PERSPECTIVE

Complexity Theory and Organization Science

Philip Anderson
Amos Tuck School, Dartmouth College, Hanover New Hampshire 03755-9000

Abstract

Complex organizations exhibit surprising, nonlinear behavior. Although organization scientists have studied complex organizations for many years, a developing set of conceptual and computational tools makes possible new approaches to modeling nonlinear interactions within and between organizations. Complex adaptive system models represent a genuinely new way of simplifying the complex. They are characterized by four key elements: agents with schemata, self-organizing networks sustained by importing energy, coevolution to the edge of chaos, and system evolution based on recombination. New types of models that incorporate these elements will push organization science forward by merging empirical observation with computational agent-based simulation. Applying complex adaptive systems models to strategic management leads to an emphasis on building systems that can rapidly evolve effective adaptive solutions. Strategic direction of complex organizations consists of establishing and modifying environments within which effective, improvised, self-organized solutions can evolve. Managers influence strategic behavior by altering the fitness landscape for local agents and reconfiguring the organizational architecture within which agents adapt.
(*Complexity Theory*; *Organizational Evolution*; *Strategic Management*)

Since the open-systems view of organizations began to diffuse in the 1960s, *complexity* has been a central construct in the vocabulary of organization scientists. Open systems are open because they exchange resources with the environment, and they are systems because they consist of interconnected components that work together. In his classic discussion of hierarchy in 1962, Simon defined a complex system as one made up of a large number of parts that have many interactions (Simon 1996). Thompson (1967, p. 6) described a complex organization as a set of interdependent parts, which together make up a whole that is interdependent with some larger environment.

Organization theory has treated complexity as a structural variable that characterizes both organizations and their environments. With respect to organizations, Daft (1992, p. 15) equates complexity with the number of activities or subsystems within the organization, noting that it can be measured along three dimensions. Vertical complexity is the number of levels in an organizational hierarchy, horizontal complexity is the number of job titles or departments across the organization, and spatial complexity is the number of geographical locations. With respect to environments, complexity is equated with the number of different items or elements that must be dealt with simultaneously by the organization (Scott 1992, p. 230). Organization design tries to match the complexity of an organization's structure with the complexity of its environment and technology (Galbraith 1982).

The very first article ever published in *Organization Science* suggested that it is inappropriate for organization studies to settle prematurely into a normal science mindset, because organizations are enormously complex (Daft and Lewin 1990). What Daft and Lewin meant is that the behavior of complex systems is surprising and is hard to

1047-7039/99/1003/0216/$05.00

predict, because it is *nonlinear* (Casti 1994). In nonlinear systems, intervening to change one or two parameters a small amount can drastically change the behavior of the whole system, and the whole can be very different from the sum of the parts. Complex systems change inputs to outputs in a nonlinear way because their components interact with one another via a web of feedback loops.

Gell-Mann (1994a) defines complexity as the length of the schema needed to describe and predict the properties of an incoming data stream by identifying its regularities. Nonlinear systems can difficult to compress into a parsimonious description: this is what makes them complex (Casti 1994). According to Simon (1996, p. 1), the central task of a natural science is to show that complexity, correctly viewed, is only a mask for simplicity. Both social scientists and people in organizations reduce a complex description of a system to a simpler one by abstracting out what is unnecessary or minor. To build a model is to encode a natural system into a formal system, compressing a longer description into a shorter one that is easier to grasp. Modeling the nonlinear outcomes of many interacting components has been so difficult that both social and natural scientists have tended to select more analytically tractable problems (Casti 1994). Simple boxes-and-arrows causal models are inadequate for modeling systems with complex interconnections and feedback loops, even when nonlinear relations between dependent and independent variables are introduced by means of exponents, logarithms, or interaction terms. How else might we compress complex behavior so we can comprehend it?

For Perrow (1967), the more complex an organization is, the less knowable it is and the more deeply ambiguous is its operation. Modern complexity theory suggests that some systems with many interactions among highly differentiated parts can produce surprisingly simple, predictable behavior, while others generate behavior that is impossible to forecast, though they feature simple laws and few actors. As Cohen and Stewart (1994) point out, normal science shows how complex effects can be understood from simple laws; chaos theory demonstrates that simple laws can have complicated, unpredictable consequences; and complexity theory describes how complex causes can produce simple effects.

Since the mid-1980s, new approaches to modeling complex systems have been emerging from an interdisciplinary invisible college, anchored on the Santa Fe Institute (see Waldrop 1992 for a historical perspective). The agenda of these scholars includes identifying deep principles underlying a wide variety of complex systems, be they physical, biological, or social (Fontana and Ballati 1999). Despite somewhat frequent declarations that a new paradigm has emerged, it is still premature to declare that a science of complexity, or even a unified theory of complex systems, exists (Horgan 1995). Holland and Miller (1991) have likened the present situation to that of evolutionary theory before Fisher developed a mathematical theory of genetic selection.

This essay is not a review of the emerging body of research in complex systems, because that has been ably reviewed many times, in ways accessible to both scholars and managers. Table 1 describes a number of recent, prominent books and articles that inform this literature; Heylighen (1997) provides an excellent introductory bibliography, with a more comprehensive version available on the Internet at `http://pespmc1.vub.ac.be/Evocobib.html`. Organization science has passed the point where we can regard as novel a summary of these ideas or an assertion that an empirical phenomenon is consistent with them (see Browning et al. 1995 for a pathbreaking example).

Six important insights, explained at length in the works cited in Table 1, should be regarded as well-established scientifically. First, many dynamical systems (whose state at time t determines their state at time $t + 1$) do not reach either a fixed-point or a cyclical equilibrium (see Dooley and Van de Ven's paper in this issue). Second, processes that appear to be random may be chaotic, revolving around identifiable types of *attractors* in a deterministic way that seldom if ever return to the same state. An attractor is a limited area in a system's state space that it never departs. Chaotic systems revolve around "strange attractors," fractal objects that constrain the system to a small area of its state space, which it explores in a never-ending series that does not repeat in a finite amount of time. Tests exist that can establish whether a given process is random or chaotic (Koput 1997, Ott 1993). Similarly, time series that appear to be random walks may actually be fractals with self-reinforcing trends (Bar-Yam 1997). Third, the behavior of complex processes can be quite sensitive to small differences in initial conditions, so that two entities with very similar initial states can follow radically divergent paths over time. Consequently, historical accidents may "tip" outcomes strongly in a particular direction (Arthur 1989). Fourth, complex systems resist simple reductionist analyses, because interconnections and feedback loops preclude holding some subsystems constant in order to study others in isolation. Because descriptions at multiple scales are necessary to identify how emergent properties are produced (Bar-Yam 1997), reductionism and holism are complementary strategies in analyzing such systems (Fontana and Ballati

Table 1 Selected Resources that Provide an Overview of Complexity Theory

Allison and Kelly, 1999	Written for managers, this book provides an overview of major themes in complexity theory and discusses practical applications rooted in experiences at firms such as Citicorp.
Bar-Yam, 1997	A very comprehensive introduction for mathematically sophisticated readers, the book discusses the major computational techniques used to analyze complex systems, including spin-glass models, cellular automata, simulation methodologies, and fractal analysis. Models are developed to describe neural networks, protein folding, developmental biology, and the evolution of human civilization.
Brown and Eisenhardt, 1998	Although this book is not an introduction to complexity theory, a series of small tables throughout the text introduces and explains most of the important concepts. The purpose of the book is to view strategic change through the lens of complexity theory. Brown and Eisenhardt contend that successful companies manage in turbulent environments by charting a strategic course that puts them on the edge of chaos, poised between order and disorder.
CalResCo Introduction to Complex Systems	Available at `http://www.calresco.org/intro.htm#def`, this website provides a comprehensive introduction and set of links that introduce many different aspects of complexity theory. It is a good, free, and relatively nontechnical starting place for exploring the topic.
Capra, 1996	Written for laymen, this book traces the development of systems theories, then discusses self-organization and the mathematics of chaos theory. Complexity theory is discussed in the context of how life emerged from inert chemical components.
Coveney and Highfield, 1995	Written for laymen, this book traces the history of thinking about computational complexity in mathematics. It discusses quite readably cellular automata, spin-glass models, neural networks, and genetic algorithms, self-organization, artificial life, and theories of brain functioning.
Cowan, Pines, and Meltzer, 1994	This is a collection of papers presented at the Santa Fe Institute's Fall 1991 workshop on integrative themes of the sciences of complexity. Although no single chapter provides a comprehensive overview, taken together, the chapters thoroughly cover the predominant themes developed by scholars of complex systems. Six chapters introduce fundamental concepts; the rest provide a number of examples of complex adaptive systems (principally, but not exclusively, biological), while four explore cellular automata, self-organized criticality and the "edge of chaos," and the concept of emergence.
Mainzer, 1994	Written at a high level but without extensive use of mathematics, this is a comprehensive overview of complex systems theory. Chapters describe different models in the context of the evolution of matter, the evolution of life, the evolution of the brain, the evolution of artificially intelligent computational systems, and the evolution of human society.
United Nations University, 1985	A collection of chapters from a very early conference on complexity theory in 1984, this book no longer captures the main lines along which complexity theory has developed. Nonetheless, the individual chapters, though quite eclectic, remain thought-provoking. This book is more a source of interesting ideas than a comprehensive introduction to the field.
Waldrop, 1992	Written for laymen, this book is a popular yet sophisticated introduction to complexity theory. It is the most readable introduction to the field for nonspecialists. In tracing the founding and early years of the Santa Fe Institute, Waldrop touches on applications of complexity theory to economic systems, Boolean networks and the *NK* model, genetic algorithms and classifier systems, self-organization and artificial life, and evolution to the edge of chaos.
Weisbuch, 1991	This is a methodological book, which introduces analytical techniques at the introductory graduate-school level. Topics include cellular automata, neural networks, simulated annealing, Boolean networks, and evolutionary population dynamics.

