

3

ON THE STRUCTURE OF BEHAVIORAL SELF-REGULATION

CHARLES S. CARVER* AND
MICHAEL F. SCHEIER†

**University of Miami, Coral Gables, Florida*
†Carnegie Mellon University, Pittsburgh, Pennsylvania

We are interested in the structure of human behavior, broadly conceived. This interest has taken us into several specific research domains, including test anxiety, social anxiety, and self-regulation of health-related behavior. But in some respects the specific explorations in diverse areas of work have been in the service of a more general interest in the structure of behavior.

The questions underlying this interest are very abstract: What is the most useful way to think about how people create actions from intentions and desires? Once people have decided to do something, how do they stay on course? What processes account for the existence of feelings, as people make their way through the world? At the fore of our thinking over the past two decades is the idea that behavior is a self-regulatory event. It is an attempt to make something happen in action that is already in mind. This general idea forms the basis of this chapter.

This chapter is organized in terms of a series of conceptual themes that we've found useful. Some of them have been central in our thinking for a long time; others have been taken up only more recently. We start simple, with basic ideas about the nature of behavior and the organization of some of the processes by which we believe behavior is regulated. We then turn to consideration of emotion—how we think it is created and how certain classes of affects differ from each other. This is followed by a discussion of

the fact that people sometimes are unable to do what they set out to do, and what follows from that problem. The next sections are more speculative. They deal with dynamic systems and catastrophe theory as models for understanding behavior and how these models may contribute to the ways in which people such as ourselves think about self-regulation.

1. BEHAVIOR IS GOAL DIRECTED AND FEEDBACK CONTROLLED

The view we take on behavior begins with the concept of goal and the process of feedback control. We see these ideas as intimately linked. Our focus on goals is very much in line with a growing reemergence of goal constructs (Austin & Vancouver, 1996; Elliott & Dweck, 1988; Miller & Read, 1987; Pervin, 1989). A variety of labels are used in this literature: for example, *current concern* (Klinger, 1975, 1977), *personal strings* (Emmons, 1986), *life task* (Cantor & Kihlstrom, 1987), and *personal project* (Little, 1983). In all these theories, there is room for individualization. That is, a life task can be achieved in many ways. People choose paths that are compatible with other aspects of their life situations (many current concerns must be managed simultaneously) and other aspects of their person-allies.

Two goal constructs that differ somewhat from those named thus far are the *possible self* (Markus & Nurius, 1986) and the *self-guide* (Higgins, 1987, 1996). These constructs are intended to bring a dynamic quality to conceptualization of the self-concept. In contrast to traditional views, but consistent with other goal frameworks, possible selves are future oriented. They concern how people think of their unrealized potential, the kind of person they might become. Self-guides similarly reflect dynamic aspects of the self-concept.

Theorists who use these various terms—and others—have their own emphases (for broader discussions, see Austin & Vancouver, 1996; Carver & Scheier, 1998), but many points are the same. All include the idea that goals energize and direct activities; these views implicitly (and sometimes explicitly) convey the sense that goals give meaning to people's lives (cf. Baumeister, 1989). In each theory there is an emphasis on the idea that understanding the person means understanding the person's goals. Indeed, in the view represented by these theories, it is often implicit that the self consists partly of the person's goals and the organization among them.

A. FEEDBACK LOOPS

How are goals used in behaving? Part of our answer is that goals serve as reference values for feedback loops. A feedback loop, the unit of cybernetic control, is a system of four elements in a particular organization

(cf. Miller, Galanter, & Pribram, 1960): an input function, a reference value, a comparator, and an output function (Figure 1).

An input function is a sensor. We will treat this function as equivalent to perception. The reference value is a second bit of information (i.e., in addition to the input function). We'll treat the reference values in the loops we're interested in as goals. A comparator is a device that makes comparisons between input and reference value. The comparison yields one of two outcomes: either the values being compared are discriminably different from one another or they're not. The comparison can vary in sensitivity. Sometimes very small discrepancies are detected; sometimes only quite large ones.

Following the comparison is an output function. We will treat this as equivalent to behavior, although sometimes the behavior is internal. If the comparison yields a "no difference," the output function remains whatever it was. If the comparison yields "discrepancy," the output changes.

There are two kinds of feedback loops, corresponding to two kinds of goals (Figure 2). In a negative or discrepancy-reducing loop, the output function is aimed at diminishing or eliminating any detected discrepancy between input and reference value. It yields conformity of input to reference. This conformity is seen in the attempt to approach or attain a valued goal.

In this view, behavior isn't for the sake of behavior, but occurs in the service of creating and maintaining conformity of input to standard. Behavior can create conformity, but disturbances from outside also can create conformity. Although disturbances often change conditions ad-




FIGURE 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.




FIGURE 2 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. *Note:* From *On the self-regulation of Behavior*, by C. S. Carver and M. F. Scheier, 1998, New York: Cambridge University Press; Copyright 1998 by Cambridge University Press. Reprinted with permission.

versely (enlarging a discrepancy with the reference value), they also can change conditions favorably (diminishing a discrepancy). In the first case, recognition of a discrepancy prompts a change in output, as always. In the second case, the disturbance preempts the need for an output adjustment, because the system sees no discrepancy. Thus no output adjustment occurs.

The second kind of feedback loop is a positive or discrepancy-enlarging loop (Figure 2). The reference value here is not one to approach, but one to avoid. Think of this as an “anti-goal.” A psychological high-level example is a feared possible self. Other, more concrete examples would be traffic tickets, public ridicule, and being fired from your job. A positive loop senses present conditions, compares them to the anti-goal, and tries to enlarge the discrepancy. For example, a rebellious adolescent who wants to be different from his parents senses his own behavior, compares it to his parents’ behavior, and tries to make his own behavior as different from theirs as possible.

The action of discrepancy-enlarging processes in living systems is typically constrained in some way by discrepancy reducing loops (Figure 3). To put it differently, avoidance behaviors often lead into approach behaviors. An avoidance loop creates pressure to increase distance from the anti-goal. The movement away occurs until the tendency to move away is captured by the influence of an approach loop. This loop then serves to pull the sensed input into its orbit. The rebellious adolescent, trying to be different from his parents, soon finds other adolescents to *conform* to, all of whom are deviating from their parents.

