Personality and Social Psychology Review
2002, Vol. 6, No. 4, 304–315

Control Processes and Self-Organization as Complementary Principles Underlying Behavior

Charles S. Carver
Department of Psychology
University of Miami

Michael F. Scheier
Deparment of Psychology
Carnegie Mellon University

This article addresses the convergence and complementarity between self-regulatory control-process models of behavior and dynamic systems models. The control-process view holds that people have a goal in mind and try to move toward it (or away from it), monitoring the extent to which a discrepancy remains between the goal and one's present state and taking steps to reduce the discrepancy (or enlarge it). Dynamic systems models tend to emphasize a bottom-up self-organization process, in which a coherence arises from among many simultaneous influences, moving the system toward attractors and away from repellers. We suggest that these differences in emphasis reflect two facets of a more complex reality involving both types of processes. Discussion focuses on how self-organization may occur within constituent elements of a feedback system—the input function, the output function, and goal values being used by the system—and how feedback processes themselves can reflect self-organizing tendencies.

A new (or, at least, newly visible) family of ideas has arisen for consideration as potentially useful to personality and social psychologists. People in our disciplines are increasingly exploring how ideas from bodies of thought such as dynamic systems, connectionism, and catastrophe theory may apply to phenomena of their interest. Sometimes these ideas are treated as being revolutionary; other discussions treat them instead as being complementary to ideas already in use (e.g., Carver, 1997; Carver & Scheier, 1998). Indeed, several authors have noted that aspects of these new ideas resonate well to themes that have long been embedded in psychology under different names (e.g., Read, Vanman, & Miller, 1997; Vallacher & Nowak, 1997).[1]

In this article we proceed from the position that the bodies of thought previously listed provide new angles on what is known, provide conceptual bases for new methods and new questions, but do not suggest a need to abandon pre-existing ways of thinking. They do, however, induce us to re-examine many phenomena we have taken for granted. They may thereby lead us to somewhat different views of things we thought we already understood.

This article focuses on some implications of dynamic systems thinking for understanding the self-regulation of action. We have long taken a particular theoretical viewpoint on this topic. As it happens, a lot of what we believe about the occurrence of action can also be described in the language of dynamic systems theory. There are, however, differences of emphasis between our usual view and the dynamic-systems view, which can make the two seem competitors, rather than candidates for merger. This article focuses on one difference in emphasis, while still highlighting what we regard as the fundamental compatibility and complementarity between the models.

Preparation of this article was facilitated by NCI grants CA64710, CA78995, CA84944, and CA84944, and by NHLBI grants HL65111 and HL65112. We thank Sheri Johnson, Robin Vallacher, and Andrzej Nowak for comments on earlier drafts.

Requests for reprints should be sent to Charles S. Carver, Department of Psychology, University of Miami, Coral Gables, FL 33124–2070. E-mail: Ccarver@miami.edu

[1]We portray here a contrast between "older" and "newer" models, but there is a question as to whether newness is really the issue. The real issue may instead be depth of conceptualization. Most theories in personality and social psychology aim at the surface of the phenomena under study and are framed in constructs appropriate to that surface description. The ideas under discussion in this special issue address deeper underlying dynamics, and correspondingly are framed in more abstract terms. Indeed, it is often noted that dynamics patterns of particular kinds emerge independent of the nature of the system in which they are embodied (e.g., Amazeen, Amazeen, & Turvey, 1998). This is a claim that was also made, of course, about general systems theory (von Bertalanffy, 1968), which many would regard as an "older" model—albeit a very abstract model.

Some Old Ideas: Feedback Control of Behavior

We begin by briefly outlining the "old" concepts we have used throughout our careers. These ideas trace at least to Cannon's (1932) description of homeostatic processes and Wiener's (1948) description of cybernetics—communication and control processes that occur in both living and artificial systems.[2]

The logic of cybernetic feedback control is an old model in the sense that the principle has been understood for a long time. It is also old in having been used for a long time in psychology, principally areas of psychology with a physiological focus. Whether it should be seen as an old model in personality and social psychology is less clear. On the one hand, it was argued nearly a half century ago that the cybernetic principle applies to behaviors of interest to personality and social psychologists (e.g., MacKay, 1956, 1963; Miller, Galanter, & Pribram, 1960; Powers, 1973). On the other hand, explicit use of the principle in these areas has been limited.

Nonetheless, there are some indications that the essence of cybernetic logic has seeped into mainstream personality–social psychology. Consider the recent re-emergence of the goal construct (in various forms) as a focus for theory (e.g., Austin & Vancouver, 1996; Carver & Scheier, 1998, 1999a; Elliott & Dweck, 1988; Higgins, 1987, 1996; Markus & Nurius, 1986; Miller & Read, 1987; Pervin, 1989; Ryan, Sheldon, Kasser, & Deci, 1996). The goal construct is often invoked by itself, with little mention (if any) of what processes are entailed in the pursuit and attainment of a goal. However, the goal construct actually provides an ideal entry point into the logic of cybernetic self-regulation.

Feedback Processes

A goal is a reference point (a target) toward which a person wishes or intends to move. In the cybernetic view, moving toward a goal requires identifying one's location on the relevant variable (a process that is typically termed a perceptual input function), determining where that location is with respect to the goal (by means of a comparison between input and goal), and making changes to diminish the gap between present location and goal (a process that is typically termed an output function). These processes together form a discrepancy reducing feedback loop (Figure 1).[3]

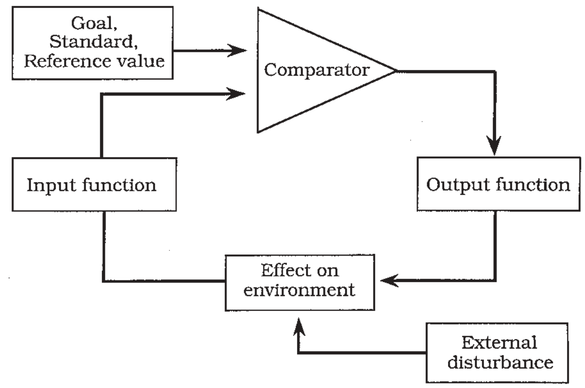

Figure 1. Schematic depiction of a typical feedback loop, the basic unit of cybernetic control. In a discrepancy reducing loop, a sensed input value is compared to a reference value, and adjustments are made in an output function (if necessary) to shift the sensed value closer to the standard, thereby countering pre-existing deviations and deviations introduced by external disturbances.

Although discrepancy reducing loops are the most common feedback processes in living systems, there is another kind of loop that enlarges discrepancies between present states and the reference value. To avoid confusion, we will label the reference values of such loops *anti-goals* (Carver & Scheier, 1998, 1999a). Two examples at the level of abstraction of personality–social psychology are the feared self (Markus & Nurius, 1986) and the values held by disliked social groups (Carver & Humphries, 1981). Self-regulation regarding an anti-goal entails escaping, withdrawing from, or avoiding it. The actions of discrepancy enlarging loops are often constrained by discrepancy reducing loops. Thus, a person can often avoid an anti-goal by approaching a goal that is at some remove from the anti-goal (cf. Carver, Lawrence, & Scheier, 1999; Elliott, 1999; Higgins, 1996).