1999). Fifth, complex patterns can arise from the interaction of agents that follow relatively simple rules. These patterns are "emergent" in the sense that new properties appear at each level in a hierarchy (Holland 1995). Sixth, complex systems tend to exhibit "self-organizing" behavior: starting in a random state, they usually evolve toward order instead of disorder (Kauffman 1993).

The Evolution of Modern Complexity Theory

As Simon (1996) has pointed out, these ideas have deep historical roots; the intellectual ferment reviewed by the works in Table 1 represents the third wave of interest in complex systems this century. First, the years after World War I saw an explosion of interest in holism and gestalt theories. Then, cybernetics and general systems theory emerged after World War II, fueled by the success of wartime feedback-control devices and accelerated by the development of computers. These intellectual movements meant to replace reductionism with an appreciation for modeling interactions instead of simplifying them away. Presaging today's most enthusiastic "science of complexity" boosters, one of the founders of cybernetics declared in 1955: "Science is at last giving serious attention to systems that are intrinsically complex" (Ashby 1981, p. 219).

Cybernetics emphasized coordination, regulation, and control using feedback loops (Ashby 1956). General systems theory (Forrester 1961, von Bertalanffy 1968) attempted to elucidate deep principles underlying all types of systems whose components are linked by feedback loops. Both influenced the intellectual revolution that swept organization theory in the 1960s and ushered in a view of organizations as open systems (Katz and Kahn 1978). The systems design school of organization theory (Haberstroh 1965) was based on a characterization of systems as a collection of black boxes connected by input–output loops. A new breed of systems analysts designed work processes and organizational control systems around the ideas of systems theory and cybernetics (Beniger 1986). Pointing out that social organizations are more loosely coupled than most physical systems, Weick (1979) introduced a theory of organizing based on loosely coupled subassemblies called "double interacts," behavioral cycles linking the behavior of two people in a set of feedback loops. Systems dynamics models continue to inform a broad stream of contemporary research (e.g., Samuel and Jacobsen 1997, and Sterman and Witterberg in this issue), and nonlinear dynamical systems theory has been used to study many complex, nonlinear behavior patterns (Epstein 1997).

According to Simon (1996), the third wave of theories about complex systems is rooted in a new understanding of equilibrium that emerged in the late 1960s. Catastrophe theory (Thom 1975) explained how in some deterministic systems, a small shift in a parameter could send the system to a very different equilibrium. Chaos theory explored how some dynamical systems that appear to be random are, in fact deterministic (Thietart and Forgues 1995). Typically, such systems take the value of a variable in time t, stretch it, then fold it to produce a new value at time $t + 1$ (Cohen and Stewart 1994). The stretching operation magnifies small initial differences, while the folding operation constrains the range of values to a relatively small volume of the state space. It is usually impossible to forecast the exact value of a chaotic system in nature, because small measurement errors between two 'apparently identical values at time t can lead to large differences at time $t + 1$. However, such systems are in equilibrium around a strange attractor, a limited region of the state space within which the system stays permanently. Similarly, nonchaotic dynamical systems that appear to be random ("white noise") may in fact have underlying structural time trends ("colored noise") as Dooley and Van de Ven's paper (this issue) discusses. Colored noise can have important effects on outcomes, such as the risk that a population will become extinct (Heino 1998, Ripa and Lundberg 1996), and it should not be confused with chaos or randomness.

Cybernetics, general systems theory, catastrophe theory, and chaos theory all address deterministic dynamical systems, systems where a set of equations determine how a system moves through its state space from time t to time $t + 1$. Another way of modeling complex behavior examines regularity that emerges from interaction of individuals connected together in *complex adaptive systems* (CASs). The hallmark of this perspective is the notion that at any level of analysis, order is an emergent property of individual interactions at a lower level of aggregation. Although there is no universally accepted paradigm for describing CASs (Gell-Mann 1994b), four elements characterize models that have particularly interesting implications for organization theorists.

Agents with Schemata. First, to model an outcome at a particular level of analysis, one assumes that the outcome is produced by a dynamical system comprised of agents at a lower level of aggregation (Holland and Miller 1991). For example, in a model of an organization, agents might be individuals, groups, or coalitions of groups. Each agent's behavior is dictated by a *schema,* a cognitive structure that determines what action the agent takes at time t, given its perception of the environment (at time t, or at time $t - k$ if theoretical considerations suggest applying a lag structure). Different agents may or may not have different schemata (depending on one's theory), and schemata may or may not evolve over time. Often, agents' schemata are modeled as a set of rules, but schemata may be characterized in very flexible ways. For example, an agent may select one rule from a suite of possible rules, or it may invoke fuzzy rules, or its cognitive structure may be represented by a neural network (described in more detail later in this article).

Self-Organizing Networks Sustained by Importing Energy. Second, agents are partially connected to one another, so that the behavior of a particular agent depends on the behavior (or state) of some subset of all the agents in the system. In systems dynamics models, *variables* are connected to one another by feedback loops; in CAS models, *agents* are connected to one another by feedback loops. Each agent observes and acts on local information only, derived from those other agents to which it is connected. Unlike cybernetic control theories, no single component dictates the collective behavior of the system: such systems self-organize (Drazin and Sandelands 1992). Maintaining a self-organized state requires importing energy into the system (Prigogine and Stengers 1984).

Coevolution to the Edge of Chaos. Third, agents coevolve with one another. Each agent adapts to its environment by striving to increase a payoff or fitness function over time (Holland and Miller 1991). Each individual's payoff function depends on choices that other agents make, so each agent's adaptive landscape—mapping its behavior to its realized outcomes—is constantly shifting (e.g., Levinthal 1997). The equilibrium that results from such coevolution is dynamic, not static: small changes in behavior at time t can produce small, medium, or large changes in outcomes at time $t + 1$, according to a power law (see Morel and Ramanujam, this issue). Unlike chaotic equilibria, where small changes in behavior frequently cause large changes in outcomes, power–law equilibria lie at the edge of chaos (Kauffman 1993).

Recombination and System Evolution. Fourth, complex adaptive systems evolve over time through the entry, exit, and transformation of agents. New agents may be formed by recombining elements of previously successful agents. Furthermore, the linkages between agents may evolve over time, shifting the pattern of interconnections, the strength of each connection, and its sign or functional form. CASs can contain other complex adaptive systems, as, for example, organisms have immune systems (Gell-Mann 1994a).

CAS models represent a genuinely new way of simplifying the complex, of encoding natural systems into formal systems. Instead of making nonlinear systems tractable by reducing them to a set of causal variables and an error term, CAS models typically show how complex outcomes flow from simple schemata and depend on the way in which agents are interconnected. Rather than assuming that aggregate outcomes represent a homeostatic equilibrium, they show how such outcomes evolve from the efforts of agents to achieve higher fitness. By not forcing scholars to understand all the parts of a complex system in a holistic way, they allow investigators to focus on an agent in its local environment. It becomes possible to grasp complex behavior by varying assumptions about the schemata, connections, fitness functions, or population dynamics that characterize the agents. CAS models afford exciting new opportunities for analyzing complex systems without abstracting away their interdependencies and nonlinear interactions. This is particularly important for organizational scholars because interdependency is central to modern conceptions of what an organization is (Barnard 1938, Thompson 1967).

In the next section, I discuss each of these four features in greater detail, examining how each contributes to organization theory. Then I turn to a discussion how these ideas might lead to new ways of modeling organizational phenomena, and conclude by assessing how complexity theory creates new directions for research in strategic management.

Key Elements of Complex Adaptive Systems Models

Agents with Schemata

Most conceptual and empirical models employed by scientists studying organizations use a set of independent variables to explain variation in one or more dependent variables. Typically, outcomes at one level are explained by causal drivers at the same level of analysis. CAS models take a different approach. They ask how changes in the agents' decision rules, the interconnections among agents, or the fitness function that agents employ produce different aggregate outcomes. These models are inherently multilevel, because order is considered an emergent property that depends on how lower-level behaviors are aggregated. Accordingly, they respond well to contemporary calls for more integrative, cross-level research in organization science (Rousseau and House 1994).

CAS models and ordinary causal models are complements, not rivals. It is not necessary for scholars to adopt one or the other as the best way to analyze organizations. Causal theories and tests that relate variables on the same level identify important aggregate regularities and factors that help create them. CAS models build on this foundation, explaining observed regularities as the product of structured, evolving interactions among lower-level units. Good CAS models should not only explain established findings, but successfully predict new aggregate regularities and aggregate-level causal relationships.