Our use of the word orbit in the last paragraph suggests a metaphor that may be useful for anyone to whom these concepts do not feel terribly




FIGURE 3 The effects of positive feedback systems are often bounded or constrained by negative 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 its desired condition. *Note:* From *On the Self-Regulation of Behavior*, by C. S. Carver and M. F. Scheier, 1998, New York: Cambridge University Press; Copyright 1998 by Cambridge University Press. Reprinted with permission.

intuitive. You might think of feedback processes as metaphorically equivalent to gravity and anti-gravity. The negative feedback loop exerts a kind of gravitational pull on the input it is controlling, pulling that input closer to its ground zero. The positive loop has a kind of anti-gravitational push, moving sensed values ever farther away. Don’t forget, though, that this is a metaphor. More is involved here than a force field.

It’s worth noting that the situations people confront are often more complex than the one shown in Figure 3. Often there are several potential values to move toward. For this reason, one positive value won’t always capture or constrain all the avoidance attempts. Thus, if several people are trying to deviate from a mutually disliked reference point, they may diverge from one another. For example, one adolescent trying to escape from his parents’ values may gravitate to membership in a rock band, whereas another may gravitate to the army. Presumably the choice of what direction to approach will depend, in part, on the fit between the available reference values and the person’s preexisting values and, in part, on the direction the person took initially to escape from the anti-goal.

Some years ago we argued that the comparator of a psychological feedback loop is engaged by increases in self-focused attention (Carver,

1979; Carver & Scheier, 1981, 1990). Indeed, the similarity in function between manipulations of self-focus and the elements of the feedback loop was one thing that attracted us to the feedback model in the first place. Self-focused attention leads to more comparisons with salient standards (Scheier & Carver, 1983), and it enhances behavioral conformity to salient standards. The standards have ranged from instructions to personal attitudes to subjective norms (for reviews, see Carver & Scheier, 1981, 1998). On the avoidance side, self-focus has led to more rejection of attitudinal positions held by a negative reference group (Carver & Humphries, 1981) and to stronger reactance effects (Carver & Scheier, 1981).

The literature of self-awareness is not the only one that fits the picture of feedback loops, however (Carver & Scheier, 1998). For example, social comparison processes can easily be viewed in these terms: People use upward comparisons to help them pull themselves toward desired goals; people use downward comparisons to help them force themselves farther away from (upward from) those who are worse off than they are.

B. REEMERGENT INTEREST IN APPROACH AND AVOIDANCE

Our interest in the embodiment of these two different kinds of feedback processes in behavior is echoed in the recent reemergence of interest in two modes of regulation in several other literatures. One of these derives from a group of theories that are biological in focus. The research base of these theories ranges from animal conditioning and behavioral pharmacology (Gray, 1982, 1987b) to studies of human brain activity (Davidson, 1992a, 1992b; Tomarken, Davidson, Wheeler, & Doss, 1992). The theories incorporate the idea that two systems (sometimes more) are involved in regulation of behavior.

One system, handling approach behavior, is called the behavioral activation system (Cloninger, 1987; Fowles, 1980), behavioral approach system (Gray, 1987a, 1990), behavioral engagement system (Depue, Krauss, & Spont, 1987), or behavioral facilitation system (Depue & Iacono, 1989). The other, dealing with withdrawal or avoidance, is usually called the behavioral inhibition system (Cloninger, 1987; Gray, 1987a, 1990), although it is sometimes termed a withdrawal system (Davidson, 1992a, 1992b). The two systems are generally regarded as independent, because they're believed to be regulated by different brain mechanisms.

Another literature with a dual-motive theme derives from self-discrepancy theory (Higgins, 1987, 1996; Higgins, Bond, Klein, & Strauhan, 1986). This theory holds that people relate their perceptions of their actual selves to several self-guides, particularly ideals and oughts. Ideals are qualities the person desires to embody—aspirations, hopes, positive wishes for the self. Living up to an ideal means attaining something desired. An ideal is clearly an approach goal. We believe it is *purely* an approach goal.

Oughts, in contrast to ideals, are defined by a sense of duty, responsibility, or obligation. An ought is a self that one feels compelled to be, rather than intrinsically desires to be. The ought self is a positive value in the sense that people try to conform to it. However, living up to an ought also implies acting to *avoid a punishment*—self-disapproval or the disapproval of others. In our view, oughts are more complex structurally than ideals. Oughts intrinsically imply both an avoidance process and an approach process. Their structure thus resembles what was illustrated earlier in Figure 3. Recent work has demonstrated the avoidance aspect of the dynamics behind the ought self (Higgins & Tykocinski, 1992).

II. HIERARCHICALITY AMONG GOALS

Another theme in the translation of goals into behavior reflects the obvious fact that some goals are broader in scope than others. How to think about the difference in breadth is not always easy to put your finger on. Sometimes it's a difference in temporal commitment. Sometimes, though, it's more than that. It's a difference in the goal's level of abstraction.

A. PREMISE: GOALS CAN BE DIFFERENTIATED BY LEVELS OF ABSTRACTION

The notion that goals differ in their level of abstraction is easy to illustrate. You might have the goal of being an honorable person or a self-sufficient person. These goals are at a relatively high level of abstraction. You also may have the goal of avoiding a person at work who gossips or of making dinner for yourself. These are all at a lower level of abstraction. The first set concerns being a particular kind of *person*, the second set concerns completing a particular kind of *action*. You could also think of goals that are even more concrete than the latter set, such as the goal of walking quietly to your office and closing the door without being heard or the goal of cutting vegetables into a pan. These goals (which some would call plans or strategies instead of goals) are closer to specifications of individual acts than were the second set just described, which were more summary statements about the desired outcomes of intended action patterns.

How should we think about this difference in abstraction among goals? As you may have noticed, the examples used to illustrate concrete goals relate directly to the examples of abstract goals. We did this to point out that abstract goals are linked to concrete goals in a hierarchy of levels of abstraction. William Powers (1973) argued that a hierarchical organization of feedback loops underlies the self-regulation of behavior. Because

feedback loops imply goals, this argument also constituted a model of hierarchical structuring among the goals involved in creating action.

