attempt to make the current self conform to the desired self. The latter involves creating a closer match between present self-image and a desired self-image. Once this happens, the person will seek out feedback in subsequent encounters to verify the newly altered current self. Until then, however, the more positive feedback will lack confirming value, and won't be sought out.




Figure 3.4. (A) Efforts at self-verification through social feedback from others is an attempt to cause perceptions of social reality to conform to one's current self-image. (B) In many cases there are also (independent) efforts to cause perceptions of one's current self to conform to a desired self-image. Usually, however, these efforts involve different sorts of behavior, because the goals are quite different.

Social Comparison and Feedback Control

Another literature that fits with the argument that behavior follows feedback principles is the literature of social comparison. Social comparison theory (Festinger, 1950, 1954; Suls & Wills, 1991; Wood, 1989, 1996) stems partly from the observation that much of social reality can't be verified objectively. It's easy to measure a room or your weight, but comparable measurement techniques don't exist to tell you whether a painting is good art or whether you have a good life.

Festinger argued that the ambiguity of subjective aspects of the world leads people to search out comparison information from others. Sometimes people use social comparison to validate the levels of their abilities (e.g., Conolley, Gerard, & Kline, 1978; Festinger, Torrey, & Willerman, 1954; Jones & Regan, 1974; Suls, Gastorf, & Lawhon, 1978; Wood, 1989). In other cases people use it to assess whether their attitudes, emotions, or actions are correct or appropriate (e.g., Gerard, 1963; Goethals & Nelson, 1973; Radloff, 1961; Schachter, 1951, 1959). In both cases, ambiguity leads people to seek others and get information from them directly or indirectly, to create a consensual validation of reality. Once a consensus has been reached, there's also an implicit (and sometimes explicit) pressure to conform to the consensually defined value.

Social comparison that pertains to evaluating the appropriateness of attitudes, emotions, and actions bears a striking structural resemblance to the processes we've been discussing. That is, you determine a reference value (the consensually defined standard) and then conform to it. This aspect of social comparison resembles the gravitational metaphor we




Figure 3.5. (A) Social comparison regarding attitudes involves discerning a group value and moving toward it, regardless of the direction of one's initial deviation from it. (B) Social comparison regarding ability involves a similar pressure to conform to the group's norm, but also a second pressure toward higher ability (gray arrows). These pressures work in the same direction for some people (1 and 2), but in opposing directions for others (3 and 4), causing internal conflict.

mentioned in Chapter 2. The standard of appropriateness provides an attraction point, and the person is drawn to it. Although much of the early interest in social comparison focused on how groups act to enforce conformity, overt pressure isn't required to cause conformity (Asch, 1951, 1956). Indeed, it isn't even necessary that the group be present, only that their views be known (Wicklund & Duval, 1971).

An important contribution of the social comparison model is to point out a way in which reference values for self-evaluation are made salient or, indeed, are created. That is, the first process in social comparison is to determine the standard of appropriateness or correctness. Only then can you conform to it. It's undeniable that a great deal of human social behavior relies on consensual definitions about what's appropriate or correct, ranging from issues of morality to issues of fashion. Social comparison theory highlights the importance of this process. Indeed, this analysis is usually assimilated directly to the discussion of reference groups, groups to which people refer or compare themselves in order to determine the norms to conform to (e.g., Newcomb, 1958).

Social comparison with respect to ability is a little more complex (see Figure 3.5). The information-gathering aspect of the process, prompted by ambiguity, is much the same as with attitudes, but there exist two pressures rather than one (Festinger, 1950, 1954). There's pressure

to conform to the group's ability level, but there's also an upward pressure -- a desire to have higher ability (represented by the gray arrows). For anyone with ability below the group's norm (persons 1 and 2 in panel B), these two tendencies act in concert with one another, both creating pressure to move higher, more closely approximating the norm. However, for a person with above average abilities (persons 3 and 4), the two tendencies conflict. The group conformity pressure tends to promote a reduction in displayed ability, whereas the upward pressure promotes the opposite.

An illustration of this point is the experiences of a group who play racquetball with one another on a regular basis. There's pressure within the group for everyone to be equivalent in ability, so the games remain enjoyable. Players who are less competent are encouraged to improve, and players toward the top of the group's distribution are discouraged from using their full ability. But there's also pressure to get better, even for the ones at the top. By following this pressure, people sometimes gravitate out of their groups into other ones, where both pressures again work in the same direction instead of promoting conflict.

The term *upward comparison* refers to comparisons made with people who are better off than oneself. *Downward comparisons* are those made with people who are worse off than oneself (Taylor & Lobel, 1989; Wills, 1981). These terms have meaning when applied to social comparison of ability, and also in other contexts that seem to follow the same principle as ability comparisons. For example, suppose you've been diagnosed with a serious disease. How are you doing? To answer this question, you need a comparison group. But what group to use? A group of people who've recovered from the disease would be an upward comparison. Patients with more severe cases of the disease would be a downward comparison.

There's evidence that patients engage in both kinds of comparisons (Helgeson & Taylor, 1993; see also Buunk & Gibbons, 1997). They tend to compare downward in gauging how well they're doing, and they prefer to affiliate with patients who are doing better than they are. This upward comparison can provide them with a sense of hope or motivation to improve, and it can also provide an opportunity to learn more about how to help themselves (Butler, 1992).

For present purposes, what's most important about this illustration is that this kind of upward comparison suggests a gravitational model, a model of pulling oneself closer to a desired point, a model of the conformity tendency that's created by a discrepancy-reducing feedback loop.

When you make adaptive upward comparisons, you think of the better-off group as a goal that's relevant to you, and try to *pull yourself toward it* (see also Collins, 1996; Lockwood & Kunda, 1997). This phenomenon seems very amenable to understanding in terms of discrepancy-reducing feedback processes.

SUMMARY

In this chapter we've outlined a number of phenomena that seem to embody the processes that characterize discrepancy-reducing feedback loops. Some of what we've discussed is at a micro level (evidence of comparisons); some is at a macro level (evidence of conformity tendencies). The research domains have been diverse enough to establish a broad base for the principle that much of human behavior is a matter of isolating a point of reference, then trying to conform to it (and sometimes to convince others of it). Not everything fits this picture, though. For example, social comparison can be upward or downward. Downward comparisons don't fit here. They're about something else. Their story comes next.