Routinely, CAS models characterize agents as following a set of rules (Gell-Mann 1994). Rule-based models are also common in organization theory (Carley 1995), but representing human actors in this way is problematic. Institutional theorists have shown that rules are often rationalized myths (Meyer and Rowan 1977). Individual goals and intentions may be only loosely related to behavior (March and Olsen 1976), and rules may well be inferred from behavior instead of causing behavior (Weick 1979). Scholars who view organizations as natural systems have shown that rules often do not govern actions; rules can change without behavioral consequences, and behavior can change without modifications to rule systems (Scott 1992).

Agents in CAS models need not be the prisoners of a fixed set of rules. In Boulding's 1956 arrangement of general systems according to their complexity, social systems are distinguished by the fact that symbol-processing actors who share a common social order organize information from the environment into a knowledge structure. Simon (1996) distinguished recipes (feedback-triggered sequences of specific activities) from blueprints, images of the environment that attempt to capture its salient complexity. In routinized situations, actors can employ recipes, but in the face of greater uncertainty, problem-solving responses based on these blueprints are necessary (March and Simon 1958). In social psychology, such blueprints are termed "schemata" (Rumelhart 1984).

The social order that characterizes Boulding's social systems arises from interactions among agents (Mead 1934). In the symbolic interaction perspective, individual actors struggle for control of shared interpretation; roles and rules are negotiations and gambits in the struggle to define meaning. Because agents in CAS models can be endowed with schemata more complicated than simple rule systems, they can capture this struggle over meaning. In so doing, they can help integrate interpretivist and positivist views of organizations (Lee 1990), by providing clear ways for scholars in the interpretive tradition to describe the meanings that actors in a particular complex historical situation construct together.

For Gell-Mann (1994a), the characteristic that distinguishes a complex adaptive system from evolving yet nonadaptive systems, such as galaxies, is that it condenses environmental regularities into schemata. Furthermore, Gell-Mann argues that complex adaptive systems encode their environments into *many* schemata that compete against one another internally. An organizational example is the competing interpretive structures inside Intel that allowed internal selection processes to produce organizational-level adaptation, as the firm abandoned memory chips to focus on microprocessors (Burgelman 1994).

Because agents can possess multiple competing schemata at any one time, CAS models embody Campbell's (1974) idea that evolution occurs through a nested hierarchy of selective systems. Evolving actors develop vicarious selective systems so that they can experiment and fail without being killed; for example, animals have inherited instinctive pattern-recognition systems that let them identify potential predators and flee. Such indirect selective systems are nearly universal, because animals that fall heir to them from their ancestors are more likely to survive. Because schemata can evolve more rapidly than agents can, complex adaptive systems enjoy similar selective advantages when they allow schemata to compete and reinforce those that seem to be associated with favorable outcomes. In contrast, evolutionary game theory (Weibull 1995) models repeated games played over and over again by actors who are preprogrammed to implement a single strategy. Only the disappearance of actors whose strategies fail can alter the distribution of strategies in a population.

The notion that society is a "marketplace of ideas" is commonplace, but models of organizations in which knowledge structures compete with one another and evolve are rare. One promising avenue of inquiry that the CAS perspective opens up is exploration into how ideas, initiatives, and interpretations form an internal ecology within an organization. As McKelvey (1997) points out, organizational scholars have emphasized macroevolution (within organizational populations) at the expense of microevolution (within organizations). As CAS modeling concepts diffuse within organization science, we can expect to see more attention paid to the coevolution of actors and their knowledge structures (A. Lewin et al. 1999).

Paul et al. (1996) provide an interesting example of a model incorporating the simultaneous evolution of agents and their schemata. In their organization, nine agents employing different decision rules must all contribute to an aggregate decision. Action is taken only if all nine agents' recommendations are congruent. A decision mechanism is specified that controls which agents' output is required to activate other agents. Paul et al. specify a fitness function that the organization tries to meet and a feedback function that compares the outcome of each decision opportunity with the performance objective. Through a competitive bidding process, agents that contribute to successful decisions are more likely to be utilized in future decisions. Variation is introduced into this model by endowing new agents with hybrid decision rules that combine aspects of the decision rules that high-performing agents have employed. Models such as this are an important building block for future organizational research that adopts a CAS perspective.

Self-Organizing Networks Sustained by Importing Energy

As Drazin and Sandelands (1992) point out, systems that consist of independent actors whose interactions are governed by a system of recursively applied rules naturally generate stable structure. They self-organize; pattern and regularity emerge without the intervention of a central controller. When we observe order in a system, they argue, we should search for a set of rules that explain how connections between agents at time t influence connections at time $t + 1$. Rules generate structure because the state that is the output of one application of rules becomes the input for the next round. As Weick (1979) notes, managers often get in the way of activities that have their own self-regulation, form, and self-correcting tendencies.

Self-organization, or "autogenesis," is the natural result of nonlinear interaction, not any tendency of individual agents to prefer or seek order (Fontana and Ballati 1999). When the interactions of large numbers of components involve positive feedback loops, some behaviors self-amplify, quickly crowding out others. Groups of components become locked into self-reinforcing feedback cycles that lead to predictable collective behavior. Interacting microscopic entities form macroscopic structures that simplify the input structure of other macroscopic structures (Drazin and Sandelands 1992).

Self-organization only occurs in open systems that import energy from the outside (Prigogine and Stengers 1984). The second law of thermodynamics states that closed systems degenerate to a fixed-point equilibrium characterized by maximum disorder. In contrast, a "dissipative structure" is an organized state that arises when a system is maintained far from thermodynamic equilibrium because energy is constantly injected into it.

Organizations are dissipative structures that can only be maintained when members are induced to contribute energy to them (Barnard 1938). Social entities always self-organize as long as their members contribute work; this is why informal structures emerge and persist in a way that is remarkably robust to changes in the formal organizational structure. Those with influence and/or authority turn the heat up or down on an organization by recruiting new sources of energy (e.g., members, suppliers, partners, and customers), by motivating stakeholders, by shaking up the organization, and by providing new sets of challenges that cannot be mastered by hewing to existing procedures. Generally, the more turbulent an organization's environment is, the more energy must be generated to keep the system above the threshold beyond which self-organization is sustained.

When we observe complex aggregate structures, such as multinational corporations or the economic web of Silicon Valley, we need not search for complex building blocks. A defining feature of complexity is that self-organization is a natural consequence of interactions between simple agents. Paradoxically, scholars have abstracted away nonlinear interactions for the sake of analytical tractability, even though such interactions are the key to the emergence of pattern. When there are too few components or not enough interactions among them, patterns tend not to emerge (Weick 1979). Instead of making nonlinear systems tractable by modeling complex building blocks with few interactions, we can make them understandable by modeling simple building blocks with many interactions.

Order requires, however, that the number of interactions stay within an upper boundary as well as a lower one. Order arises in complex adaptive systems because their components are partially, not fully, connected. Systems in which every element is connected to each other in a feedback loop are hopelessly unstable (Simon 1996). Instead, CASs tend to form a decompositional hierarchy, in which elements are loosely coupled with one another (Simon 1996). Most components receive inputs from only a few of the system's other components, so change can be isolated to local neighborhoods.

When all organizational actors are interconnected with one another, either decay results (if feedback loops quickly dampen out change) or chaos ensues (if changes keep reverberating throughout the system). In complex adaptive systems, agents only act on information available in their immediate environments, from those few agents connected to them in a feedback loop. Simple mathematical models that capture pattern formation tend to rely on local activation and long-range inhibition of behaviors (Bar-Yam 1997). For example, in studies of Boolean networks (Kauffman 1993, analyzed in detail by McKelvey in this issue), order emerges when agents respond to inputs from just two other agents. A two-input system is homeostatic; when the system is perturbed from its attractor (a limited range of the network's state space that it occupies), it tends to return. When networks are more densely connected, so that each agent's behavior is influenced by the outputs of three or more agents, homeostasis collapses: stable states are delicate, because small changes to a few elements can send them careening off to a new attractor. (The only exception occurs when each element's schemata are so fixed that it produces the same output for virtually every combination of inputs it receives from other agents).

Applied to organizations, CAS models will require scholars to specify the pattern of connections among agents, not the pattern of connection among variables (as in path models). This has been done extensively in studies that draw upon social network theory and analytical techniques (Wasserman and Faust 1994). Social network analyses typically examine the structure of a network at a single point in time, but Frank and Fahrbach's paper (this issue) shows how they can be viewed as evolving dynamical systems. Although social network analysis is a well-established technique in organization science, most studies of networks focus on the presence or absence of a particular kind of tie between actors (Mizruchi 1994).

For CAS models, establishing the presence or absence of a tie between actors must be treated as a preliminary step. CAS models require us to specify how the behavior of an actor at time t influences the behavior at time t (or time $t + 1$ if there is a lag) of others with whom the actor has ties. A number of simulation models have taken this step (Burton and Obel 1998). However, there is no accepted, standard way to model organizational or interorganizational networks in the abstract, and the outcomes of many simulations are sensitive to small changes in the assumed structure of connections among actors. CAS modeling will require a melding of these two approaches, using empirical observations of behavioral ties at several time points as the basis for simulating how different types of networks will evolve.