His general line of thinking ran as follows: In a hierarchical organization of feedback systems, the output of a high-level system consists of the resetting of reference values at the next lower level of abstraction. To put it differently, higher order or superordinate systems "behave" by providing goals to the systems just below them. The reference values specified are more concrete and restricted as one moves from higher to lower levels. Control at each level reflects regulation of a quality that contributes to the quality controlled at the next higher level. Each level monitors input at a level of abstraction appropriate to its own functioning, and each level adjusts output so as to minimize its discrepancies. It is not assumed that one processor handles functions at various levels of abstraction. Rather, structures at various levels handle their concerns simultaneously.

Powers focused particularly on low levels of abstraction. He said much less about the levels that are of most interest to us, except to suggest labels for several levels whose existence makes intuitive sense. Programs are activities involving conscious decisions at various points. Sequences, the next level down, run off directly once cued. The level above programs is principles, qualities that are abstracted from (or implemented by) programs. These are the kinds of qualities represented by trait labels. Powers gave the name "system concepts" to the highest level he considered. Goal representations there include the idealized overall sense of self, relationship, or group identity.

A simple way of portraying this hierarchy is in Figure 4. This diagram omits the loops of feedback processes, using lines to indicate only the links among goal values. The lines imply that moving toward a particular lower goal contributes to the attainment of some higher goal (or even several at once). Multiple lines to a given goal indicate that several lower-level action qualities can contribute to its attainment. As indicated previously, there are goals to "be" a particular way and goals to "do" certain things (and at lower levels, goals to create physical movement).

B. ACTION IDENTIFICATION

Although the Powers hierarchy per se has not been studied empirically, another theory that strongly resembles it—Vallacher and Wegner's (1985) action identification theory—has been. This model is framed in terms of how people think about their actions, but it also conveys the sense that how people think about their actions is informative about the goals by which they are guiding their actions.

People can identify a given action in many different ways, and the act identifications can vary in level of abstraction. High-level identifications




FIGURE 4 A hierarchy of goals (or of feedback loops). Lines indicate the contribution of lower level goals to specific higher-level goals. They also can be read in the opposite direction, indicating that a given higher order goal specifies more concrete goals at the next lower level. The hierarchy described in text involves goals of "being" particular ways, which are attained by "doing" particular actions. *Note:* From *On the Self-Regulation of Behavior*, by C. S. Carver and M. F. Scheier, 1998. New York: Cambridge University Press. Copyright 1998 by Cambridge University Press. Reprinted with permission.

are abstract (e.g., becoming more cultured), lower-level identifications get more and more concrete (e.g., attending a ballet, listening to sounds, and watching people move while you sit quiet and still). Low-level identifications tend to convey a sense of "how" an activity is done; high-level identifications tend to convey a sense of "why."

Although the Vallacher and Wegner (1985) model is hierarchical, it doesn't specify what qualities define various levels: It simply assumes that where there is a potential emergent property there is the potential for differing levels of identification. On the other hand, the examples used to illustrate the theory tend to map onto the levels of the Powers hierarchy: sequences of acts, programs of actions (with variations of smaller-scale and larger-scale programs), and principles of being. Thus, work on action identification tends to suggest the reasonableness of these particular levels of abstraction in thinking about behavior.

C. MULTIPLE PATHS TO HIGH-LEVEL GOALS, MULTIPLE MEANINGS IN CONCRETE ACTION

Although the hierarchy we are discussing is in some ways very simple, it has implications for several issues in thinking about behavior (for a broader treatment, see Carver & Scheier, 1998). It is implicit here that goals at any given level can often be achieved by a variety of means at

lower levels. This flexibility is particularly apparent at upper levels of the hierarchy, where the goals are abstract. This permits one to address the fact that people sometimes shift radically the manner in which they try to reach a goal when the goal itself has not changed. This happens commonly when the emergent quality that is the higher order goal is implied in several lower order activities. For example, a person can be helpful by writing a donation check, picking up discarded for a recycling center, volunteering for a charity, or holding a door open for someone else.

Just as a given goal can be obtained via multiple pathways, so can a specific act be performed in the service of diverse goals. For example, you could buy someone a gift to make him or her feel good, to repay a kindness, to put him or her in your debt, or to satisfy a perceived holiday-season role. Thus, a given act can have strikingly different meanings, depending on the purpose it's intended to serve. This is an important subtheme of this view on behavior: Behavior can be understood only by identifying the goals to which behavior is addressed. This isn't always easy to do, either from an observer's point of view (cf. Read, Druiian, & Miller, 1989) or from the actor's point of view.

D. GOAL IMPORTANCE: GOALS AND THE SELF

Another point made by the notion of hierarchical organization concerns the fact that goals are not equivalent in their importance. The higher you go into the organization, the more fundamental to the overriding sense of self are the qualities encountered. Thus, goal qualities at higher levels would appear to be intrinsically more important than those at lower levels.

Goals at lower levels are not necessarily equivalent to one another in importance, however. Just as it's sometimes hard to tell what goal underlies a given behavior, it can also be hard to tell from a behavior how important is the goal that lies behind it. In a hierarchical system there are at least two ways in which importance accrues to a concrete goal. The more directly a concrete action contributes to attainment of some highly valued goal at a more abstract level, the more important is that concrete action. Second, an act that contributes to the attainment of several goals at once is more important than an act that contributes to the attainment of only one goal.

Relative importance of goals returns us to the concept of self. In contemporary theories the self-concept has several aspects: one is the structure of knowledge about your history; another is knowledge about who you are now; another is the self-guides or images of potential selves used to guide movement from the present into the future (which may also be working models). A broad implication of this sort of theory is that the self is partly the person's goals.

III. FEEDBACK CONTROL AND CREATION OF AFFECT

We shift now to another aspect of human self-regulation: emotion. Here we add a layer of complexity to the feedback model which differs greatly from the complexity represented by hierarchicality. Again the fundamental organizing principle is feedback control, but now the control is over a different quality.