The view of behavior based on feedback principles describes one mechanism that would account for how goals are attained by actions. The mechanism behind this view presumes the existence of a capacity to represent a goal, a channel of informational input, and pathways to exert influence on present conditions. These elements are simple, but it is important to note that they represent a kind of structure as well as function. To in-

[2]Wiener coined the term *cybernetic* (derived from the Greek word for "steersman") to refer to the logic underlying these process. Today most people probably think this term refers to computers or the internet, but its initial referent was the structure of control processes. The fact that cybernetic principles are embodied in computers, among other artificial systems, has led some to mistakenly infer that the cybernetic model derives from, and thus depends on, the computer as a root metaphor. Not so. Indeed, the argument has been made that the opposite is true—that computer design initially followed from an understanding of control processes in living systems (Powers, 1978).

[3]Reducing of sensed discrepancies has long had a prominent role in social psychology, in terms of behavioral conformity to norms (Asch, 1955) and in theories of cognitive consistency (Festinger, 1957; Heider, 1946; Leckey, 1945; for a more recent example see Higgins, 1996). Mechanisms invoked for such phenomena have varied, sometimes being quite compatible with feedback models, sometimes less so.

voke these principles is to suggest that some physical elements engage in activities that correspond to the component functions of the loop.

We note in passing that a single feedback loop can manage self-regulation of a single quality, but it is too simple by itself to create physical action. That problem was addressed (Powers, 1973) in the idea that control of abstract behavior involves a hierarchy of loops, such that the output at a given level consists of re-setting reference values at the next lower level (for related discussions, see Carver & Scheier, 1998, 1999; Vallacher & Wegner, 1987). This permits abstract qualities to dissociate into constituent elements iteratively through several layers, resulting eventually in the patterns of muscle contractions that yield physical action (cf. Rosenbaum, Carlson, & Gilmore, 2001).

Neurobiological Motive Systems

We believe these ideas about discrepancy reduction and discrepancy enlargement stand on their own as useful principles regarding behavior. It is interesting, however, that the general sense of these ideas is echoed in a family of neurobiological motive theories that currently are receiving considerable attention from personality and social psychologists (e.g., Cloninger, 1998; Davidson, 1998; Davidson, Jackson, & Kalin, 2000; Depue & Collins, 1999; Depue & Zald, 1993; Fowles, 1993; Gray, 1994a, 1994b; Lang, 1995). Most of these theories assume the existence of one set of mechanisms that manage appetitive motivation and approach behavior, and a relatively separate set of mechanisms that manage aversive motivation and avoidance behavior.

There are obvious similarities between the overall effects of these two sets of mechanisms and the overall functions of the two classes of feedback loops described just before. In each case there is an approach function and an avoidance function, though the two lines of thought frame the sets of functions in different ways. These similarities seem worthy of further attention (Carver, 2001, in press; Carver, Sutton, & Scheier, 2000; Montant, 2000).

Priorities and Scheduling

One more issue must be raised here regarding this view of goal-directed action.[4] Most theorists assume that people normally have many goals in mind, at various levels of activation, which compete for attentional resources (Bargh, 1997; Cooper & Shallice, 2000; Norman & Shallice, 1986; Shallice, 1978; Simon, 1967). At a given moment, the goals presumably form a priority queue, with the goal that currently has the highest priority being the one that is actively pursued at that moment. Because priorities fluctuate with changing circumstances, people shift the focus of their behavior repeatedly. They do one thing for a while, then something else (cf. Atkinson & Birch, 1970; Murray, 1938).

The idea of a prioritization process among multiple goals raises many issues (Carver, in press). At present, however, we wish to make only this simple point: people typically have many goals in mind, which are manifested in overt action sequentially, through some sort of scheduling mechanism.

[4]In other contexts, other issues are also important. One of them which is tangential to present concerns is that confidence versus doubt influences the intensity of efforts at goal attainment (or anti-goal avoidance). When people are confident of eventual success, they try hard to succeed. When people are doubtful, they try less hard. Given enough doubt, they give up and stop trying at all. We have elsewhere explored the application of catastrophe theory to such phenomena (Carver & Scheier, 1998, 1999a).

Dynamic Systems, and Links To Feedback Models

The principle of feedback control just described is a dynamic model. However, it's not the kind of model most people associate with the label *dynamic systems*. The next section outlines a few of the basic ideas from the dynamic-systems view. Once some of those ideas are in place, we consider some of the relations between the dynamic-systems model and the feedback model (see also Carver & Scheier, 1998, chapters 14 and 17).

Attractors and Repellers

The general view identified with terms such as dynamic systems, complexity, and chaos holds that the behavior of a system reflects the multiplicity of forces operating on (and within) the system. Because we never know all the forces that are operating, and because the relations between and among forces are often nonlinear (Brown, 1995; Kelso, 1995; Vallacher & Nowak, 1997), the behavior of a complex system is hard to predict over anything but the very short term.

Dynamic-systems models focus on how systems change over time. The behavior of a system over time is often portrayed as a continuous tracing of a line in a two-dimensional space. Time is represented in the continuing trace, and other variables are embodied in the x and y axes. The term *phase space* refers to the set of possible states of the system. As the system changes states over time, it traces a *trajectory* in its phase space—a path of the successive states it occupies. When a system's behavior is observed over an extended period, it sometimes happens that certain regions of the phase are occupied often, others occasionally, and others never. An area of phase space the system occupies or approaches more frequently than others is called an *attractor.* An attractor exerts a kind of metaphorical pull on the system, bringing the system's behavior close to it.

Each attractor has a *basin*, its region of attraction. Trajectories that enter the basin move toward that attractor. A shift from one attractor to another is called a phase transition or phase change. In systems with more than one attractor, the system's trajectory typically approaches each one periodically but is never fully captured by any of them. Plotting the behavior of such a system over time shows a tendency to approach the various attractors, but often unpredictably. Shifts from one attractor to another may even seem random.

A phase space also contains regions called *repellers*, regions that are hardly ever occupied. Indeed, these regions seem to be actively avoided. That is, forcing the system near the focal point of a repeller would lead to a rapid escape from that region of phase space. Indeed, it can appear very much as though the forces that are in operation on and within the system normally prevent the system from ever entering that region.

Goals As Attractors

There are many ways to apply the concepts of phase space and attractor to human behavior. The application we make here is to suggest that goals represent attractors for behavior over time (Carver & Scheier, 1998; Pribram, 1991; Vallacher & Nowak, 1997). The fit of these two concepts to each other is near-perfect. People spend much of their time trying to keep their behavior in line with their goals. In the same way, dynamic systems spend much of their time near attractors. A similar parallel exists between anti-goals and repellers: both are actively avoided.

The description of a system's behavior over time in its phase space also has a link to the idea of prioritizing and scheduling among goals, which was briefly described earlier. The phase diagram of a system that is multi-stable (has several attractor regions) can be seen as a way of portraying the behavior of a person who has several goals, which are taken up and put aside in some sequence. A map of this person's behavior over time, approaching one goal, then another (while avoiding anti-goals), greatly resembles the movement of a system in a complex attractor landscape (Figure 2).

Dynamic-systems language brings along with it some labels for parameters of the system's behavior, which may also be useful in talking about goals. For example, one can describe (or specify) a system's degree of resistance to perturbation when it is in an attractor basin (essentially, the person's resistance to distraction when in pursuit of a particular goal), or the width of a given attractor basin (the ease with which the person is induced into pursuit of this goal when not presently doing so). The formal nature of the dynamic-systems model lends itself naturally to thinking in quantitative terms, which is desirable. On the other hand, precisely how to operationalize these parameters in behavioral terms poses a challenge to the researcher's creativity.