Such models will also gain explanatory power when scholars take into account how a continuous injection of energy is necessary to sustain a pattern of interactions in a network. Most simulations abstract away the problem of how to energize the making, breaking, and maintenance of ties. They specify a particular pattern of interactions without assigning to each interaction a probability of occurrence related to the effort that agents allocate to it. Self-organization does not occur absent a continual flow of energy into a system, yet studies of how managers energize organizations have been divorced from inquiries into how pattern and structure emerge and evolve. The effort level of organizations waxes and wanes as managers propel them into new domains, bring new challenges and goals to the attention of members, make and break connections internally and externally, alter reward systems, and manipulate symbols. Understanding the causes and consequences of injecting energy into an evolving network of agents is an important topic for further research.

Coevolution to the edge of chaos.

Complex adaptive systems theories presume that the adaptation of a system to its environment emerges from the adaptive efforts of individual agents that attempt to improve their own payoffs. Consistent with notions of bounded rationality (March and Simon 1958), agents are presumed unable to forecast the system-level consequences of their individual choices, and so they optimize their own fitness, not that of the organization. Each agent is adaptive if its actions can be assigned a value (payoff, fitness) and the agent behaves so as to increase this value over time (Holland and Miller 1991). The landscape on which agents adapt continually shifts, because the payoffs of individual agents depend on the choices that other agents make (Levinthal 1997; McPherson and Ranger-Moore 1991). Agents (and clusters of agents that form stable subsystems) coevolve with one another, because changes in the distribution of behaviors among agents change individual fitness functions, and such shifts in turn alter behaviors (Robertson and Grant 1996). Local adaptations lead to the formation of continually evolving niches, so complex adaptive systems operate far from the equilibrium of globally optimal system performance (Holland and Miller 1991).

Morel and Ramanujam (this issue) argue that the apparent disequilibrium facing coevolving adaptive agents is actually a dynamic equilibrium. Bak (1996) has proposed that all complex adaptive systems evolve to a "critical state" that differs from traditional definitions of equilibrium. In an ordinary equilibrium state, small changes in the state of a system are self-correcting; the system quickly adjusts, and settles back into its attractor state(s). In the state of self-organized criticality, a dynamic equilibrium prevails, such that small changes in behavior can have small, medium, or large impacts on the system as a whole, according to a power law. Power laws can have different functional forms, but they imply that larger system changes occur exponentially less frequently than smaller ones do: in general, $y = \beta x^{\alpha}$ where y is the frequency of a change and x is the magnitude of the change. Bak's well-known experiment illustrating self-organized criticality involves dropping grains of sand onto a sand pile. The addition of a single grain usually causes very small landslides or cascades, but it also generates avalanches of all sizes, roughly with a frequency $1/x$ where x is the size of the avalanche.

Why would systems of interacting, coadapting agents evolve to this state? For example, what would explain the findings of Jørgensen et al. (1998), who contend that the relationship between body size and abundance for species in an ecosystem is a power law, and the frequency with which observed changes exceed a given change follows a power law? As Morel and Ramanujam discuss in this issue, many mathematical models generate power-law outcomes; such distributions are the stationary state of any stochastic process where the probability of an event is proportional to the number of times it has occurred in the past. The fact that a system generates changes that follow a power-law distribution does not demonstrate that it has evolved to the critical state and remains poised there. Two different lines of reasoning have led different scholars to conjecture that evolution drives complex adaptive systems to this state.

Kauffman (1995) argues that all complex adapting systems evolve to the edge of chaos, the point where small and large avalanches of coevolutionary change cascade according to a power law, because this state gives them a selective advantage: Systems that are driven to (but not past) the edge of chaos out-compete systems that do not, he suggests. Using the image of an adaptive landscape (discussed in detail by McKelvey and by Levinthal and Warglien, this issue), Kauffman suggests that if small changes in behavior lead only to small cascades of coevolutionary change, the system's performance can never improve much. On the other hand, if small changes in behavior lead to wildly different fitness levels (as occurs in chaotic environments), systems can reach extraordinary fitness peaks but cannot remain on them. The slightest change in behavior will send the system tumbling off its peak, perhaps plunging into a region of very low fitness. It is in the intermediate region that maximum system fitness will be found, Kauffman contends. Because fitness peaks tend to be located near each other, evolutionary search and selection can function efficiently. However, the occasional large coevolutionary cascade associated with small changes in behavior allows the system to leap to higher fitness peaks than it would likely locate through evolutionary refinement.

Building on these ideas, Brown and Eisenhardt (1998) have suggested that the most effective organizations evolve strategies that lie at the edge of chaos. Like Weick, (1979, p. 215), they argue that organizations can continue to exist only if they maintain a balance between flexibility and stability. Additionally, they contend that the strategic equilibrium over time for an organization is a combination of frequent small changes made in an improvisational way that occasionally cumulate into radical strategic innovations, changing the terms of competition fundamentally.

The second line of reasoning leading to the conjecture that complex adaptive systems naturally evolve to the edge of chaos follows Bak's observation that this state is the outcome of evolutionary processes that alter the fitness of the least-fit element of the system. Bak suggests that selection frequently replaces the weakest agent in a collectivity with one drawn randomly from a pool of candidates. Organizations tend to replace their least efficient members; the least effective firms in an industry tend to go bankrupt and be replaced by new entrants, and the most poorly adapted species in an ecosystem tends to become extinct. Ordinarily, a new element drawn randomly will have higher fitness than the weak one it replaced, setting off a cascade of coevolutionary adaptation. Bak demonstrates that in a wide variety of circumstances, these cascades follow a power law. It is this line of reasoning that leads both McKelvey and Morel and Ramanujam (this issue) to call for more investigation into organizational systems that evolve through the selecting-out of their weakest elements.

The idea that a system such as an organization will experience many small changes punctuated by infrequent, irregular, massive changes is familiar in organization theory (Gersick 1991). Most punctuated-equilibrium models set forth by organizational scholars rely on arguments that inertia builds up over time until the degree of misfit between an organization and its environment provokes a crisis (e.g., Cyert and March 1963, Gresov et al. 1993, Romanelli and Tushman 1994). Complexity theory does not invoke inertia to explain punctuated equilibrium. Rather, it suggests that a pattern over time of large and small changes is what one would expect from a system of coevolving agents subjected to selection pressures (as illustrated by Morel and Ramanujam's model, this issue). This does not invalidate inertial theories of punctuated equilibrium; it simply suggests a rival explanation that holds even in situations where inertia is weak.

The conjecture that agents coevolve on a fitness landscape to a state poised between order and chaos is an intriguing one, and Kauffman's "*NK*" adaptive landscape models have been very influential. Nearly a quarter of the 56 papers submitted for this special issue drew on Kauffman's models, far more than relied on any other approach to complex systems. McKelvey's paper (this issue) criticizes the *NK* model cogently and suggests improvements to it. Additionally, students of complex systems must come to grips with the problematic nature of the fitness function concept. In Kauffman's adaptive landscape metaphor (borrowed from Wright 1931), fitness is depicted as the z-axis on a three-dimensional landscape. Agents are depicted as climbing uphill toward higher fitness. In biology, where this mental image originated, fitness is a relatively unambiguous construct: the more offspring an organism contributes to the next generation, the fitter it is. However, organizational fitness is a much more complex affair.

Scholars who viewed organizations as natural systems noted early on that organizations must pursue maintenance goals as well as output goals (Scott 1992). Early institutional theorists (e.g., Selznick 1957) produced many rich case studies showing how organizations turned away from their original goals in response to environmental demands. Later institutional theorists (e.g., Meyer and Rowan 1977) pointed out that organizations must optimize much more than a single numeraire (e.g., profit) to survive and grow. Friedlander and Pickle (1968) showed that organizations that perform well on a criterion preferred by one constituency tend to do poorly on a criterion favored by another. Simon (1996) argued that even simple output goals are complex and multifaceted.

Agents at any level of analysis face far more complicated adaptive landscapes than CAS models have envisioned to date. Hill-climbing toward higher fitness on one measure may cause performance to deteriorate on others. The image of a rugged adaptive landscape presumes that conflicting selection pressures can somehow be aggregated into a single measure of performance. In reality, organizations and the individuals in them juggle a host of conflicting expectations and assessments that create a payoff function too difficult to assess and optimize (March and Simon 1958). Fitness is a complex combination of returns to exploitation, returns to exploration, and returns to reputation, market position, and capabilities built from past adaptations (A. Lewin et al. 1999). Additionally, many organizations fall considerably short of the frontier defining the highest fitness attainable, and the actions of firms move this frontier, leading to a cascade of changes within and among actors. Consequently, the adaptive landscape metaphor that underlies present-day studies of agents coevolving to the edge of chaos must not be pushed too far.