What are feelings and what makes them exist? Many have analyzed the information that feelings provide and situations in which affect arises (see, e.g., Frijda, 1986; Lazarus, 1991; Ortony, Clore, & Collins, 1988; Roseman, 1984; Scherer & Ekman, 1984). The question we address here is slightly different: What is the internal mechanism by which feelings arise?

A. THEORY

We have suggested that feelings arise as a consequence of a feedback process (Carver & Scheier, 1990). This process operates simultaneously with the behavior-guiding function and in parallel to it. One way to describe this second function is to say it's checking on how well the behavior loop is doing at reducing its discrepancies. Thus, the input for this second loop is a representation of the *rate of discrepancy reduction in the action system over time*. (We focus first on discrepancy-reducing loops, turning later to enlarging loops.)

We find an analogy useful here: Because action implies change between states, consider behavior analogous to distance. If the action loop deals with distance, and if the affect-relevant loop assesses the progress of the action loop, then the latter loop is dealing with the psychological equivalent of velocity, the first derivative of distance over time. To the extent this analogy is meaningful, the perceptual input to this loop should be the first derivative over time of the input used by the action loop.

We don't believe this input creates affect by itself, because a given rate of progress has different affective consequences under different circumstances. As in any feedback system, this input is compared against a reference value (cf. Frijda, 1986, 1988). In this case, the reference is an acceptable or desired rate of behavioral discrepancy reduction. As in other feedback loops, the comparison checks for a deviation from the standard. If there is one, the output function changes.

We suggest that the result of the comparison process at the heart of this loop (the error signal generated by the comparator) is manifest phenomenologically in two forms: one is a hazy and nonverbal sense of expectancy—confidence or doubt; the other is affect, feeling—a sense of positiveness or negativness.

H. RESEARCH EVIDENCE

At least a little evidence has accumulated to support the idea that affect originates in a velocity function. Hsue and Abelson (1991), who came independently to this idea, reported two studies of velocity and satisfaction. In one, subjects read descriptions of paired hypothetical scenarios and indicated which one they would find more satisfying. For example, they chose whether they would be more satisfied if their class standing had gone from the 30th percentile to the 70th over the past 6 weeks or if it had done so over the past 3 weeks.

Some comparisons were of positive outcomes; some negative. Given positive outcomes, subjects preferred *improving* to a high outcome over a *constant* high outcome; they preferred a fast velocity over a slow one; and they preferred fast brief changes to slower larger changes. When the change was negative (e.g., salaries got worse), subjects preferred a constant low salary to a salary that started high and fell to the same low level; they preferred slow falls to fast falls; and they preferred large slow falls to small fast falls.

We conducted a study that conceptually replicates aspects of these findings, but with an event that was personally experienced rather than hypothetical (Lawrence, Carver, & Scheier, 1999). We manipulated success feedback on an ambiguous task over an extended period. The patterns of feedback converged such that block 6 was identical for all subjects at 50% correct. Subjects in a neutral condition had 50% on the first and last block, and 50% average across all blocks. Others had positive change in performance, starting poorly and gradually improving. Others had negative change, starting well and gradually worsening. All rated their mood before starting and again after block 6 (which they did not know ended the session). Those whose performances were improving reported better moods; those whose performances were deteriorating reported worse moods, compared to those with a constant performance.

Another study that appears to bear on this view of affect, although not having this purpose in mind, was reported by Brunstein (1993). It examined subjective well being among college students over the course of an academic term as a function of several perceptions, including perception of progress toward goals. Of particular interest at present, progress at each measurement point was strongly correlated with concurrent well being.

C. CRUISE CONTROL MODEL

Ours is essentially a "cruise control" model of affect. That is, the system we've postulated functions much the same as the cruise control on a car. If you're going too slowly toward some goal in your behavior, negative affect arises. You respond by putting more effort into your action, trying to speed

up. If you're going faster than you need to, positive affect arises and you coast. A car's cruise control is very similar. You come to a hill, which slows you down. Your cruise control responds by feeding the engine more gas to bring the speed back up. If you come across the crest of a hill and roll downhill too fast, the system pulls back on the gas and drags the speed back down.

The analogy is intriguing in part because it concerns an electromechanical regulation of the very quality we believe the affect system is regulating: velocity. It is also intriguing to realize that this analogy incorporates a similar asymmetry in the consequences of deviating from the set point. That is, both in your car's cruise control system and in your behavior, going too slow requires investment of greater effort and resources. Going too fast does not. It requires only pulling back on resources. That is, your cruise control doesn't apply your brakes; it just cuts back on the gasoline. In this way it permits you to coast back to your velocity set point. In the same fashion, you don't respond to positive affect by trying to make it go away, but just by easing off.

Does positive affect lead people to withdraw effort? There is a little information on this, but not much. Melton (1995) found that people in a good mood performed worse than control subjects on syllogisms. A variety of ancillary data led him to the conclusion that the people in good moods did worse because they were expending less effort. To us, this looks like coasting.

D. AFFECT FROM DISCREPANCY-ENLARGING LOOPS

When we began this section we said we would restrict ourselves at first to discrepancy-reducing loops. Thus far we've done that, dealing only with issues that arise in the context of approach. Now we turn to attempts to distance oneself from a point of comparison, attempts to not be or not do, discrepancy-enlarging loops.

It should be apparent from our earlier discussion that behavior toward avoidance goals is just as intelligible as behavior toward approach goals. But what about the affective accompaniments to avoidance loops? The affect theory described here rests on the idea that positive affect results when a behavioral system is making rapid progress in *doing what it is organized to do*. The systems considered thus far are organized to close discrepancies. There's no obvious reason, however, why the principle shouldn't apply just as well to systems with the opposite purpose. If the system is making rapid progress doing what it's organized to do, the result should be positive affect. If the system is doing poorly at what it's organized to do, the result should be negative affect.

