

Organizations can be characterized as complex systems composed of adaptive and intelligent agents. Organizational adaptation occurs through restructuring and learning. Organizations can be modeled using a dual-level model in which restructuring is modeled as a simulated annealing process and individual learning is modeled using a stochastic learning model and boundedly rational agents. Such a model is presented, and its behavior is illustrated using a virtual experiment where the type of organizational adaptation is varied. Results suggest that the organizational design and performance relationship may be chaotic, despite the simple rules of change. Simple restructuring rules lead to a wide range of emergent organizational structures that increases with individual adaptation. Organizations locate good designs (through chance and slow change) regardless of the agents' intelligence; however, emergent designs depend on adaptability. Design features are not systematically related to performance; rather, small initial differences in design and environment can affect the emergent behavior.

Modeling Organizational Adaptation as a Simulated Annealing Process

KATHLEEN M. CARLEY
DAVID M. SVOBODA
Carnegie Mellon University

Studies of organizational adaptation often point out that organizational "knowledge" exists on at least two levels: in the minds of the personnel and in the design of the organization (Huber 1991). Adaptation can occur at both the individual (March and Simon 1958; Hastie 1986; Levitt and March 1988) and the design level (Eccles and Crane 1988; Lawrence and Lorsch 1967; Hannan and Freeman 1977; DiMaggio and Powell 1983). However, most studies of organizational learning do not make a distinction between these levels (Kim 1993), or they consider only one of these two modes of adaptation (Lant 1994). At the individual level, organizational learning can be seen as involving the accumulation of experience by personnel within the organization (Carley 1991a, 1992). At the design level, Eccles and Crane (1988) in their comprehensive analysis of investment banking suggest that over time, "reorganizations initiated by top

SOCIOLOGICAL METHODS & RESEARCH, Vol. 25 No. 1, August 1996 138-168
© 1996 Sage Publications, Inc.

management resemble the annealing process used in crystal formation" (p. 143). Herein, organizational adaptation is explored from a combined individual and design perspective, using a simulated annealing model of organizational adaptation in which the agents within the organization are themselves complex adaptive agents.

Organizational theorists have long recognized the value of formal models (mathematical, logical, computational) for examining organizational behavior in general and organizational adaptation in particular (Cyert and March 1963; see also Lant [1994] and Carley [1995b] for a review of this area). Computational models are an important, and appropriate, device for building theory about adaptive agents and organizational change. In the past decade, the use of computational models to explore issues of social and organizational change has grown tremendously. These models demonstrate the value of computational techniques for studying complex adaptive phenomena. For example, it has been demonstrated that simple interaction among individuals has a strong and nonlinear impact on the effective generation of social sanctions, cues, and norms (Macy 1990; Horgan 1994), the evolution of social stability, and information diffusion (Carley 1991b, 1995a); that environmental and institutional factors such as payoffs, population dynamics, and population structure influence the evolution of cooperation in a discontinuous fashion (Axelrod and Dion 1988); and that structure can, under certain conditions, mitigate the effect of individual learning and turnover on organizational performance (Carley 1992) and alter decision outcomes (Cohen, March, and Olsen 1972). Many of these models demonstrate the ability of complex social and organizational phenomena to emerge from simple processes (Prietula and Carley 1994; Padgett forthcoming).

The models of social or organizational adaptation tend to fall into two camps—single-agent learning models and genetic models. The learning models use either stochastic models of learning (Macy 1990, 1991; Carley 1991a, 1992), strategic adaptation (Lant 1994), or more complex artificial intelligence models of human cognition and learning (Carley, Park, and Prietula 1993; Carley and Prietula 1993; Verhagen and Masuch 1994). This research assumes that social and organizational change either results from amalgamating individual learning or is the result of the organization (or society) as a single agent learning to respond to its environment. The primary problem with these learn-

ing models as models of organizational adaptation is that they ignore the simultaneous role of individual cognition and institutional factors and organizational strategies in effecting organizational adaptation.