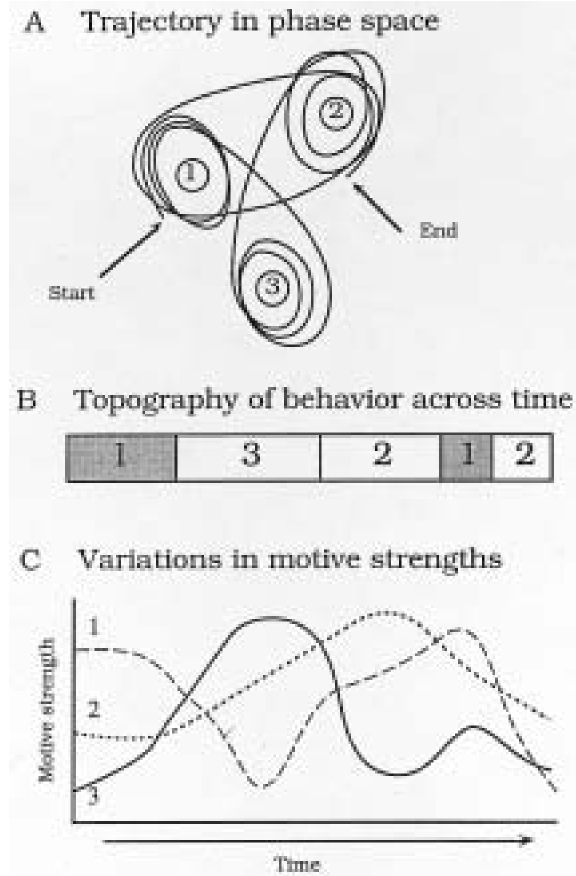

Figure 2. Three views of behavior shifting over time. (A) The trajectory of a system in phase space with 3 attractors, each shift representing a phase transition. (B) The topography of the observable behavior that occurs across the same span of time. (C) Variations in the strengths of the 3 motives that underlie and foster the behaviors, over the same span of time (cf. Atkinson & Birch, 1970). From C. S. Carver and M. F. Scheier, *On the self-regulation of behavior*, Copyright 1998, Cambridge University Press. Reprinted with permission.

A Difference in Emphasis

It should be clear from the foregoing that the dynamic systems view of behavior and the control-process view of behavior are in many ways complementary to one another. However, there are also some differences in emphasis in the ways these models have been applied to psychological phenomena. Recall that in describing the feedback-based view earlier, we said that goals are presumed to be represented within the nervous system. When a goal is focal, the process of moving toward it is presumed to occur top-down, guided by that representation and managed and controlled by some sort of executive intentional process.

In contrast to this, the dynamic-systems principle of regions of attraction does not rely on assumptions about top-down management, or even structure. Rather, attractors are said to arise from the intrinsic dynamics of the system as it operates in its world over an extended period of time. Complex systems are said to have a *self-organizing* character (Kelso, 1995; Prigogene & Stengers, 1984). The various forces interweave in ways that are not determined by any one of them alone, but rather by their mutual influences on each other. Patterns spontaneously emerge.

Self-Organization

The principle of self-organization is a key idea for a dynamic-systems picture of ongoing behavior. As such, it warrants closer examination. The principle of self-organization suggests that forces operating within and on the system converge to create an organization of sorts, an emergent set of properties that was not there before, and might not come to exist without the confluence of forces. No single force drives this emergence. It arises instead from a dynamic that involves multiple precipitators.

This can be a difficult argument to assimilate. To those unfamiliar with the principle (and especially those skeptical of it), the idea even seems reminiscent of a well-known cartoon from the *American Scientist* by Sidney Harris, in which a scientist portrays the key process of his theory with the phrase "and then a miracle occurs." The idea of self-organization can seem a little like that. This may be one reason why people sometimes question whether the attractor landscape is a meaningful conceptual entity or is instead just a metaphor—descriptively useful, but not explanatory (for more discussion see, e.g., Aslin, 1993, especially p. 394; Carver & Scheier, 1999b; Rosenbaum, 1998).[5]

It is perhaps easier to accept the notion of self-organization in inanimate systems than in living systems. Consider an example. The presence of heat, humidity, and low air pressure over warm water sometimes gives rise to thundershowers. Sometimes, however, what appear to be the same elements give rise instead to a highly organized phenomenon known as a hurricane. The hurricane is not created from the top down. It is emergent, a phenomenon that depends on the interweaving of a set of contributing influences. Further, what triggers the emergence of self-organization is hard to pinpoint. Even in what appear to be ideal circumstances, a hurricane sometimes fails to form.

In living systems, additional forces doubtlessly are at work beyond those operating on inanimate phenomena. Living systems incorporate pressures toward internal coherence of various sorts (see Nowak, Vallacher, Tesser, & Borkowski, 2000). Without pressures toward coherence at the physiological level, living systems would fairly quickly become non-living systems. Pressures toward coherence presumably also operate on the kinds of processes that become manifest as psychological phenomena. As psychological systems operate over time under that kind of constraint, they evolve and self-organize.

What about behavior? Does the principle of self-organization characterize human behavior? Indeed, what does it mean to say that behavior displays self-organization? It actually might mean several different things. Indeed, some apparent instances of self-organization can be found within the framework of the goal-and-feedback model. In particular, self-organizing tendencies may well occur in the perceptual input process, in goal selection, and in the behavioral output process. The next sections consider several such possibilities.

Self-Organization in Perception

Self-organization may take place within the perceptual input function (cf. Lewis, 1996). Self-organization may induce a particular perception, influencing how the person views a situation and thus how the person responds to it. This characterization of self-organization in perception actually isn't far from the common assumption that information contributing to perceptions is organized as it is brought to consciousness, in a way that depends both on the external stimuli themselves and on forces operating within the system (cf. Pashler, Johnston, & Ruthruff, 2001). Nor is it far from the idea that principles of organization (e.g., Gestalt principles) induce coherence in perception (Read et al., 1997).

This application of the principle of self-organization is also very compatible with arguments made by people in social psychology and elsewhere concerning connectionist networks that satisfy multiple constraints simultaneously (Kunda & Thagard, 1996; Read et al., 1997; Schultz & Lepper, 1996; Thagard, 1989). These models often (though not always) focus on how situations are perceived or construed. The multiple constraints within the perceiving system serve to induce an overall coherence within the system. How such a model "settles out" can be difficult to predict from its initial conditions and inputs, however. This suggests that a self-organizing process is taking place, rather than a more-simple additive process.

By influencing the perception that emerges, self-organizing in the input function can influence how the person behaves. Self-organization in perceptual input

[5]Of course, one could just as readily apply Harris's phrase to the bringing to mind of a goal, and go on to ask whether goals are meaningful as explanatory devices. Indeed, McFarland and Bösser (1993) have done this (for a different kind of critique see Wegner, 2002).

can even induce a new goal unexpectedly. If you suddenly realize there is a 20-dollar bill on the sidewalk in front of you (where previously you saw only leaves), the goal of bending over and picking it up may suddenly pop into mind. Indeed, that goal may immediately displace the goal of walking to the food court to have lunch, which was guiding the activity you were engaged in until that moment.