That much would seem to be fully comparable across the two types of systems. We see, however, a difference in the affect qualities involved (see




FIGURE 5 Two sorts of affect-creating systems and the affective dimensions we believe arise from the functioning of each. Discrepancy-reducing systems are presumed to yield affective qualities of sadness or depression when progress is below standard and happiness or elation when progress is above standard. Discrepancy-enlarging systems are presumed to yield anxiety when progress is below standard and relief or contentment when progress is above standard. *Note.* From *On the Self-Regulation of Behavior*, by C. S. Carver and M. F. Scheier, 1998, New York: Cambridge University Press. Copyright 1998 by Cambridge University Press. Reprinted with permission.

Figure 5). In each case there's a positive pole and a negative pole, but the positives aren't quite the same, nor are the negatives.

Our view of this difference derives partly from the insights of Higgins and his colleagues (Higgins, 1987, 1996). Following their lead, we suggest that the affect dimension relating to discrepancy reducing loops is (in its purest form) the dimension that runs from depression to elation. The affect dimension that relates to discrepancy-enlarging loops is (in its purest form) the dimension from anxiety to relief or contentment. As Higgins and his colleagues note, dejection-related and agitation-related affect may take several forms, but these two dimensions capture the core qualities behind them. The connections drawn in Figure 5 between affect quality and type of system are compatible not just with the Higgins model, but also with certain other theories. For example, Roseman (1984, p. 31) has argued that joy and sadness are related to appetitive (moving-toward) motives, whereas relief and distress are related to aversive (moving-away-from) motives.

E. MERGING AFFECT AND ACTION

How does the mechanism creating affect influence *action*? A more basic question (which takes us to the same end) is this: We've treated affect as the error signal of a feedback loop, but what's the *output function* of that loop? If the input function is a perception of rate of progress, the output function must be an *adjustment* in rate of progress.

Some adjustments are straightforward—go faster. Sometimes it's less so. The rates of many "behaviors" (higher order activities) aren't defined by the literal pace of physical action. Rather, they're defined in terms of

choices among potential actions, or even potential *programs* of action. For example, increasing your rate of progress on a reading assignment may mean choosing to spend a weekend working rather than playing. Increasing your rate of kindness means choosing to do an action that reflects that value. Thus, adjustment in rate must often be translated into other terms, such as concentration or reallocation of time and effort.

It should be apparent, however, that the action system and the rate system work in concert with one another. Both are involved in the flow of action. They influence different *aspects* of the action, but both always are involved.

It's interesting that the functions that we've just described are roughly comparable to two functions typically ascribed to motivation. In effect, we seem to have arrived at saying that the action loop handles most of what is sometimes called the *directional* function of motivation (the choice of an action from among many options, keeping an action on the track intended), and that the affect loop is handling the *intensity* function of motivation (the vigor, enthusiasm, effort, concentration, or thoroughness with which the action is pursued). Our linking of affect with the intensity aspect of motivation is a consequence of structural assumptions we began with, rather than a principled decision. However, this link is certainly consistent with statements of many theorists who have emphasized the intimate connection between emotion and motivation.

This view of the nature and origin of affect raises a number of issues—far more than we can address here. We limit ourselves to a few of the more provocative and interesting ones.

F. IS THIS REALLY A FEEDBACK SYSTEM?

Our view on affect is that it results from a comparison process in a feedback loop. This view has a counterintuitive implication. If affect is created the way we say it is, it is a signal that the rate of progress isn't right and should be adjusted. This implies that although the organism tries to minimize pain, it does *not*, in general, try to maximize pleasure.

Minimizing pain is straightforward. Negative feelings reflect a negative discrepancy in rate, indicating a problem. Things aren't moving forward fast enough. The normal response is to try harder. If this happens, the negative affect ceases to exist. Thus, people try to minimize pain.

Maximizing pleasure is trickier. Positive feelings reflect a positive discrepancy in rate. This means things are going better than they need to, and the experience feels good. To a system whose goal is controlling rate, however, discrepancies are to be reduced. If there really is a feedback loop of the sort we have proposed, neither negative *nor* positive affect is a state the system wants to see. Either quality of affect (either deviation from the

standard) would represent an "error" and lead to changes in output that would reduce it.

If this is truly a feedback system, an overshoot of the reference value should lead to a self-corrective attempt to return to the reference value. Put more concretely, this view argues that people who exceed the desired rate of progress will slow subsequent effort in this domain. They are likely to "coast." The result in subjective experience would be that the positive affect from the overshoot is not sustained for very long.

Why should there be a natural tendency that would cause positive affect to be short lived? A plausible basis lies in the idea that behavior is hierarchically organized and has multiple current concerns. People typically are working toward several goals more or less simultaneously. To the extent that movement toward goal attainment is more rapid than expected in one domain, it lets the person shift effort toward strivings in another domain at no cost. To continue an unnecessarily rapid pace in the first domain may increase positive affect, but by diverting efforts from other goals, it may increase the potential for negative affect in other domains.

G. SHIFTS IN STANDARDS

Reference values in this system can differ across categories of behavior, and they can shift through time and experience. That is, as people accumulate experience in a given domain, adjustments can occur in the pacing they expect and demand of themselves. There is a recentering of the system around the past experience by changing the reference value.

Sometimes the adjustment is downward, toward less stringent pacing. One consequence of this is a more favorable balance of positive to negative affect across a given time span. In other cases, the adjustment is upward. This has the side effect of decreasing the potential for positive affect and increasing the potential for negative affect.

These changes don't happen quickly or abruptly. Shifting the reference value downward is not the first response when you have trouble keeping up a demanding pace. First you try harder to keep up. Only more gradually, if this fails, does the standard shift to accommodate. Similarly, an upward shift in reference value is not the immediate response when your rate exceeds the standard. The more typical response is to coast for a while. Only when the overshoot is frequent does the standard shift upward.

We believe that such adjustments in standard occur automatically, but slowly. Such adjustments *themselves* appear to reflect a self-corrective feedback process, as the person reacts to insufficient challenge by taking on a more demanding pace, and reacts to too much challenge by scaling back (see Figure 6). This feedback process is slower than the ones we've focused on thus far. Rather, there is a gradually accumulating shift.