Recombination and System Evolution

As Simon (1996) has pointed out, any adaptive entity contains an adaptive inner environment; complex adaptive systems are nested hierarchies that contain other complex adaptive systems. These subsystems are therefore themselves subject to evolutionary pressures. Every aspect of a complex adaptive system—agents, their schemata, the nature and strength of connections between them, and their fitness functions—can change over time: new ones may appear, old ones may become extinct, and existing ones may survive in a fundamentally new form. Models of organizational life that build on CAS theories need not simply endow agents with schemata, connections, and adaptive behavior. They can also allow these elements to evolve.

A fundamental aspect of complex adaptive systems is that they allow local behavior to generate global characteristics that then alter the way agents interact (Burkhart, 1996). Actions not only proceed along feedback loops, they can also change these loops. In traditional causal models, the relationship between variables is presumed to be fixed, but in CAS models, the evolution of the network that links agents is an important object of theorizing and empirical observation in its own right. In organization science, studies of how social ties are broken and reconstituted (e.g., Palmer 1983, Stearns and Mizruchi 1986) provide particularly useful insights for scholars who wish to view organizations as complex adaptive systems.

Additionally, complex adaptive systems can evolve when new agents or new schemata are introduced. They may be drawn from a pool of candidates outside the system, or they can be generated by recombining elements of existing agents or schemata. In a number of CAS models (e.g., Holland 1995), the schemata of the most successful agents in a system are copied and then spliced together into a new schema, a process deliberately patterned after the recombination of chromosomes that takes place in biological reproduction. Levinthal and Warglien's paper (this issue) discusses recombination as a fundamental requisite for adaptation on rugged fitness landscapes.

In organization science, insights into the generation of novelty through recombination have been generated at several different levels of analysis. Technological innovations recombine elements of previous innovations (Fleming 1998, Kogut and Zander 1992). Groups, teams, and task forces integrate the ideas and attitudes of their members, and are arenas in which new ideas emerge from the interaction of their members. Joint ventures generate novelty by recombining skills and processes inherited from their parents. In some corporate mergers, a new entity can emerge that blends elements from several formerly independent companies. At the industry level, technological convergence can lead to the formation of new organizational communities that recombine elements of what were formerly distinct populations. These streams of research provide a rich foundation for modeling organizations as complex systems that evolve through the recombination of agents or their schemata.

Toward New Models in Organization Science

The development of systems theories that led to the open systems revolution in organization science was fueled by the development of new computing technologies. Procedural computer programming languages naturally accommodated models that linked variables together in complex feedback loops. Similarly, the study of complex adaptive systems has been facilitated by the emergence of new computational technologies. A technological shift toward distributed, decentralized computing gained momentum during the 1980s as local area networks diffused, and accelerated dramatically as the invention of the World Wide Web protocol caused an explosion in Internet access during the 1990s. Intellectual models that link individual adaptive agents linked together in networks of interaction have grown hand-in-hand with modular computing architectures that link independent processors and small programs together the same way.

Three new types of computer models have been used extensively to study complex adaptive systems: cellular automata, neural networks, and genetic algorithms. Each is the subject of its own large and growing literature, and therefore this article will not describe them in detail. Organizational scholars who are interested in complex adaptive systems need to understand at a high level how each has been used to model complexity, because each contributes ideas that can serve as building blocks for new approaches to modeling organizational life.

Cellular Automata. In a cellular automaton, each adaptive agent occupies a position on a lattice (a cell), surrounded by a set of neighboring agents (see Gutowitz 1991, for an introduction). The state of each cell depends on the state of those considered to be its neighbors. The key elements chosen by the modeler are the shape of the lattice (e.g., whether it is two- or three-dimensional, and whether cells are depicted as squares or some other regular shape), the choice of states a cell can occupy, the decision rule used to determine a cell's state, and the neighborhood that each agent observes in order to apply its decision rule. Lomi and Larsen (1997) built a cellular automata model that is a variant of John Conway's "Game of Life" (see Bar-Yam 1997 for a description), to show that the effects of density at founding on subsequent mortality rates are sensitive to assumptions about how the organizations are connected. Interestingly, of all cellular automata studied to date, only the "Game of Life" self-organizes into a critical state (Bak 1996).

Nowak and Vallacher (1998) provide an excellent overview of the strengths and limitations of cellular automata models for studying complex social systems. Cellular automata give the researcher great flexibility in specifying decision rules, so complex interactions can be modeled. However, each cell is constrained to interact with the same number of neighbors as every other cell; clearly, in organizations some individuals have many more ties than others do. Because they impose a rigid geometric structure on the pattern of ties in the network, cellular automata are not well-suited for modeling situations where only one actor in a neighborhood has a tie to an actor outside the neighborhood. An individual has to adopt a whole set of ties that map to its location; it cannot mix and match from the set of ties made feasible by its location. Because each cell can be assigned to a separate program or computer processor, it is possible for each to have its own unique schema; in practice, however, most cellular automata models to date have assumed that each agent applies the same decision rule as all other agents.

Neural Networks. A neural network model (also called a connectionist model) consists of a set of nodes called "neurons" that are connected to one another (see Anderson 1995 for an overview). As opposed to the geometric pattern imposed by cellular automata, any set of connections can be modeled. Each node uses an equation specified by the modeler to determine whether it should be activated, based on incoming signals from other neurons and its previous pattern of activation. Each connection between two nodes has a weight, influencing how strongly a signal from one node enters into the activation equation of the other. These weights change in response to experience, in a fashion determined by the modeler. Typically, modelers train a network by mapping a set of independent variables to a set of neurons, and feeding the network observations on these independent variables along with associated observations on the dependent variables. By modifying the strength of the connections among neurons, the network evolves a set of connections that is able to predict the value of the dependent variable, given a set of values for the independent variables.

Neural networks have been used extensively in business and economic forecasting applications, and have been proposed as a basis for studying network forms of organizing (Heydebrand 1989). They provide an interesting alternative to rule-based models for depicting agents' schemata. Future studies of organizations as complex adaptive systems might well model them as a network of partially connected, coevolving neural networks.

Genetic Algorithms and Classifier Systems. Genetic algorithms have principally been employed in computer science and operations research to solve optimization problems that were once considered intractable (see Holland 1995 for an overview; see Bruderer and Singh 1996 for an application to organizations). An agent is modeled as a set of one or more instruction strings, consisting either of a rule table or an automaton. Each string is assigned a fitness. Each period, strings are copied in proportion to their observed performance, and a new set of strings is generated by combining subsets of these strings at some random crossover point. Such strings are employed in classifier systems, which consist of rules. Rules compete with one another to post messages, but a rule can only post a message if it is activated by other messages. When a rule succeeds in posting its message, it transfers some of its fitness to the rules that contributed to activating it. As a result, each rule's strength depends on its past usefulness to other rules and the payoffs it receives from them.

Classifier systems have been used to manage bidding processes across a supply chain (Roy 1998). By allowing rules to bid and compete with one another, these models are able to evolve adequate solutions to problems involving very complex nonlinear dynamics. Genetic algorithms have also been used to breed useful software programs from chunks of competing code. Their implementation of a procedure for recombining rules to generate new rules is an important idea for organizational scholars to consider. More problematically, this technique relies on assigning fitness functions to strings; assigning fitness functions to organizations or their components is less straightforward. As Morel and Ramanujam note in this issue, genetic algorithms are appropriate when there is a known function to maximize; coevolution, in contrast, implies a constantly shifting definition of what an organization is trying to optimize.

Cellular automata, neural networks, and classifier systems that employ genetic algorithms all contribute important concepts for modeling complex adaptive systems. It is doubtful, however, that any of these approaches will supplant standard causal modeling, as scholars come to grips with the nonlinear dynamics of organizations. What is needed is an approach that melds empirical observation with the computer's power to simulate the many possible paths through which complex networks of interacting agents can evolve.

Segel (1995) draws an interesting contrast between models, which attempt to enhance our understanding of a system by representing it in terms of mathematical equations, and simulations, which attempt to reproduce through a computer program how a system behaves in a given set of circumstances. A modeler strives for simplicity at the expense of realism, while a simulator strives for realism at the expense of simplicity. It seems apparent that simulation is an essential tool for modeling a set of complex, changing interactions over time. Simulation in organizations has made enormous strides (see Carley 1995 for an overview), and modern object-oriented programming methods provide a natural way to model agents, their schemata, and their interconnections (Zeggelink et al. 1996). Moss et al. (1998) describe an object-oriented programming language, SDML (strictly declarative modeling language), that represents agents as models of cognition within organizational structures. Other models such as Santa Fe's SWARM simulation system (profiled on the Internet at `http://www.santafe.edu/projects/swarm/`) may also simplify the problem of modeling organizations that have rich interactions among their components.

Organization science has historically advanced by combining theoretical with empirical research. A limitation of simulations is that many equally plausible structures can lead to very different predictions, and a given outcome can be explained equally well by a host of simulations with very different assumptions. The power of simulation as a technique is its ability through many iterations to explore a variety of paths through which a system might evolve, given a structure of partially connected, coevolving agents that possess changing schemata. Such simulations need not be based on abstract specifications of how agents behave and interact; they may be seeded with real data from real actors.