The genetic models use genetic algorithms to model social and organizational adaptation (Crowston 1994, forthcoming; Holland and Miller 1991). This research focuses on social and organizational change purely as the result of evolutionary processes such as permutation and crossover. There are several problems with these genetic models as models of organizational adaptation. First, they generally do not track the performance of specific individuals in the organization and cannot examine individual-level effects such as agent learning and turnover. Second, genetic algorithm models of organizational or social adaptation typically represent either the organization or an organizational strategy by a string of numbers. Permutation of strings and crossover among them result in either new organizations or new designs. If each string represents a new organization, then the assumption being made is that organizations as such do not adapt but rather die and spawn new organizations. That is, it is the population of organizations and not the specific organization that adapts. This assumption is fine under some conditions; however, it belies the fact that organizations, such as the YMCA (Zald and Denton 1963), do adapt sometimes quite dramatically. In contrast, when each string represents a design, then the assumption being made is that organizations adapt by generating new populations of designs. The difficulty is that there is no way of mapping designs to organizations and so no way of knowing how a specific organization adapted. If the argument is made that the population of designs as a whole represents a single organization, then it is being assumed that the organization has the capability of simultaneously employing and evaluating multiple designs, which is patently not the case. An exception here is the work of Padgett (forthcoming) in which strategies are directly tied to individuals by specifying a physical landscape. The limitations of these two standard approaches suggest the need for a new model of organizational adaptation.

In this article, the issue of organizational adaptation is addressed theoretically by using a computational model in which both the personnel learn from experience and the organization adapts by altering its design. This research couples individual learning with organi-

zational restructuring. At the individual level, learning is carried out using a standard stochastic learning model for boundedly rational agents based on work in cognitive psychology. At the organizational level, adaptation is carried out using a simulated annealing model. Simulated annealing is a recently developed method for solving difficult combinatorial optimization problems (Kirkpatrick, Gelatt, and Vecchi 1983; Rutenbar 1989).

In the following section, the basic theoretical stance is presented by describing the features considered necessary for modeling organizational adaptation when organizations are viewed as complex systems composed of complex adaptive agents. This is followed by a description of simulated annealers. The fit of annealers to the theoretical conception of organizational adaptation is described. The potential value of this type of model for increasing our understanding of organizational change is illustrated via a virtual experiment. The focus of this article is on demonstrating the potential values of simulated annealing to organizational theory; the results from the virtual experiment should be treated as indicative rather than conclusive. This virtual experiment looks at the interplay between individual learning and organizational restructuring on performance.

ADAPTIVE ORGANIZATIONS

THEORETICAL STANCE

Organizational adaptation should contain elements of change due to executive decisions (what will hereafter be referred to as restructuring strategies) and due to personnel gaining experience. When organizations choose a strategy for change, only a limited number of strategies are considered, and they are not selected by comparing them consistently one against the next (Allison 1971; Cohen and March 1974; Axelrod 1976). Rather, a strategy is considered, and if it looks as if it will improve things or meet some goal, it is adopted with little consideration as to whether or not this strategy more than all other available strategies actually moves the organization closer to the long-term goal. In other words, each organizational decision is a satisficing rather than an optimizing decision (Simon 1944; March and

Simon 1958). Sometimes, the organization is even willing to take risks and adopt a strategy that is thought to be poorer in the short term in the hopes that it will have long-term benefits. Once a strategy is adopted, it is put into effect, and it leads to a basically nonreversible change in the organization's design. Within computer science, such a change process would be referred to as destructive state modification.

Organizations are trying to locate the optimal design for their environment and will make changes in their design accordingly. Organizational strategies for altering organizational design may take on a number of forms (e.g., for example, personnel hiring procedures, downsizing procedures, task reassignment procedures, and personnel reassignment procedures). What the organization considers to be the optimal design depends on what it perceives as important (e.g., minimizing salary, maximizing the number of widgets produced, or maximizing decision accuracy). Essentially, the organization is seen as trying to minimize or maximize some cost function having to do with performance. In a nonchanging environment, the organization will eventually find a design that results in better, if not optimal, performance. In theory, if the set of strategies considered is sufficient, if there is an optimal design for that environment, if the organization does not change its cost function, and if the environment does not change, then the organization should eventually be able to locate the optimal design. For organizations, when they are founded or when a new technology comes into play or when the environment shifts radically, the organization will make a great number of strategic changes, some of which may be high risk. As the organization matures, the number of such changes, particularly those that are high risk, decreases. That is, over time, organizations become more staid and less likely to adopt changes that do not have obvious foreseeable benefits; they become locked into a certain way of doing business. This staidness has been characterized in terms of competency traps and in terms of locking into a technology that is too expensive to radically change.