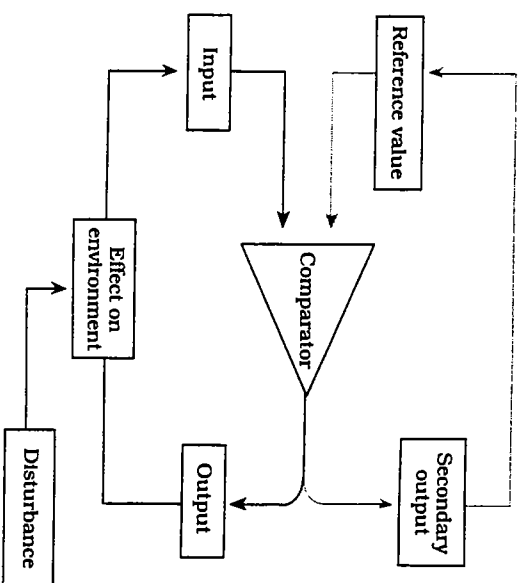


FIGURE 6 A feedback loop (in this case, the postulated velocity loop) acts to create change in the input function to shift it toward the reference value. Sometimes an additional process is in place as well (gray lines), which adjusts the reference value in the direction of the input. This additional process is presumed to be weaker or slower; thus, the reference value is stable relative to the input value. *Note:* From *On the Self-Regulation of Behavior*, by C. S. Carter and M. F. Scheier, 1998, New York: Cambridge University Press. Copyright 1998 by Cambridge University Press. Reprinted with permission.

It is also of interest (and once again counterintuitive) that these shifts in reference value (and the resultant effects on affect) imply a mechanism within the organism that functions to actively prevent the too frequent occurrence of positive feeling, as well as the too frequent occurrence of negative feeling. That is, the (bidirectional) shifting of the rate criterion over time would tend to control pacing of behavior in such a way that affect continues to vary in both directions around neutral, roughly the same as it had before.

Such an arrangement for changing the standard thus would not result in maximization of pleasure and minimization of pain. Rather, the affective consequence would be that the person experiences more or less the same range of variation in affective experience over extended periods of time and circumstances (cf. Myers & Diener, 1995). The organization as a whole would function as a gyroscope serving to keep us floating along within the framework of the affective reality we're familiar with. It would provide for a continuous recalibration of the feeling system across changes in situation. To use a different image, it would repeatedly shift the balance point

of a psychic lecturer-totter, so that rocking in both directions remains possible.

II. COMPARISON WITH BIOLOGICAL MODELS OF BASES OF AFFECT

It is useful to compare this model with the group of biologically focused theories we mentioned earlier. The theories are quite similar to one another in many ways, but in other ways they differ. These theories all incorporate the idea that two systems (or more) are involved in the regulation of behavior. Many assume further that the two systems underlie affect. In situations with cues of impending reward, the activity of the approach system creates positive feelings. In situations with cues of impending punishment, the avoidance system creates feelings of anxiety.

Data from a variety of sources fit this picture. Of particular interest is work by Davidson and collaborators, involving EEG recordings assessing changes in activation in response to affective inducing stimuli. Among the findings are these: Subjects exposed to films inducing fear and disgust (Davidson, Ekman, Saron, Senulis, & Friesen, 1990) and confronted with possible punishment (Sobotka, Davidson, & Senulis, 1992) show elevations in *right* frontal activation. In contrast, subjects with a chance to obtain reward (Sobotka et al., 1992), subjects presented with positive emotional adjectives (Cacioppo & Petty, 1980), and smiling 10-month-olds viewing their approaching mothers (Fox & Davidson, 1988) show elevations in *left* frontal activation. From findings such as these, Davidson (1992a, 1992b) concluded that neural substrates for approach and withdrawal systems (and thus positive and negative affect) are located in the left and right frontal areas of the cortex, respectively.

The logic of these models thus far resembles the logic of our model. At this point, however, theories diverge. The question on which they diverge concerns the regulatory processes involved in, and affects that result from, *failure to attain reward* and *successful avoidance of punishment*. Gray (1987b, 1990) holds that the avoidance system is engaged by cues of punishment and cues of frustrative nonreward. It thus is responsible for negative feelings in response to either of these types of cues. Similarly, he holds that the approach system is engaged by cues of reward or cues of escape from (or avoidance of) punishment. It thus is responsible for positive feelings in response to either of these types of cues. In his view, then, each system creates affect of one hedonic tone (positive in one case, negative in the other), regardless of its source (see Figure 7). This view is consistent with a picture of two unipolar affective dimensions, each linked to a distinct behavioral system. A similar position has been taken by Lang, Bradley, and Cuthbert (1992).

Our position is different. It argues for an approach system and an avoidance system, in which—in each case—affect is a product of either doing well or doing poorly. Thus, it implies two bipolar dimensions: one tied to approach, the other to avoidance (Figure 7). We think the frustration and eventual depression that result from failure to attain desired goals involve the approach system (for similar predictions, see Clark, Watson, & Mineka, 1994, p. 107; Cloninger 1988, p. 103). Sadness and depression involve *reduced* activity in the approach system, as the pursuit of goals diminishes. A parallel line of reasoning suggests that relief, contentment, tranquility, and serenity relate to the avoidance rather than the approach system, reflecting low levels of activity in that system.

Less information exists about neurophysiological bases of these affects than about anxiety and happiness. With regard to relief–tranquility, we know of no data at all. With respect to depression, limited evidence exists. Henriques and Davidson (1991) found that clinically depressed persons had less activation in left frontal areas than nondepressed controls. In contrast, there was no evidence of a difference in right frontal activation. This pattern has since been replicated by Allen, Iacono, Depue, and Arbisi (1993). Recall that Davidson viewed baseline measures as representing susceptibility rather than ambient affect per se. Thus, this finding tenta-




FIGURE 7 Gray's view of affect (horizontal groupings) ties positive affects to the effects of a behavioral activation system (BAS), as results of occurrence of reward and avoidance of punishment. It ties negative affects to the effects of a behavioral inhibition system (BIS), as results of frustrative nonreward and occurrence of punishment. Our view (vertical groupings, as in Figure 5) ties the dimension of elation–depression to an approach system and the dimension of anxiety–relief to an avoidance system, each of which thus has properties somewhat different from those assumed by Gray. *Note.* From *On the Self-Regulation of Behavior*, by C. S. Carter and M. F. Scheiner, 1998, New York: Cambridge University Press. Copyright 1998 by Cambridge University Press. Reprinted with permission.

tively seems to suggest that depressed persons are vulnerable to depression through deficits in their approach system. This set of findings seems quite compatible with our position.