Another example of self-organization in perception concerns self-perception rather than perception of the outside world. Nowak et al. (2000) have argued that self-evaluations represent an emergent and self-organizing experience. They argued that self-evaluations emerge from repeated iterations of influence within an interconnected array of disparate bits of self-knowledge (though the simulations by which they tested this idea employed cellular automata rather than connectionist networks). The effect across multiple iterations seems to reflect a self-organizing tendency within the system as a whole (again, quite different from an additive process).

Self-Organization in Output

A second possibility for self-organization in behavior is that self-organization may be involved in the process of executing an action. It is often argued that actions "self-assemble" as they are taking place, such that the action actually executed is not precisely of the form that was initially intended (Kugler & Turvey, 1987). If your hands are full as you reach to close a door, you may find yourself pushing the door with one of the books you are holding, or with your knee, when you actually intended to use your hand or arm. This possibility seems consonant with views of motor control in which an executive may call for an action of some general form but leave the details of execution to lower order parts of the nervous system (cf. Gazzaniga, Ivry, & Mangun, 1998; Turvey, 1990).

There is also evidence that the form of physical movements is influenced by the immediately preceding movements. There appears to be a conservation of aspects of movement from one action to another (Vaughan, Mattson, & Rosenbaum, 1998). That is, just as the priming of a perceptual category renders its subsequent use more likely, people tend to re-use a just-used quality of motion, if it is a reasonably common way of executing the intended act. This lets the motor system make double use of the computations that were needed to do the first action. This phenomenon seems to represent self-organization, in the sense that it is not intentional.

There also appear to be cases (perhaps many of them) in which an intended action incorporates behavioral qualities that are not intended, but which nonetheless slip into the stream. For example, there is evidence that activating an interpretive concept in memory can cause the corresponding behavioral quality to emerge in the person's subsequent actions (Bargh, Chen, & Burrows, 1996, Experiment 2; Carver, Ganellen, Froming, & Chambers, 1983, Experiment 2; see also Prinz, 1997). These studies used an initial task to activate a cognitive concept. This in turn seemed to partly activate an output specification for the corresponding behavior. The behavioral quality then was incorporated into subsequent actions. Thus, after experiencing the priming of an aggressive interpretive schema, people were more punitive in a teaching task (Carver et al., 1983). After priming of the stereotype of the elderly (even subliminally), people walked more slowly down a corridor (Bargh et al., 1996).

The idea that activation of interpretive qualities can influence behavior has widespread implications. Whatever you read, hear on the radio, see on the street—whatever is activating any information in your mind that has a behavioral side to it—may influence in small ways the things you do in the next few minutes. The intrusion of behavioral qualities through such paths may be very common. Such phenomena certainly seem to reflect a kind of self-organization in action.

Self-Organization In Shifts Between Goals

The cases described in the preceding section involved only unintentional ripples on the surface of acts that were being done as intended. However, it is also possible for self-organization in output to have larger effects. Sometimes the act as it unfolds eventually differs entirely from the act that had been intended. In effect, the overall goal undergoes a substitution, a substitution that seems to have involved self-organization. Norman (1981) described a particularly vivid example of this, in which a person starts to change clothes to go out to dinner, but completes the sequence with a different ending—removes his clothes and goes to bed. This sort of thing presumably is most likely to happen when one part of an action sequence (the removal of clothes) belongs not just to the intended sequence but also to a different sequence, and the person somehow is shunted from the one to the other.

This side-tracking of goal-directed action can be very striking. Not all cases in which one activity gives way to a different one are so striking, but the shift generally seems to reflect the same dynamic: That is, the ongoing action fits the second activity as easily as the first one, and a cue of some sort induces the second one to emerge in place of the first one.

Here is another example that should be familiar to many readers. It is fairly common to be having a conversation with a colleague about nothing in particular, vaguely work related, but with no goal other than social exchange over lunch or coffee, catching up on what the other person has been up to. Without anticipation or forewarning, it turns out that you have both been thinking about a topic or issue from different but mutually helpful angles. Almost before you know it, you are planning a study together. There presumably is a point where you realize the casual conversation has become something else. But the process of getting to that something else—of changing goals—has been a self-organizing one.

Another illustration that may feel familiar to some people is how the transactions that take place in the context of a friendship (or professional relationship) sometimes induce a change to a higher level of intimacy. For a time the exchanges occur on a level that comfortably meets the goals associated with the existing relationship, then (sometimes abruptly) there is a shift such that a new "layer" of interaction appears, with different goals. Again, this shift gives the appearance of being based in self-organization.[6]

One might even argue that a self-organizational process intrudes briefly any time one goal is completed and another is begun. When one goal is attained, the priority management system (whatever it consists of) presumably offers up the next candidate waiting in line, but that may not be all that happens at that moment. Perhaps at the same time there is a brief chaotic sorting of options (cf. Vallacher & Kaufman, 1996), including a scan for possibilities that are not presently in the queue. Perhaps the decision about what to take up next itself reflects a multiple constraint satisfaction process. These characterizations suggest a process of self-organization in the decision about what goal to adopt next.

This character is distinctly present in the principle of "emergence" in action identification theory (Vallacher & Wegner, 1987). Emergence occurs when an activity is occurring smoothly at a relatively lower level of identification and the person drifts to a higher-level understanding of the nature of the same behavior. The specific higher-level identification that emerges is determined by the interweaving of many potential influences, rather than being predetermined. The process of emergence thus seems clearly to involve a self-organization (Vallacher & Nowak, 1997).

Self-Organization and Apparent Complexity

Self-organization tendencies such as those just described can lead to interesting problems of inference. When an outsider observes a stream of behavior, there often appears to be an organized complexity to it. As observers, we tend to interpret that complexity as intentional and planful (indeed, we often make inferences about the actor's personality from it, cf. Jones & Davis, 1965). But did the actor intend the complexity we observed? Does the actor even know the complexity is there? Sometimes, but not always. Sometimes complexity arises on the surface of behavior from the operation of forces other than the actor's intentions. The result can be the appearance of a goal where there actually is none.

Consistent with this general line of thought, Newtson (1993, 1994) has argued that dyadic interaction leads to emergent qualities in the behavioral stream that are not actually goals of either interactant. People in a conversation don't usually intend to engage in back and forth interruptions and hand-offs to each other, but that's what most conversations turn out to look like. People don't usually plan to mirror each other in postural adjustments in the course of the conversation, but such mirroring seems to occur naturally. In Newtson's view, such patterns are emergent from the dyadic process itself.

A conceptually similar line of thought is often taken in discussing gaits of locomotion, both in quadrupeds (Turvey, 1990) and in humans (Diedrich & Warren, 1998). Gaits reflect emergent self-organizations, rather than intended complexity. Horses don't intend particular patterns of foot placements (though it's unclear how anyone sees inside the mind of a horse to be sure). Rather, the patterns seem dictated by the physics involved in creating biological springs of varying torque while moving forward. Various gaits involve tradeoffs between speed and energy efficiency (which translates to comfort). Thus, when a horse changes gaits, it presumably is regulating speed and comfort (though there are other possibilities, cf. Holt, 1998). What appears on the surface, however, is a shift from one very complicated pattern of leg movements to another one.