Within the organization, individuals have skills and work on tasks. As individuals work on tasks, they gain experience (March 1981). Individuals, however, have cognitive and physical limitations and so act in, at most, a boundedly rational fashion (Simon 1955, 1956; Carley and Newell 1994). Individuals may enter or leave the organization, may move about in the organization and so report to or manage

different people, and may change what tasks they are doing and hence their access to different resources and information. Individuals can learn from experience, but what they learn and its applicability in future situations are affected by their position in the organization and the tasks they do. Changes in the individual's position alter what is learned and the applicability of experience.

SIMULATED ANNEALING AND ORGANIZATIONAL ADAPTATION

Simulated annealing is a heuristic solution strategy for combinatorial optimization problems, originally introduced in 1982 by Kirkpatrick et al. (1983; for an overview, see Rutenbar 1989). Simulated annealing is a computational analog of the physical process of annealing a solid. In a physical system, the goal is to find a configuration of atomic particles (state) that minimizes the energy (cost) of the system. Annealing is done by allowing the system to reach thermodynamic equilibrium at each temperature in the schedule of decreasing temperatures. The system is heated to a state that admits many alterations and then cooled slowly until it freezes in a good configuration. This procedure alters an unordered solution into a highly optimized solution.

Simulated annealing techniques are particularly valuable for combinatorial optimization problems in which there are a large variety of variables. Such problems try to minimize some function $f(x)$ such that x is determined given a number of parameters p_1 through p_n . Many such problems are NP-complete—that is, known techniques for obtaining an exact solution require an exponentially increasing number of steps as the problem becomes larger and so cannot be exactly solved in a reasonable time. Consequently, heuristic techniques need to be used. However, heuristic techniques such as simulated annealing are not guaranteed to generate the correct answer, to find the true optimum. For these combinatorial optimization problems, there is a performance surface that is the value of $f(x)$ by the n different parameters. This surface has many peaks and valleys, and it may be difficult to locate the lowest valley. Simulated annealing can be thought of as a technique for rolling a ball along this surface to locate the lowest valley. The ball starts out randomly in some location (the initial state) and rolls downhill. Importantly, now and then, this process is per-

turbed, and the ball is moved to a different location in the landscape and then allowed to roll downhill again. The uphill moves are generally carefully controlled so that the ball does not get near the lowest valley only to be randomly plucked up and moved.

Simulated annealing captures this process by proposing a new state, determining the change in cost from the old to the new state, using the Metropolis criterion to determine whether or not to adopt the new state, making the change if called for, and repeating this process until the temperature reaches "freezing." According to the Metropolis criterion, a new state will always be accepted if it decreases the cost and will be accepted with some probability that decreases as the temperature decreases. The probability of making a cost-increasing move is defined by the boltzman equation. Let ΔF be the change in performance, T the temperature, and P the probability of accepting the cost-increasing move. The boltzman equation can be written as $P = P_0 e^{-\frac{\Delta\Phi}{T}}$. The initial probability of accepting the cost-increasing move, P_0 , is referred to as the initial partition.

What new state is proposed depends on the move set. The move set is a set of operations that the annealer can take to generate a new state. A new state is proposed by applying a move chosen at random from the set of possible moves. The annealer is in only one state at a time (i.e., it employs destructive state modification as the new accepted state replaces the old). Cost is calculated using a user-specified cost function. Initial temperature and the temperature cooling schedule are set relative to the problem domain. The cooling schedule is a sequence of decreasing temperatures such that the initial effective temperature is sufficient to admit random search in the performance space. At the coldest temperature, the state is in its close-to-final form, and few moves are allowed.