IV. CONFIDENCE AND DOUBT, PERSISTENCE AND GIVING UP

In describing the genesis of affect, we suggested that one mechanism yields two subjective readouts: affect, and a hazy sense of confidence versus doubt. We turn now to a consideration of confidence and doubt—expectancies for the immediate future. We focus here on the *consequences* of this sense of confidence and doubt.

We've often suggested that when people experience adversity in trying to move toward their goals, they periodically experience an interruption of efforts, to assess in a more deliberative way the likelihood of a successful outcome (e.g., Carver & Scheier, 1981, 1990). In effect, people suspend the behavioral stream, step outside it, and evaluate in a more deliberated way than occurs while acting. This may happen once or often. It may be brief or it may take a long time. In this assessment, people presumably depend heavily on memories of prior outcomes in similar situations. They may also consider such things as additional resources they might bring to bear (cf. Lazarus, 1966) or alternative approaches to the problem. People use social comparison information (Wills, 1981; Wood, 1989) and attributional analysis of prior events (Wong & Weiner, 1981).

How do these thoughts influence the expectancies that emerge? In some cases, when people retrieve "chronic" expectancies from memory, the information already is expectancies, summaries of products of previous behavior. For some cases, however, people bring to mind possibilities for the situation's evolution. For these possibilities to influence expectancies, their consequences must be evaluated. We suggest they are briefly played through mentally as behavioral scenarios (cf. Taylor & Pham, 1996). This should lead to conclusions that influence the expectancy: ("If I try approaching it this way instead of that way, it should work better.") "This is the only thing I can see to do, and it will just make the situation worse.")

It seems reasonable that this mental simulation engages the same mechanism as handles the affect creation process during actual overt behavior. When your progress is temporarily stalled, playing through a scenario that is confident and optimistic yields a higher rate of progress than is currently being experienced. The affect loop thus yields a more optimistic outcome assessment than is being derived from current action. If the scenario is negative and hopeless, it indicates a further reduction in progress and the loop yields further doubt.

A. ENGAGEMENT VERSUS GIVING UP

The expectancies that result, whatever their source, are reflected in behavior (Figure 8). If expectations are for a successful outcome, the person returns to effort toward the goal. If doubts are strong enough, the result is an impetus to disengage from further effort and potentially from the goal itself (Carver & Scheier, 1981, 1990, 1998, 1999; see also Klingler, 1975; Kukla, 1972; Wortman & Brehm, 1975). This theme—divergence in behavioral response as a function of expectancies—is an important one, applying to a surprisingly broad range of literatures (see Chapter 11 of Carver & Scheier, 1998).

Sometimes the disengagement that follows from doubt is overt, but disengagement sometimes takes the form of mental disengagement—off-task thinking, daydreaming, and so on. Although this sometimes can be useful (self-distraction from a feared stimulus may permit anxiety to abate), it also can create problems. If there is time pressure, mental disengagement can impair performance, as time is spent on task-irrelevant thoughts. Consistent with this, interactions between self-focus and expectancies have been shown for measures of performance (Carver, Peterson, Follansbee, & Scheier, 1983; Carver & Scheier, 1982).

Often mental disengagement can not be sustained, because situational cues force a re confrontation of the task. In such cases, the result is a phenomenology of repetitive negative rumination, which often focuses on




FIGURE 8 Flow-chart depiction of self-regulatory possibilities, indicating that action sometimes continues unimpeded toward goal attainment, that obstacles to goal attainment sometimes induce a sequence of evaluation and decision making, and that if expectancies for eventual success are sufficiently unfavorable, the person may disengage from further effort. *Note.* From *On the Self-Regulation of Behavior*, by C. S. Carver and M. F. Scheier, 1998. New York: Cambridge University Press. Copyright 1998 by Cambridge University Press. Reprinted with permission.

self-doubt and perceptions of inadequacy. This cycle is both unpleasant and performance impairing.

A number of writers—both from an earlier tradition of cognitive-attentional theories of anxiety (e.g., Sarason, 1975; Wine, 1971, 1980) and some more recent (Ingram, 1990)—have equated this phenomenology of negative rumination with the term “self-focus.” We would argue that this label is misleading (see also Pyszczynski, Greenberg, Hamilton, & Nix, 1991). Why? Because self-focus does not always produce interference. As described earlier, self-focus is in many cases associated with *task* focus, as the experience of attention inward engages the feedback process underlying effort. Indeed, among confident subjects in the studies just mentioned, this is what occurred even under conditions of adversity. Self-focus led to better task effort, as manifest both in overt behavior and in contents of consciousness. Only among doubtful subjects did self-focus lead to performance impairment or to negative rumination.

B. IS DISENGAGEMENT GOOD OR BAD?

Is the disengagement tendency good or bad? The answer is that it is both and neither. On the one hand, disengagement (at some level, at least) is an absolute necessity. Disengagement is a natural and indispensable part of self-regulation. If we are ever to turn away from efforts at unattainable goals, if we are ever to back out of blind alleys, we must be able to disengage, to give up and start over somewhere else.

The importance of disengagement is particularly obvious with regard to concrete, low-level goals: We must be able to remove ourselves from literal blind alleys and wrong streets, give up plans that have become disrupted by unexpected events, even spend the night in the wrong city if we've missed the last plane home. The tendency is also important, however, with regard to more abstract and higher-level goals. A vast literature attests to the importance of disengaging and moving on with life after the loss of close relationships (e.g., Orbach, 1992; Stroebe, Stroebe, & Hansson, 1993; Weiss, 1988). People sometimes must even be willing to give up values that are deeply embedded in the self, if those values create too much conflict and distress in their lives.