An excellent illustration of how surface complexity can arise from a set of simple goal-directed activities, each running its own course, comes from a very different source: the literature of robotics. This is a good example partly because the design of the robot in the example ensures that no high-order goal exists (see Brooks, 1999, for a broad treatment of issues regarding the issue of apparent complexity in artificial systems, including this one). This robot travels down the halls of a laboratory, enters rooms, and locates, picks up, and retrieves empty soda cans (but not full ones). The goal of retrieving empty cans is not part of the robot's design. It has only specific modules controlling such action qualities as staying away from walls, approaching cylindrical objects, closing grips when an object is between its grippers, ungripping if the object is too heavy, retracing its path when the grippers are full. The robot doing these actions goes down the hall, enters of-

[6]Shifts in goal such as those discussed here presumably depend on shifts in perceptual input, or at least on the occurrence of particular cues that induce the alternate goal. Thus, in principle, it might be argued that the examples just described entail the same processes as the earlier example involving the twenty dollar bill, in which a new goal was induced by the changed perception.

fices, notices cans, reaches for them, closes on them, lifts them if they are light (empty), and returns with them. It retrieves empty cans. But it has no intention of doing so, and no knowledge of doing so.[7]

Differing Emphases Reconsidered

Before embarking on this discussion of self-organization tendencies, we said that the feedback-based model of behavior we've used and dynamic-systems models in psychology have somewhat different emphases. The feedback model tends to assume a top-down process, in which a goal is represented internally and efforts toward it are monitored and controlled by some sort of executive intentional process. In contrast, the dynamic-systems model tends to assume a bottom-up self-organizing process. The preceding sections explored how self-organizing tendencies might be fit into processes of the top-down intentional model. This section turns the comparison on its head, and considers how the feedback model can be fit to the dynamic-systems view.

Let's begin with the matter of top-down control. Although our use of the feedback model has assumed a top-down goal-based mechanism, that is not a requisite for the existence of a feedback loop.[8] Feedback loops exist in non-living systems, and in systems involving populations of living beings. Such loops do not always include every element shown in Figure 1. In particular, there is not always the distinct representation of a reference value. Rather, a feedback function appears to exist whenever an output creates an influence that ultimately affects subsequent output.

Consider, for example, a case in which neurons A and B project to each other, with A having an excitatory effect and B an inhibitory effect. A stimulates B, causing B to stimulate A. Because B's effect is inhibitory, it reduces A's stimulation of B. If A's level of activity falls, so does B's inhibitory effect, thereby inducing A's activity to increase. This arrangement is viewed as being a feedback loop (Lewis, 1998; Prosser, 1973). There is no reference value here, however, and no top-down control.

Another commonly used example of feedback uses an ecosystem comprising (for instance) a colony of rabbits and a food source. The size of the food source controls the size of the colony by determining how many animals can live on it. The population converges on the value that is the system's carrying capacity. If the population exceeds that value, some animals starve and the population falls. If the population falls below that value, there is surplus food and the population rises. There is no explicit reference value for population size, though, and no top-down control. This loop does not have the goal of stabilizing the size of the colony. Stabilizing the size of the colony is simply a consequence of the relations among the processes that form the loop. In cases such as this it is reasonable to refer to the function of the loop, but not to the purpose of the loop.[9]

Feedback systems may also depart from the depiction of Figure 1 in another way: the entity that exerts a given influence at one time in the loop (input or output) may not be the same entity as exerts that influence at another time. For example, in economic systems purchasing may be done by different people in two iterations of a feedback cycle in which purchasing plays a role. In sustaining the feedback loop, however, what matters is that the component processes—the input and output— continue to occur, not whether they occur via the same physical elements from moment to moment.

Examples such as these raise the question of whether the physical structures that underlie feedback control of behavior in the human nervous system should be ascribed only a function, rather than a purpose, and the question of whether different physical el-

[7]As an aside, we have found the literature of artificial agents to be a very interesting one. It is growing rapidly and converging with cognitive science. It provides a window on what is plausible (though with no assurance that living systems work the same way as artificial ones). Sources we have found useful, though by now somewhat dated, are articles by Beer (1995) and Maes (1990, 1994; see also Brooks, 1999, 2002). Much of this work proceeds by creating task-achieving modules, each involving both perception and action qualities, adding complexity by adding layers of modules. Brooks (1999) argues that this strategy is a rough analog to evolution, in which basic capabilities remain in place and newer ones are added. Of particular interest regarding the present discussion is the fact that robots such as those of Brooks have no executive or central representation of overall goals or of the world. Each layer operates autonomously. The goals of layers often are compatible (e.g., locomote, avoid hitting things, explore), but sometime conflicts arise (e.g., if locomotion would lead to hitting things) and are handled by a prearranged arbitration scheme (giving the layers different priorities so that a "higher" one can override a "lower" one). In contrast to Brooks, Maes (1990b, 1994) believes autonomous agents should have goals. Her approach shares with that of Brooks, however, the perception-action link, the "distributed" quality of action, and the absence of an executive, with decisions emerging in a continuous way from changing activation levels of specific modules.

[8]It is, of course, possible that feedback processes underlying behavior are a specialized subset of feedback processes, with constraints that are more narrow than those applicable to other cases.

[9]The distinction between function and purpose is an important one. As noted earlier, it is easy to ascribe a goal where none exists. There are several possibilities between simple physics and true goal-directed behavior. As noted by McFarland (1995; McFarland & Bösser, 1993), an entity that arrives at a place solely because of the structure of the environment is quite different from an entity that represents a goal and tries to reach it. For example, a ball placed on a surface that slopes toward a basin will move to the bottom of the basin. That, however, is a consequence of the force of gravity. It would be misleading to ascribe a goal to the ball and to interpret the ball's behavior as goal-directed (though, interestingly enough, the ball and basin are often used as a metaphor for an attractor). The feedback effects in the examples discussed in this section appear to be somewhere between the purely mechanical behavior of a ball on a sloping surface and purposive behavior.

ements create the influence over time (cf. Tononi & Edelman's, 1998, suggestion that a "dynamic core" of neural activity—an organized grouping that varies in composition from one time to another—underlies consciousness). In a view that permitted such variations, a feedback loop in the nervous system need comprise only a semi-structured pattern of influences, which continues to operate for a while (albeit perhaps a long while) because it happens to acquire a degree of stability. Such a description does not seem terribly different from the description of an attractor arising within a system of influences among diverse elements over a period of time. Thus, it seems quite possible to fit the feedback model to the dynamic-systems model, as well as vice versa.

Merging Self-Organization and Top-Down Control

One more way in which the two viewpoints might be interwoven uses the two views at difference stages in the evolution of a behavior. There is reason to suspect there are self-organizing, emergent rhythms and cycles in human behavior, which exert influences that have not been well appreciated. There is reason to suspect that people drift or stumble into patterns of action they hadn't done before. Yet it also seems reasonable to suggest that as emergent patterns stabilize over repeated occurrences, information about their nature is coded into memory in a form that permits the patterns themselves to be invoked for re-creation by an intentional process. To put it differently, a bottom-up process of self-organized pattern development may consolidate in a way that leaves an entry point for top-down control.

Does such consolidation occur? Clearly something like this happens in the process of skill learning. Something changes, as behaviors—even self-organized coordinations—are repeated over and over. Indeed, there is evidence that different parts of the brain are involved to different degrees when a behavior is relatively new versus being well practiced (Gazzaniga et al., 1998). There may be two modes of creating behavior at work, then, one of which operates bottom-up, the other top-down. Executive use of compiled capabilities can't happen without a solid record of what the capabilities are; one way for such a record to exist would be through an earlier emergence and consolidation of lower-order self-organized coordinations.