R. Lewin et al. (1998) suggest that modeling complex adaptive systems involves identifying agent characteristics, the dimensions of relationships among the agents, and the figures of merit that govern their coevolution. Qualitative field-based work can produce candidate parameters for each of these elements, they argue. Connectionist models that incorporate actual data have already been developed (Read and Miller 1998, see Varkas-Duong 1998 for a critique). Simulation allows us to see what emerges when agents whose behavior and cognitive structure we assess empirically interact with one another through a set of connections that we assess empirically, but that can change over time as part of the model.

What might a future empirical study look like, that introduces a new way of thinking about modelling complexity? It may well try to develop and test a theory that tries to explain an empirical regularity observed in standard causal-modeling research. Instead of asking which other independent variables seem to be significantly and causally related to the outcome, it will ask what model of interacting might lead to the observed outcome in dynamic equilibrium, and what other outcomes would be predicted from such a model.

For example, organizational scholars were to observe that the longer an organization has been in existence, the less likely it is to pioneer a radical innovation. One approach to building on this observation has been to gather data on what organizational characteristics seem to be correlated with pioneering behavior. A different approach might model radical innovation as the outcome of interaction among a variety of organizations that pursue better technical performance in coevolutionary competition with one another. The empirical data one might gather from an actual population would include:

(1) Who are the agents? How many organizations compete in this space, and what are their salient demographic characteristics?

(2) What are the agents' schemata? A researcher might use survey or observational or even archival data to model a set of competing cognitive structures that determine what innovations each agent pursues and how it reacts to the efforts of other agents. Each firm might be modeled as a set of (perhaps fuzzy) decision rules, or one might train a neural network to mimic each firm's response to a given set of inputs.

(3) How are the agents connected? How do these connections change over time?

(4) What payoff functions do these agents pay attention to? What tradeoffs are they willing to make among different types of payoffs?

(5) How do the actions of one agent affect the payoffs of others? What is the payoff structure of the evolutionary game they appear to be playing?

The first step of the study would be to build a model based on theoretical assumptions that incorporates these empirical data. The second step would be to demonstrate that the model can simulate the trajectory of innovation in the population observed to date. The third step would be to make novel predictions based on the model that move beyond received theory. The fourth step would be to make predictions about what outcomes would be observed were key elements of the model to change. What would happen, for example, were the pattern of connections or the number of agents or the payoff structure to be altered significantly? A fifth step would be to build several competing models that are all consistent with observed data, but that lead to different predictions. By applying data from the evolving population or from other populations, we can test the predictions of different models, and progress theoretically by discarding models whose predictions do not hold, perhaps synthesizing new models from their most successful elements.

Such a modeling approach requires a combination of data that has never been collected in a single study. It also requires familiarity with simulation techniques, although the type of toolkit represented by the SDML and SWARM systems promises to make sophisticated simulation technology more broadly accessible. It is unlikely that a single investigation that captures all four key features of complex adaptive systems will be carried out in the short run. Rather, complex adaptive systems thinking will penetrate organization science through a series of middle-range theories that gradually build up a new generation of models that have testable implications. The purpose of this special issue is to accelerate the development of such models, and the papers it contains provide important conceptual foundations and research directions for scholars interested in coming to grips with the nonlinear dynamics of organizational life.

New Directions in the Strategic Management of Organizations

In addition to suggesting new ways to model nonlinear, dynamic behavior in organizations, complex adaptive system theory has rich implications for the strategic management of organizations. A combination of institutional and technological factors has created a trajectory since the end of World War II toward greater social connectedness. New technologies have expanded the geographic and product/market scope of many enterprises, while factors such as deregulation and global trading institutions have broken down old barriers that once isolated organizational populations. The environment that organizations face is characterized by many interactions among organizations and institutions, creating complex, nonlinear relationships between actions and outcomes.

Such environments are hypercompetitive (D'Aveni 1994; Illinitch et al. 1998); their nonlinearity leads both to unpredictable behavior and a rapid rate of change, because changes in one agent's behavior reverberate to influence others in a chain reaction. Unlike systems with fixed-point or cyclical equilibria, theirs is a more dynamic equilibrium in which actions can lead to small, medium, or large cascades of adjustment. For this reason, the aim of organizations' strategy is to evolve temporary advantages more rapidly than competitors can (Brown and Eisenhardt 1998). As McKelvey (this issue) points out, complexity theory is particularly relevant for organizations facing rates of external change that exceed their internal rate change.

In environments far from equilibrium, where cascades of change are constantly playing out and overlapping with one another, adaptation must be evolved, not planned. Adaptation is the passage of an organization through an endless series of organizational microstates that emerge from local interactions among agents trying to improve their local payoffs. The task of those responsible for the strategic direction of an organization is not to foresee the future or to implement enterprise-wide adaptation programs, because nonlinear systems react to direction in ways that are difficult to predict or control. Rather, such managers establish and modify the direction and the boundaries within which effective, improvised, self-organized solutions can evolve (Meyer et al. 1998). They set constraints upon local actions, observe outcomes, and tune the system by altering the constraints, all the while raising or lowering the amount of energy injected into the dissipative structure they are managing.

Brown and Eisenhardt's application of complexity theory to strategic management suggests that single business units achieve rapid evolutionary progress through improvisational moves based upon a few rules, responsibilities, goals, and measures. Synergies among business units follow when every unit has a distinct role (with none as central controller), and collaboration is focused on a few key areas. Evolution proceeds most rapidly, they argue, when senior managers effect small, cheap probes in a characteristic rhythm, recombining the elements of a portfolio of modular business units, so that novelty is deliberately generated without destroying the best elements of past experience.

Managers seeking to design and tune such systems, which succeed through their superior capability to evolve complex adaptive behaviors, have two principal levers at their disposal. First, they can alter the fitness landscape for local agents; second, they can reconfigure the organizational architecture within which agents adapt. In both cases, the strategist operates on agents indirectly, taking advantage of the tendency for myriad local interactions to self-organize into a coherent pattern. Rather than shaping the pattern that constitutes a strategy (Mintzberg and Waters 1982), managers shape the context within which it emerges (Burgelman 1991).

By altering the fitness landscape on which individual agents are trying to adapt, strategists can change both the trajectory of emergent behavior and the diversity of behaviors in an organization's repertoire. As Levinthal and Warglien note (this issue), managing a fitness landscape involves much more than the design of incentive mechanisms. Altering the reward system is an important way in which managers can shape the flow of behaviors in an organization, and in particular, they must counteract the tendency of managers to favor exploiting short-term opportunities over exploring riskier ones that have a more distant and less certain payoff (March 1991). Reward systems must also overcome classic agency problems that arise in the course of interaction, such as compensating individuals for team production, encouraging agents to exert themselves on behalf of their principals, and overcoming social dilemmas. In addition, however, fitness landscapes can be altered by operating on the context within which a reward system operates.

The most direct way in which executives operate on the context of reward systems is by choosing the organization's domain. A firm's fitness landscape depends directly on the strategic choice of its niche, its way of making a living. By altering the shape or location of the niche, executives propel their organizations into arenas that channel emergent behavior in novel directions.

Because a fitness landscape results from nested coevolution at several different levels (A. Lewin et al. 1999), executives operate upon an indirect ecological system within which ideas and initiatives compete by managing variation, selection, and retention among schemata (Anderson 1999). An agent's interpretation of the fitness landscape depends on the schemata s/he is following, and the nature and distribution of schemata in an organization is subject to managerial influence. Social cues interact with reward systems to help agents decide which signals deserve their limited attention.

In addition, as noted above, organizational fitness is an ambiguous construct. Consequently, managers can highlight different goals and measures at different times, and can alter the tradeoff that agents make among conflicting dimensions of fitness. In this way, they can encourage more exploration or more exploitation, by emphasizing different aspects of performance. In many organizations, managers operate several different performance measurement systems, and by altering the degree to which these systems overlap or diverge, executives can adjust the ruggedness of the adaptive landscape that agents face.

In his or her role as organizational architect, the strategist influences the extent of improvisation, the nature of collaboration, the characteristic rhythm of innovation, and the number and nature of experimental probes by changing structure and demography (Meyer et al. 1998). When agents are added to, deleted from, or recombined within a network, a coevolutionary cascade results; in dynamic equilibrium, some of these cascades will result in large-scale adaptation, allowing a continuous series of small changes to generate evolution in a punctuated equilibrium. A similar pattern of adjustment can occur when ties between agents are made, broken, or altered in strength and sign. Managers of complex systems can only dimly foresee what specific behaviors will emerge when an organization's architecture is changed. Instead of relying on foresight, they rely on evolution; changes that produce positive cascades of change are retained, while those that do not are altered.

Future research in strategic management must give executives guidelines to follow in evolving networks of agents. For example, Burt's (1980) analysis of social networks suggests that creating structural holes is a way to generate more novelty and innovation. There does not yet exist a theory that will help managers predict even generally the type of emergent outcome that will result from altering the configuration of a network in a particular way, and crafting such a theory should be important element of the research agenda in strategy.

In a similar vein, managers can indirectly influence the emergence of adaptive behavior by altering the distribution of agents in a network. Changing the demography of an organization will alter the pattern of behavior that emerges from it (Lawrence 1997). Although a link between diversity and innovative behavior seems reasonably well-established, much more research will be required to help strategists think about how to guide the strategic evolution of an enterprise by making specific types of demographic changes.