As with most processes in self-regulation, however, the choice between continued effort and giving up presents opportunities for things to go awry. It's possible to stop trying prematurely, thereby creating potentially serious problems for oneself (Carver & Scheier, 1998). It's also possible to hold onto goals too long and prevent oneself from taking adaptive steps toward new goals. However, both continued effort and giving up are necessary parts of the experience of adaptive self-regulation. Each plays an important role in the flow of behavior.

C. HIERARCHICALITY AND IMPORTANCE CAN IMPEDE DISENGAGEMENT

Disengagement sometimes is precluded by situational constraints. However, another, broader aspect of this problem stems from the idea that behavior is hierarchically organized, with goals increasingly important higher in the hierarchy and thus harder to disengage from.

Presumably disengaging from concrete values is often easy. Lower order goals vary, however, in how closely they are linked to values at a higher level, and thus how important they are. To disengage from low-level goals that are tightly linked to higher-level goals causes discrepancy enlargement at the higher level. These higher order qualities are important, even central to one's life. One cannot disengage from them, or disregard them, or tolerate large discrepancies between them and current reality without reorganizing one's value system (Greenwald, 1980; Kelly, 1955; McIntosh & Martin, 1992; Millar, Tesser, & Millar, 1988). In such a case, disengagement from even very concrete behavioral goals can be quite difficult.

Now recall again the affective consequences of being in this situation. The desire to disengage was prompted by unfavorable expectancies. These expectancies are paralleled by negative affect. In this situation, then, the person experiences negative feelings (because of an inability to make progress toward the goal) and is unable to do anything about the feelings (because of an inability to give up). The person simply stew in the feelings that arise from irreconcilable discrepancies. This kind of situation—commitment to unattainable goals—seems a sure prescription for distress.

D. WATERSHEDS, DISJUNCTIONS, AND BIFURCATIONS AMONG RESPONSES

An issue that bears some further mention is the divergence of the behavioral and cognitive responses to favorable versus unfavorable expectancies that is part of this model. We've long argued for a psychological watershed among responses to adversity (Carver & Scheier, 1981). One set of responses consists of continued comparisons between present state and goal, and continued efforts. The other set consists of disengagement from comparisons and quitting. Just as rainwater falling on a mountain ridge ultimately flows to one side of the ridge or the other, so do behaviors ultimately flow to one of these sets or the other.

Our initial reason for taking this position stemmed largely from the several demonstrations that self-focused attention creates diverging effects on both information seeking and behavior, as a function of expectancies of success (Figure 9). We aren't the only ones to have emphasized a disjunction among responses, however. A number of others have done so, for reasons of their own.




FIGURE 9 Persistence at an insoluble problem as a function of performance expectancy regarding the target task and self-directed attention. *Note:* Adapted by combining data from Experiments 1 and 2 from "Reassessment and Giving Up: The Interactive Role of Self-Directed Attention and Outcome Expectancy," by C. S. Carver, P. H. Blaney, & M. E. Scheier, 1979, *Journal of Personality and Social Psychology*, 37, pp. 1859-1870. Copyright 1979 by the American Psychological Association. Adapted with permission.

An early model that emphasized the idea of a disjunction in behavior was proposed by Kukla (1972). Another such model is the reluctance-helplessness integration of Wortman and Brehm (1975): the argument that threats to control produce attempts to regain control and that perceptions of loss of control produce helplessness. Brehm and his collaborators (Brehm & Self, 1989; Wright & Brehm, 1989) have more recently developed an approach to task engagement that resembles that of Kukla (1972), although their way of approaching the description of the problem is somewhat different. Not all theories about persistence and giving up yield this dichotomy among responses. The fact that some do, however, is interesting. It becomes more so a little bit later on.

V. DYNAMIC SYSTEMS AND HUMAN BEHAVIOR

Recent years have seen the emergence in psychology of new (or at least newly prominent) ideas about how to conceptualize natural systems. Several labels attach to these ideas: chaos; dynamic systems theory; complexity; catastrophe theory. A number of introductions to various aspects of this body of thought have been written, some of which include applications to psychology (e.g., Brown, 1995; Gleick, 1987; Thelen & Smith, 1994; Vallacher & Nowak, 1994, 1997; Waldrop, 1992). In this section we sketch some themes that are central to this way of thinking and indicate places

where we think the themes apply meaningfully to subjects of our own interest.

Dynamic systems theory (or chaos theory) is deterministic (despite the contrary implication of the word chaos). It holds that the behavior of a system reflects the forces operating on (and within) it. It also emphasizes that the behavior of a complex system over anything but a brief time is very hard to predict. Why?

A. NONLINEARITY

One reason is that the system's behavior may be influenced by the forces operating on and within it in nonlinear ways. Thus, the behavior of the system—even though highly determined—can appear random. This determinism in principle but unpredictability in practice underlies the label *chaotic*.

Many people are used to thinking of relationships between variables as linear. Dynamic-systems thinking asserts that many relationships are not linear. Familiar examples of nonlinear relationships are step functions (ice turning to water and water turning to steam as temperature increases), threshold functions, and floor and ceiling effects. Other examples of nonlinearity are interactions (Figure 10). In an interaction, the effect of one predictor on the outcome differs as a function of the level of a second predictor. Thus the effect of the first predictor on the outcome is not linear. The second predictor is thereby acting as a *control parameter*—a factor "that hold[s] potential for changing the intrinsic dynamics of a system" (Vallacher & Nowak, 1997).




FIGURE 10 Interactions indicate that the effect of one variable differs as a function of the level of another variable. On the left is the interaction between self-focus and expectancy that was shown in Figure 9. This indicates that the effect of self-focus is nonlinear (right panel). Its impact reverses at some point along the distribution of the variable confidence—doubt. *Note:* From "Themes and Issues in the Self-Regulation of Behavior" by C. S. Carver & M. E. Scheier, in R. S. Wyer, Jr. (Ed.), *Advances in Social Cognition*, 1999, Mahwah, NJ: Erlbaum. Copyright 1999 by Lawrence Erlbaum Associates. Reprinted with permission.