Two Mode Models

The idea that there exist two modes of processing is now fairly widespread. Many cognitive psychologists, having struggled with issues raised by differences between the symbolic approach (which tends to assume an executive) and connectionism (which does not), have turned to the idea that thinking involves two processes rather than one. For example, Smolensky (1988) argued that a *conscious processor* is used for effortful reasoning and following of programs of instructions. An *intuitive processor*, which manages intuitive problem solving, heuristic strategies, and skilled or automatic activities, relies on connectionist processes. In this view, influence can come from the top downward and also from the bottom upward. This general idea recurs in several other theories (see Holyoak & Spellman, 1993; Sloman, 1996).

There are also two-mode models in personality and social psychology (see Chaiken & Trope, 1999; Smith & DeCoster, 2000). An example that is particularly consonant with the preceding description is Epstein's (1985, 1994) cognitive–experiential self-theory. This theory posits two systems of processing. The rational system operates primarily consciously, uses logical rules of inference, and operates relatively slowly and effortfully. The experiential system is intuitive and crudely differentiated and functions relatively automatically, relying on heuristics and readily available information. Another example is Smith and DeCoster's model (2000), which explicitly suggests that an associative processor follows connectionist principles, and a rule-based processor follows symbolic principles.

By implication, at least, two-mode models also provide for two distinguishable sources of automaticity in behavior. Some instances of automaticity occur because the behavior is newly self-organizing and thus is not reflecting the person's intentions (see also Bechara et al., 1997). But automaticity can also accrue from the repetition of intentional acts (Bargh, 1997). Given enough repetition, control of an activity may drop out of consciousness altogether (Smolensky, 1988).

Thus, there may actually be three distinguishable ways in which behavior occurs. First, dynamic processes at a low level may automatically (without intent) produce emergent patterns. After repeated emergence, the pattern may become an attractor. After further recurrences, the pattern emerges more easily (cf. Kelso's 1995 discussion of diminished variability in the actions of professionals), suggesting that something is consolidating in memory.

Second, the consolidated pattern—the emergent quality at the higher level—becomes recognizable to a rational or conscious mode of processing. At some point it can be invoked by the conscious processor as a guide to behavior, even if the ability to execute it from the top down is not entirely smooth. With enough practice, the patterned information can be used top-down effectively (as well as re-emerge bottom-up). With more top-down use, the behavior becomes more accurate (matches more closely the patterns initially induced bottom-up). Having become accurate, it now begins to become more automatic.

The third form follows from this increasing automaticity. Gradually the awareness of calling for the action fades, as there is less and less need to invoke it explicitly. The more it fades, the more its use may become a matter of management by the intuitive mode. An interesting question is how this third way of acting differs from the first, bottom-up emergent process.

Self-Organization As Response Organization

We wish to point to one more way in which the term self-organization has been used in the past, which appears to bear on issues of organization and consolidation of response potentials. Quite some time ago a physicist named Donald MacKay (1956) made a set of suggestions about how an information-flow model of human behavior might be constructed. He made these suggestions in the form of speculations about the design of a self-organizing system that might have some of the same functional properties as a person.

He began with a feedback system, but quickly went beyond that, addressing the need for the existence of "inventive behavior." In that context he suggested principles by which the artificial system he had in mind might self-organize, so that its outputs were optimally fit to the situations it encountered. His sense of the term self-organization thus dealt specifically with structures within the system that would create an organization of (self-organize) first-choice outputs.

The key step in his thinking was to devise a principle by which the system could discover what aspects of its environment were predictable, learning to organize for itself the responses that would be appropriate to the patterns of events that it encountered over time. The control structures he posited permit the system to learn which responses are successful in particular circumstances—in effect, displaying the "law of effect" within the system. The system he sketched out yielded a hierarchy of "organizers," whose activity develops an array of behavioral responses to match the structure of the environmental events, provided those events have some degree of regularity.

MacKay's stated intent in that exercise was merely to provide an initial effort at simulating how organisms learn to predict the world and orchestrate adaptive responses to it. However, our impression is that he did a good deal more than that. His description appears to have foreshadowed core themes in connectionism, as he described how the activity of the internal organizing system as a whole symbolizes what is being perceived in the environment. Indeed, extrapolating the idea of organizers into a hierarchical model of action control also suggests a way of thinking about how self-assembly of action occurs at lower levels.

Concluding Comment

This article has addressed the question of how dynamic-systems ideas may be of value in thinking about goal directed action. We began by pointing out a structural similarity between aspects of the two models. From that point on, we focused primarily on ways in which the self-organization tendencies that are suggested by dynamic-systems thinking might be fit into a model of feedback control (though some attention was also given to fitting the model of feedback control to the principle of self-organization).

Focusing on self-organization tendencies makes several things apparent regarding goal directed action. For one, it seems clear that self-organization is a constant and ever-present feature of the perceptions that underlie action. It also seems clear that there are several ways in which the actions that people carry out often emerge in a form other than the form intended. Sometimes the departure from intention is very minor, not even really noticeable; sometimes the departure is more substantial, leading the person in new and unexpected directions. These also appear to suggest self-organizing influences.

Even beyond the idea that intentions are not always faithfully executed, it is clear that not all aspects of behavior are actively controlled (Kelso, 1995; Turvey, 1990). Sometimes complexity and patterning occur for reasons that do not involve internal mechanisms, and in such cases an attempt to specify an internal mechanism is actually misleading. At other times, however, the complexity of behavior is quite intentional, and the behavior is being consciously regulated toward desired ends. A full account of human behavior must allow for both of these possibilities.

A final comment returns us to the point made in the introduction of this article. The themes of a control-process view on action and a dynamic systems view are not identical, but they appear to be complementary rather than contradictory. The existence of self-organizing tendencies does not mean that a top-down goal-and-feedback view of behavior is wrong—only that it is incomplete. The proper question is not which view is useful, but how can their strengths complement each other. A melding of the strengths of these views may offer us the potential for a more complete understanding of the structure and dynamics of human behavior than we now have.

References

Amazeen, P. G., Amazeen, E. L., & Turvey, M. T. (1998). Dynamics of human intersegmental coordination: Theory and research. In D. A. Rosenbaum & C. Collyer (Eds.), *Timing of behavior: Neural, psychological, and computational perspectives* (pp. 237–259). Cambridge, MA: MIT Press.

Asch, S. E. (1955). Opinions and social pressure. *Scientific American, 193*, 31–35.

Aslin, R. N. (1993). Commentary: The strange attractiveness of dynamic systems to development. In L. D. Smith & E. Thelen (Eds.), *A dynamic systems approach to development: Applications* (pp. 385–399). Cambridge, MA: MIT Press.

Atkinson, J. W., & Birch, D. (1970). *The dynamics of action*. New York: Wiley.

Austin, J. T., & Vancouver, J. B. (1996). Goal constructs in psychology: Structure, process, and content. *Psychological Bulletin*, *120*, 338–375.

Bargh, J. A. (1997). The automaticity of everyday life. In R. S. Wyer, Jr. (Ed.), *Advances in social cognition* (Vol. 10, pp. 1–61). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Bargh, J. A., Chen, M., & Burrows, L. (1996). Automaticity of social behavior: Direct effects of trait construct and stereotype activation on action. *Journal of Personality and Social Psychology, 71*, 230–244.