Conclusion

Organization theory has not yet caught up with the sophisticated tools that have emerged for analyzing the behavior of complex adaptive systems. We are not on the verge of a revolution that will render a century of organization theory obsolete, but remarkable new vistas are opening up, thanks to the melding of the science of complexity and organization theory and the increasing availability of new techniques for modeling nonlinear behavior. Those who take advantage of this opportunity will lead us to think in new ways about what kind of organizational data to gather and what kind of models to construct. They will also generate a new wave of theory in strategic management that focuses successfully managing strategic and organization change and how managers lead and influence the never ending journey of adaptation.

Organization theory has historically borrowed from a number of parent disciplines. Because complexity theory has developed along a very interdisciplinary path, it may be that in the end, organization theory contributes as much as it borrows to the development of insight into the behavior of complex systems. Many modern organizations are complex adaptive systems par excellence, and we who study them should eventually lead instead of follow efforts to understand the fundamental nature of nonlinear, self-organized structures.

Acknowledgment

Many individuals have contributed to the author's explorations in the science of complexity. He wishes to especially thank Alan Meyer, Bill McKelvey, Dan Levinthal, and Kathleen Carley for their insights, suggestions, and patience and Mike Lawless for many hours of discourse and intellectual nourishment. Special thanks go to Arie Lewin for encouraging this special issue for his insightful comments on earlier drafts of this paper, and for his editorial guidance and leadership.

References

Allison, Mary Ann, Susanne Kelly, 1999. *The Complexity Advantage: How the Science of Complexity Can Help Your Business Achieve Peak Performance.* McGraw-Hill, New York.

Anderson, James A. 1995. *An Introduction to Neural Networks.* MIT Press, Cambridge, MA.

Anderson, Philip (1999, in press). Seven levers for guiding the evolving enterprise. John Clippinger, ed. *The Biology of Business: Decoding the Natural Laws of Enterprise.* Jossey-Bass, San Francisco.

Arthur, W. Brian. 1989. Positive feedback in the economy. *Scientific American.* **262** (2) 92.

Ashby, Ross (1956), *An Introduction to Cybernetics.* London: Chapman and Hall.

———. 1981. *Mechanisms of Intelligence: Ashby's Writings on Cybernetics.* Intersystems Publications, Seaside, CA.

Bak, P. 1996. *How Nature Works: The Science of Self-Organized Critically.* Copernicus, New York.

Barnard, Chester I. 1938. *The Functions of the Executive.* Harvard University Press, Cambridge, MA.

Bar-Yam, Yaneer. 1997. *Dynamics of Complex Systems.* Addison-Wesley, Reading, MA.

Beniger, James R. 1986. *The Control Revolution: Technological and Economic Origins of the Information Society.* Harvard University Press, Cambridge, MA.

Boulding, Kenneth. 1956. General systems theory—The skeleton of science. *Management Sci.* **2** 197–208.

Brown, Shona L., Kathleen M. Eisenhardt. 1998. *Competing on the Edge: Strategy as Structured Chaos.* Harvard Business School Press, Boston.

Browning, Larry D., Janice M. Beyer, Judy C. Shetler. 1995. Building cooperation in a competitive industry: SEMATECH and the semiconductor industry. *Acad. Management J.* **38** 113–151.

Bruderer, Erhard, Jitendra V. Singh. 1996. Organizational evolution, learning and selection: A genetic-algorithm-based model. *Acad. Management J.* **39** 1322–1349.

Burgelman, Robert A. 1991. Intraorganizational ecology of strategy-making and organizational adaptation: Theory and field research. *Organ. Sci.* **2** 239–262.

———. 1994. Fading memories: A process theory of strategic business exit in dynamic environments. *Admin. Sci. Quart.* **29** 24–56.

Burkhart, Roger. 1996. The SWARM simulator: Applications in science and business. *Embracing Complexity: Exploring the Application of Complex Adaptive Systems to Business.* Ernst & Young Center for Business Innovation, Cambridge, MA.

Burt, Ronald. 1980. Autonomy in a social topology. *Amer. J. Soc.* **85** 892–925.

Burton, Richard, Børge Obel. 1998. *Strategic Organizational Diagnosis and Design: Developing Theory for Application.* Kluwer Academic Publishers, Boston, MA.

Campbell, Donald T. 1974. Downward causation' in hierarchical systems. Ayala, Francisco J. and Theodosius Dobzhansky, eds. *Studies in the Philosophy of Biology.* MacMillan, London. 179–186.

Capra, Fritjof. 1996. *The Web of Life: A New Scientific Understanding of Living Systems.* Anchor Books, New York.

Carley, Kathleen M. 1995. Computational and mathematical organization theory: Perspective and directions. *Comput. Math. Organ. Theory* **1** (1) 39–56.

Casti, John. 1994. *Complexification: Explaining a Paradoxical World Through the Science of Surprise.* HarperCollins, New York.

Cohen, Jack, Ian Stewart. 1994. *The Collapse of Chaos: Discovering Simplicity In a Complex World.* Viking, New York.

Coveney, Peter V., Roger Highfield. 1995. *Frontiers of Complexity: The Search for Order In a Chaotic World.* Fawcett Columbine, New York.

Cowan, George A., David Pines, David Meltzer. 1994. *Complexity: Metaphors, Models and Reality* Addison-Wesley, Reading, MA. 17–29.

Cyert, R. M., James G. March. 1963. *A Behavioral Theory of the Firm.* Prentice-Hall, Englewood Cliffs, NJ.

Daft, Richard L. 1992. *Organization Theory and Design.* (4th ed.) West Publishing, St Paul, MN.

———, Arie Y. Lewin. 1990. Can organization studies begin to break out of the normal science straitjacket: An editorial essay. *Organ. Sci.* **1** (1).

D'Aveni, Richard A. 1994. *Hypercompetition: Managing the Dynamics of Strategic Maneuvering.* Free Press, New York.

Dooley, Kevin J., Andrew Van de Ven. 1999—this issue. Explaining complex organizational dynamics. *Organ. Sci.* **10** (3).

Drazin, Robert, Lloyd Sandelands. 1992. Autogenesis: A perspective on the process of organizing. *Organ. Sci.* **3** (2) 230–249.

Epstein, Joshua M. 1997. *Nonlinear Dynamics, Mathematical Biology, and Social Science.* Addison-Wesley, Reading, MA.

Fleming, Lee. 1998. Exhaustion in technological evolution. Unpublished doctoral dissertation, Stanford University.

Fontana, Walter, Susan Ballati. 1999. Complexity. *Complexity* **4** (3) 14–16.

Forrester, Jay. 1961. *Industrial Dynamics.* MIT Press, Cambridge, MA.

Frank, Ken A., Kyle Fahrbach. 1999—this issue. Organization culture as a complex system: Balance and information in models of influence and selection. *Organ. Sci.* **10** (3).

Friedlander, Frank, Hal Pickle. 1968. Components of effectiveness in small organizations. *Admin. Sci. Quart.* **13** 289–301.

Galbraith, Jay R. 1982. *Designing Complex Organizations.* Addison-Wesley, Reading, MA.

Gell-Mann, Murray. 1994a. *The Quark and the Jaguar: Adventures in the Simple and the Complex.* W. H. Freeman, New York.

———. 1994b. Complex adaptive systems. George A. Cowan, David Pines, David Meltzer eds. *Complexity: Metaphors, Models and Reality:* Addison-Wesley, Reading, MA. 17–29.

Gersick, Connie J. G. 1991. Revolutionary change theories: A multilevel exploration of the punctuated equilibrium paradigm. *Acad. Management Rev.* **16** 10–36.

Gresov, Christopher, Heather A. Haveman, Terence A. Oliva. 1993 Organizational design, inertia and the dynamics of competitive response, *Organ. Sci.* **4** 181–208.

Gutowitz, Howard. 1991 *Cellular automata: Theory and experiment.* MIT Press, Cambridge, MA.

Haberstroh, Chadwick J. 1965. Organization design and systems analysis. James G. March, ed. *Handbook of Organizations.* Rand McNally, Chicago, 1171–1211.

Heino, Mikko. 1998. Noise colour, synchrony and extinctions in spatially structured populations. *Oikos* **83** 368–375.

Heydebrand, Wolf V. 1989 New organizational forms. *Work and Occupations* **16** 323–357.

Heylighen, Francis. 1997. Publications on complex, evolving systems: A citation-based survey. *Complexity* **2** (5) 31–36.

Holland, John H. 1995. *Hidden Order: How Adaptation Builds Complexity.* Addison-Wesley, Reading, MA.

———, John H. Miller. 1991. Artificial adaptive agents in economic theory. *Amer. Econom. Rev. Papers and Proceedings* **81** 365–370.

Horgan, John. 1995. From complexity to perplexity. *Sci. Amer.* **272** (6) 104.

Illinitch, Anne Y., Arie Y. Lewin, Richard D'Aveni. 1998. New organizational forms and strategies for managing in hypercompetitive environments. *Managing in Times of Disorder: Hypercompetitive Organizational Responses.* Sage Publications, Thousand Oaks, CA. Introduction xxi-xxxiv.