Bechara, A., Damasio, H., Tranel, D., & Damasio, A. R. (1997). Deciding advantageously before knowing the advantageous strategy. *Science*, *275*, 1293–1295.

Beer, R. D. (1995). A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence*, *72*, 173–215.

Brooks, R. A. (1999). *Cambrian intelligence: The early history of the new AI*. Cambridge, MA: MIT Press.

Brooks, R. A. (2002). *Flesh and machines: How robots will change us*. New York: Pantheon.

Brown, C. (1995). *Chaos and catastrophe theories* (Quantitative applications in the social sciences, no. 107). Thousand Oaks, CA: Sage.

Cannon, W. B. (1932). *The wisdom of the body*. New York: Norton.

Carver, C. S. (1997). Dynamical social psychology: Chaos and catastrophe for all. *Psychological Inquiry*, *8*, 110–119.

Carver, C. S. (2001). Affect and the functional bases of behavior: On the dimensional structure of affective experience. *Personality and Social Psychology Review, 5*, 345–356.

Carver, C. S. (in press). Pleasure as a sign you can attend to something else: Placing positive feelings within a general model of affect. *Cognition and Emotion*.

Carver, C. S., & Humphries, C. (1981). Havana daydreaming: A study of self-consciousness and the negative reference group among Cuban Americans. *Journal of Personality and Social Psychology*, *40*, 545–552.

Carver, C. S., Lawrence, J. W., & Scheier, M. F. (1999). Self-discrepancies and affect: Incorporating the role of feared selves. *Personality and Social Psychology Bulletin*, *25*, 783–792.

Carver, C. S., & Scheier, M. F. (1998). *On the self-regulation of behavior.* New York: Cambridge University Press.

Carver, C. S., & Scheier, M. F. (1999a). Themes and issues in the self-regulation of behavior. In R. S. Wyer, Jr. (Ed.), *Advances in social cognition* (Vol. 12, pp. 1–105). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Carver, C. S., & Scheier, M. F. (1999b). Several more themes, a lot more issues: Commentary on the commentaries. In R. S. Wyer, Jr. (Ed.), *Advances in social cognition* (Vol. 12, pp. 261–302). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Carver, C. S., Ganellen, R. J., Froming, W. J., & Chambers, W. (1983). Modeling: An analysis in terms of category accessibility. *Journal of Experimental Social Psychology*, *19*, 403–421.

Carver, C. S., Sutton, S. K., & Scheier, M. F. (2000). Action, emotion, and personality: Emerging conceptual integration. *Personality and Social Psychology Bulletin*, *26*, 741–751.

Chaiken, S. L., & Trope, Y. (Eds.). (1999). *Dual-process theories in social psychology.* New York: Guilford.

Cloninger, C. R. (1988). A unified biosocial theory of personality and its role in the development of anxiety states: A reply to commentaries. *Psychiatric Developments, 2*, 83–120.

Cooper, R., & Shallice, T. (2000). Contention scheduling and the control of routine activities. *Cognitive Neuropsychology, 17*, 297–338.

Davidson, R. J. (1998). Affective style and affective disorders: Perspectives from affective neuroscience. *Cognition and Emotion, 12*, 307–330.

Davidson, R. J., Jackson, D. C., & Kalin, N. H. (2000). Emotion, plasticity, context, and regulation: Perspectives from affective neuroscience. *Psychological Bulletin, 126*, 890–909.

Diedrich, F. J., & Warren, W. H., Jr. (1998). Dynamics of human gait transitions. In D. A. Rosenbaum & C. E. Collyer (Eds.), *Timing of behavior: Neural, psychological, and computational perspectives* (pp. 323–343). Cambridge, MA: MIT Press.

Depue, R. A., & Collins, P. F. (1999). Neurobiology of the structure of personality: Dopamine, facilitation of incentive motivation, and extraversion. *Behavioral and Brain Sciences, 22*, 491–517.

Depue, R. A., & Zald, D. H. (1993). Biological and environmental processes in nonpsychotic psychopathology: A neurobehavioral perspective. In C. G. Costello (Ed.), *Basic issues in psychopathology* (pp. 127–237). New York: Guilford.

Elliott, A. J. (1999). Approach and avoidance motivation and achievement goals. *Educational Psychologist, 34*, 169–189.

Elliott, E. S., & Dweck, C. S. (1988). Goals: An approach to motivation and achievement. *Journal of Personality and Social Psychology*, *54*, 5–12.

Epstein, S. (1985). The implications of cognitive–experiential self theory for research in social psychology and personality. *Journal for the Theory of Social Behavior*, *15*, 283–310.

Epstein, S. (1994). Integration of the cognitive and the psychodynamic unconscious. *American Psychologist*, *49*, 709–724.

Festinger, L. (1957). *A theory of cognitive dissonance*. Evanston, IL: Row, Peterson.

Fowles, D. C. (1993). Biological variables in psychopathology: A psychobiological perspective. In P. B. Sutker & H. E. Adams (Eds.), *Comprehensive handbook of psychopathology* (2nd ed., pp. 57–82). New York: Plenum.

Gazzaniga, M. S., Ivry, R. B., & Mangun, G. R. (1998). *Cognitive neuroscience: The biology of the mind*. New York: Norton.

Gray, J. A. (1994a). Personality dimensions and emotion systems. In P. Ekman & R. J. Davidson (Eds.), *The nature of emotion: Fundamental questions* (pp. 329–331). New York: Oxford University Press.

Gray, J. A. (1994b). Three fundamental emotion systems. In P. Ekman & R. J. Davidson (Eds.), *The nature of emotion: Fundamental questions* (pp. 243–247). New York: Oxford University Press.

Heider, F. (1946). Attitudes and cognitive organization. *Journal of Psychology, 21*, 107–112.

Higgins, E. T. (1987). Self-discrepancy: A theory relating self and affect. *Psychological Review, 94*, 319–340.

Higgins, E. T. (1996). Ideals, oughts, and regulatory focus: Affect and motivation from distinct pains and pleasures. In P. M. Gollwitzer & J. A. Bargh (Eds.), *The psychology of action: Linking cognition and motivation to behavior* (pp. 91–114). New York: Guilford.

Holt, K. G. (1998). Constraints in the emergence of preferred locomotory patterns. In D. A. Rosenbaum & C. E. Collyer (Eds.), *Timing of behavior: Neural, psychological, and computational perspectives* (pp. 261–291). Cambridge, MA: MIT Press.

Holyoak, K. J., & Spellman, B. A. (1993). Thinking. *Annual Review of Psychology*, *44*, 265–315.

Jones, E. E., & Davis, K. E. (1965). From acts to disposition: The attribution process in person perception. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 2, pp. 219–266). New York: Academic Press.

Kelso, J. A. S. (1995). *Dynamic patterns: The self-organization of brain and behavior.* Cambridge, MA: MIT Press.

Kugler, P. N., & Turvey, M. T. (1987). *Information, natural law, and the self-assembly of rhythmic movement*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Kunda, Z., & Thagard, P. (1996). Forming impressions from stereotypes, traits, and behaviors: A parallel–constraint–satisfaction theory. *Psychological Review*, *103*, 284–308.

Lang, P. J. (1995). The emotion probe: Studies of motivation and attention. *American Psychologist, 50*, 372–385.
Lecky, P. (1945). *Self-consistency: a theory of personality.* Washington, DC: Island Press.
Lewis, M. D. (1996). Self-organising cognitive appraisals. *Cognition and Emotion, 10*, 1–25.
Lewis, R. (1998). *Life* (3rd ed.). Boston: McGraw-Hill.
MacKay, D. M. (1956). Towards an information-flow model of human behaviour. *British Journal of Psychology, 47*, 30–43.
Maes, P. (1990). Situated agents can have goals. In P. Maes (Ed.), *Designing autonomous agents: Theory and practice from biology to engineering and back* (pp. 49–70). Cambridge, MA: MIT Press.
Maes, P. (1994). Modeling adaptive autonomous agents. *Artificial Life*, *1*, 135–162.
Markus, H., & Nurius, P. (1986). Possible selves. *American Psychologist*, *41*, 954–969.
McFarland, D. (1995). Opportunity versus goals in robots, animals, and people. In H. L. Roitblat & J.-A. Meyer (Eds.), *Comparative approaches to cognitive science* (pp. 415–433). Cambridge, MA: MIT Press.
McFarland, D., & Bösser, T. (1993). *Intelligent behavior in animal and robots*. Cambridge, MA: MIT Press.
Miller, G. A., Galanter, E., & Pribram, K. H. (1960). *Plans and the structure of behavior.* New York: Holt, Rinehart, & Winston.
Miller, L. C., & Read, S. J. (1987). Why am I telling you this? Self-disclosure in a goal-based model of personality. In V. J. Derlega & J. Berg (Eds.), *Self-disclosure: Theory, research, and therapy* (pp. 35–58). New York: Plenum.
Montant, M. (2000). Feedback: A general mechanism in the brain. *Behavior and Brain Science, 23*, 340–341.
Murray, H. A. (1938). *Explorations in personality.* New York: Oxford University Press.
Newtson, D. (1993). The dynamics of action and interaction. In L. D. Smith & E. Thelen (Eds.), *A dynamic systems approach to development: Applications* (pp. 241–264). Cambridge, MA: MIT Press.
Newtson, D. (1994). The perception and coupling of behavioral waves. In R. R. Vallacher & A. Nowak (Eds.), *Dynamical systems in social psychology* (pp. 139–167). San Diego, CA: Academic Press.
Norman, D. A. (1981). Categorization of action slips. *Psychological Review*, *88*, 1–15.
Norman, D. A., & Shallice, T. (1986). Attention to action: Willed and automatic control of behavior. In R. J. Davidson, G. E. Schwartz, & D. Shapiro (Eds.), *Consciousness and self-regulation: Advances in research and theory* (Vol. 4, pp. 1–18). New York: Plenum.
Nowak, A., Vallacher, R. R., Tesser, A., & Borkowski, W. (2000). Society of self: The emergence of collective properties in self-structure. *Psychological Review, 107*, 39–61.
Pashler, H., Johnston, J. C., & Ruthruff, E. (2001). Attention and performance. *Annual Review of Psychology, 52*, 629–651.
Pervin, L. A. (Ed.). (1989). *Goal concepts in personality and social psychology.* Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
Powers, W. T. (1973). *Behavior: The control of perception*. Chicago: Aldine.
Powers, W. T. (1978). Quantitative analysis of purposive systems: Some spadework at the foundations of scientific psychology. *Psychological Review*, *85*, 417–435.
Pribram, K. H. (1991). *Brain and perception: Holonomy and structure in figural processing*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
Prigogene, I., & Stengers, I. (1984). *Order out of chaos*. New York: Bantam.
Prinz, W. (1997). Perception and action planning. *European Journal of Cognitive Psychology*, *9*, 129–154.
Prosser, C. L. (1973). Central nervous systems. In C. L. Prosser (Ed.), *Comparative animal physiology* (3rd ed., pp. 634–635). Philadelphia: Saunders.
Read, S. J., Vanman, E. J., & Miller, L. C. (1997). Connectionism, parallel constraint satisfaction processes, and Gestalt principles: (Re)introducing cognitive dynamics to social psychology. *Review of Personality and Social Psychology*, *1*, 26–53.
Rosenbaum, D. A. (1998). Is dynamical systems modeling just curve fitting? *Motor Control, 2*, 101–104.
Rosenbaum, D. A., Carlson, R. A., & Gilmore, R. O. (2001). Acquisition of intellectual and perceptual-motor skills. *Annual Review of Psychology, 52*, 453–470.
Ryan, R. M., Sheldon, K. M., Kasser, T., & Deci, E. L. (1996). All goals are not created equal: An organismic perspective on the nature of goals and their regulation. In P. M. Gollwitzer & J. A. Bargh (Eds.), *The psychology of action: Linking cognition and motivation to behavior* (pp. 7–26). New York: Guilford.
Schultz, T. R., & Lepper, M. R. (1996). Cognitive dissonance reduction as constraint satisfaction. *Psychological Review*, *103*, 219–240.
Shallice, T. (1978). The dominant action system: An information-processing approach to consciousness. In K. S. Pope & J. L. Singer (Eds.), *The stream of consciousness: Scientific investigations into the flow of human experience* (pp. 117–157). New York: Wiley.
Simon, H. A. (1967). Motivational and emotional controls of cognition. *Psychological Review, 74*, 29–39.
Sloman, S. A. (1996). The empirical case for two forms of reasoning. *Psychological Bulletin*, *119*, 3–22.
Smith, E. R., & DeCoster, J. (2000). Dual-process models in social and cognitive psychology: Conceptual integration and links to underlying memory systems. *Personality and Social Psychology Review*, *4*, 108–131.
Smolensky, P. (1988). On the proper treatment of connectionism. *Behavioral and Brain Sciences*, *11*, 1–23.
Thagard, P. (1989). Explanatory coherence. *Behavioral and Brain Sciences*, *12*, 435–467.
Tononi, G., & Edelman, G. M. (1998). Consciousness and complexity. *Science, 282*, 1846–1851.
Turvey, M. T. (1990). Coordination. *American Psychologist*, *45*, 938–953.
Vallacher, R. R., & Kaufman, J. (1996). Dynamics of action identification: Volatility and structure in the mental representation of behavior. In P. M. Gollwitzer & J. A. Bargh (Eds.), *The psychology of action: Linking cognition and motivation to behavior* (pp. 260–282). New York: Guilford.
Vallacher, R. R., & Nowak, A. (1997). The emergence of dynamical social psychology. *Psychological Inquiry, 8*, 73–99.
Vallacher, R. R., & Wegner, D. M. (1987). What do people think they're doing? Action identification and human behavior. *Psychological Review*, *94*, 3–15.
Vaughan, J., Mattson, T. R., & Rosenbaum, D. A. (1998). The regulation of contact in rhythmic tapping. In D. A. Rosenbaum & C. E. Collyer (Eds.), *Timing of behavior: Neural, psychological, and computational perspectives* (pp. 195–211). Cambridge, MA: MIT Press.
von Bertalanffy, L. (1968). *General systems theory.* New York: Braziller.
Wegner, D. M. (2002). *The illusion of conscious will*. Cambridge, MA: MIT Press.
Wiener, N. (1948). *Cybernetics: Control and communication in the animal and the machine*. Cambridge, MA: MIT Press.