Jørgensen, S. E., Henning Mejer, Søren Nors Nielsen, 1998. Ecosystem as self-organizing critical systems. *Ecological Modelling* **111** 261–268.

Katz, Daniel, Robert L. Kahn, 1978. *The Social Psychology of Organizations,* 2nd ed. Wiley, New York.

Kauffman, Stuart A. 1993. *The Origins of Order: Self Organization and Selection In Evolution.* Oxford University Press, New York.

———. 1995. *At Home in the Universe: The Search for Laws of Self-Organization and Complexity.* Oxford University Press, New York.

Kogut, Bruce, Udo Zander, 1992. Knowledge of the firm, combinative capabilities, and the replication of technology, *Organ. Sci.* **3** 383–397.

Koput, Kenneth W. 1997. A chaotic model of innovative search: some answers, many questions. *Organ. Sci.* **8** 528–542.

Lawrence, Barbara S. 1997. The black box of organizational demography, *Organ. Sci.* **8** 1–22.

Lee, Allen. 1990. Integrating positivist and interpretive approaches to organization science. *Organ. Sci.* **1** (1) 342–365.

Levinthal, Daniel A. 1997. Adaptation on rugged landscapes. *Management Sci.* **43** 934–950.

———, Massimo Warglien. 1999 this issue. Landscape design: designing for local action in complex worlds. *Organ. Sci.* **10** (3).

Lewin, Arie Y., Chris P. Long, Timothy N. Carroll. 1999. The coevolution of new organization forms. Working paper.

Lewin, Roger, Theresa, Parker, Birute Regine. 1998. Complexity theory and the organization: Beyond the metaphor. *Complexity* **3** (4) 36–40.

Lomi, Allesandro, Erik R. Larsen. 1997. Density delay and organizational survival: Computational models and empirical comparisons. *Comput. Math. Organ. Theory* **3** (4) 219–247.

Mainzer, Klaus. 1994. *Thinking in Complexity: The Complex Dynamics of Matter, Mind, and Mankind.* Springer-Verlag, New York.

March, James G. 1991. Exploration and exploitation in organizational learning. *Organ. Sci.* **2** 71–87.

———, Johan P. Olsen. 1976. *Ambiguity and Choice in Organizations.* Universitetsforlaget, Bergen, Norway.

———, Herbert A. Simon. 1958. *Organizations.* John Wiley, New York.

Mead, George Herbert. 1934. *Mind, Self and Society.* University of Chicago Press, Chicago.

Meyer, Alan, Peter Frost, Karl Weick. 1998. Overture—The *Organization Science* jazz festival: Improvisation as a metaphor for organizing. *Organ. Sci.* **9** (5) 569–576.

Meyer, J., B. Rowan. 1977. Institutionalized organizations: Formal structure as myth and ceremony. *Amer. J. Soc.* **83** 340–363.

McKelvey, Bill. 1997. Quasi-natural organization science. *Organ. Sci.* **8** (4) 352–380.

———. 1999—this issue. Avoiding complexity catastrophe in coevolutionary pockets: Strategies for rugged landscapes. *Organ. Sci.* **10** (3).

McPherson, J. Miller, James R. Ranger-Moore. 1991. Evolution on a dancing landscape: Organizations and networks in dynamic blau space. *Soc. Forces* **70** 19–42.

Mintzberg, Henry, James A. Waters. 1982. Tracking strategy in an entrepreneurial firm. *Acad. Management J.* **25** 465–499.

Mizruchi, Mark S. 1994. Social network analysis: Recent achievements and current controversies. *Acta Sociologica* **37** 329–343.

Morel, Benoit, Rangaraj Ramanujam. 1999—this issue. Through the looking glass of complexity: The dynamics of organizations as adaptive and evolving systems. *Organ. Sci.* **10** (3).

Moss, Scott, Helen, Gaylard, Steve Wallace, Bruce Edmonds. 1998 SDML: A multi-agent language for organizational modeling. *Comput. Math. Organ. Theory* **4** (1) 43–70.

Nowak, Andrzej, Robin R. Vallacher. 1998. Toward computational social psychology: Cellular automata and neural network models of interpersonal dynamics. Stephen J. Read, Lynn C. Miller, eds. *Connectionist Models of Social Reasoning and Social Behavior* Lawrence Erlbaum, Mahwah, NJ. 277–311.

Ott, Edward. 1993. *Chaos in Dynamical Systems.* Cambridge University Press, Cambridge, U.K.

Palmer, Donald. 1983. Broken ties: Interlocking directorates and intercorporate coordination. *Admin. Sci. Quart.* **28** 40–56.

Paul, David L., John C. Butler, Keri E. Perlson, Andrew B. Whinston. 1996. Computationally modeling organizational learning and adaptability as resource allocation: An artificial adaptive systems approach. *Comput. Math. Organ. Theory* **2** (4) 301–324.

Perrow, Charles. 1967. A framework for the comparative analysis of organizations. *Amer. Soc. Rev.* **26** 854–866.

Prigogine, Ilya, Isabelle Stengers. 1984. *Order Out of Chaos: Man's New Dialog with Nature.* Bantam Books, New York.

Read, Stephen J., Lynn C. Miller. 1998. *Connectionist Models of Social Reasoning and Social Behavior.* Lawrence Erlbaum, Mahwah, NJ.

Ripa, J., P. Lundberg. 1996. Noise colour and the risk of population extinctions. *Proceedings of the Royal Society of London B.* **263** 1751–1753.

Robertson, Douglas S. Michael C. Grant. 1996. Feedback and chaos in Darwinian evolution. *Complexity* **2** (1) 10–14.

Romanelli, Elaine, Michael L. Tushman. 1994. Organizational transformation as punctuated equilibrium: An empirical test. *Acad. Management J.* **37** 1141–1166.

Rousseau, Denise M., Robert J. House, 1994. Meso organizational behavior: Avoiding three fundamental biases. *J. Organ. Behavior* **1** 13–31.

Roy, Beau. 1998. Using agents to make and manage markets across a supply web. *Complexity* **3** (4) 31–35.

Rumelhart, David E. 1984. Schemata and the cognitive system. Robert S. Wyer Jr., Thomas K. Srull, eds. *Handbook of Social Cognition.* Hillsdale, Lawrence Erlbaum, Hillsdale, NJ. 161–189.

Samuel, Yitzhak, Chanoch Jacobsen. 1997. A system dynamics model of planned organizational change. *Comput. Math. Organ. Theory* **3** 151–171.

Scott, W. Richard. 1992 *Organizations: Rational, Natural and Open Systems.* Prentice-Hall, Englewood Cliffs, NJ.

Segel, Lee A. 1995. Grappling with complexity. *Complexity* **1** (2) 18–25.

Selznick, Philip. 1957. *Leadership in Administration: A Sociological Interpretation.* Row, Peterson, Evanston, IL.

Accepted by Professor Arie Y. Lewin, Founding Editor-in-Chief Emeritus.

Simon, Herbert A. 1996. *The Sciences of the Artificial,* 3rd ed. MIT Press, Cambridge, MA.

Stearns, Linda Brewster, Mark Mizruchi. 1986. Broken-tie reconstitution and the functions of interorganization interlocks: A reexamination. *Admin. Sci. Quart.* **31** 522–539.

Sterman, John D., Jason Wittenberg. 1999—this issue. Path dependence, competition, and succession in the dynamics of scientific revolution. Organ. Sci. **10** (3).

Thietart, R. A., B. Forgues. 1995. Chaos theory and organization. *Organ. Sci.* **6** (1) (January-February) 19–31.

Thom, Rene. 1975. *Structural Stability and Morphogenesis; An Outline Of A General Theory Of Models.* W. A. Benjamin, Reading, MA.

Thompson, D. 1967. *Organizations in Action.* McGraw-Hill, New York.

Vakas-Duong, Deborah. 1998. Review of *Connectionist Models of Social Reasoning and Social Behavior. J. Artificial Societies Soc. Simulation.* **1** (4) `http://www.soc.surrey.ac.uk/JASSS/1/4/review2.html`.

von Bertalanffy, Ludwig. 1968. *General System Theory: Foundations, Development, Applications.* G. Braziller, New York.

Waldrop, M. Mitchell. 1992. *Complexity: The Emerging Science at the Edge of Order & Chaos.* Simon and Schuster, New York.

Wasserman, Stanley, Katherine Faust. 1994. *Social Network Analysis: Methods and Applications.* Cambridge University Press, Cambridge, U.K.

Weibull, Jörgen W. 1995. *Evolutionary Game Theory.* MIT Press, Cambridge, MA.

Weick, Karl E. 1979. *The Social Psychology of Organizing.* Addison-Wesley, Reading, MA.

Weisbuch, G. 1991. *Complex Systems Dynamics: An Introduction to Automata Networks.* Addison-Wesley, Reading, MA.

Wright, Sewall. 1931. Evolution in Mendelian populations. *Genetics* **16** 97–159.

Zeggelink, Evelien P. H., Reinier Van Oosten, Frans N. Stokman. 1996. Object oriented modeling of social networks. *Comput. Math. Organ. Theory* **2** (2) 115–138.