Simulated annealing is an appealing model for organizational adaptation due to strategic change, as its parameters directly translate into aspects of the theory previously described. States are organizational designs. The move set is the set of acceptable strategies for changing a design (e.g., a strategy or move might be to hire someone). Organizations have only one design (state) at a time. Organizations try to optimize their design relative to some cost function, even as simulated annealers try to optimize their state relative to some cost

function. Organizations do not compare all strategies and select the best; rather, they pick a strategy, and if it appears to move them closer to the goal, they satisfice and employ that strategy. Similarly, the annealer does not evaluate all moves but picks one and always moves to the new state if it decreases the cost. Organizations take more risks initially (i.e., they employ strategies that may appear to move them away from their goal in the short term) than they do later on. Similarly, the annealer takes more risks (i.e., moves to states that increase the cost) initially when temperature is high than it does later when temperature is low.

*A SIMULATED ANNEALING MODEL
OF ORGANIZATIONAL ADAPTATION*

Imagine a single organization. This organization is composed of a set of agents, a set of information on a problem, connections between agents, and connections between agents and information (see Figure 1). The organization is faced with a sequence of tasks. Each time period, the organization must make a decision for the current task. This organizational decision-making process involves a series of individual decision-making processes and requires some agents to process the information on that task and report to other agents. The organization's ability to adapt its structure (the connections) is modeled as an annealing process, such that the organization's restructuring strategies are the move set. The agents' ability to adapt is modeled as a stochastic learning process for boundedly rational agents.

An organization at time (t), O_{at} , is composed of $N_o(t)$ personnel and I_o pieces of information (which can be thought of as the raw information needed for solving the problem or some type of specific resource). The organizational design is defined by the organizational structure (who reports to whom), $S(t)$, and the resource access structure (who has access to what resources), $R(t)$. The specific restructuring strategies, or move set, are defined as a set of actions that the organization can take. The set of actions are (1) *fire* (fire n personnel such that $1 \leq n \leq N_o(t)$), (2) *hire* (hire n personnel such that $1 \leq n \leq N_{max} - N_o(t)$), (3) *retask* (delete the tie between agent i and resource s then add a tie between agent j and piece of information s , and (4) *reassign* (delete

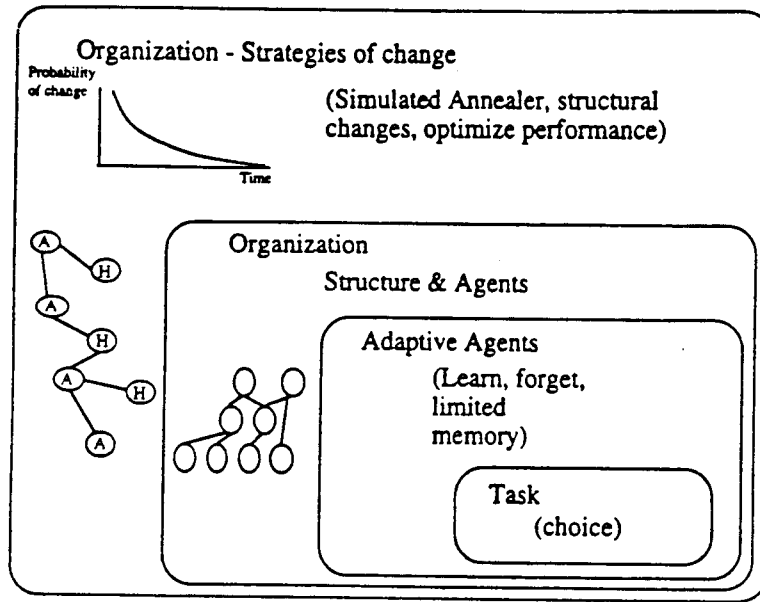


Figure 1: Overview of the Dual-Level Model

the tie between agent i and j (i reports to j) and reassign agent i to report to agent k . Exactly how many personnel are hired (or fired, retasked, or reassigned) at a time is given by a Poisson distribution.¹

The organization begins with a particular design, $S(t)$ and $R(t)$, and goes through a training period for 500 tasks. After this, its life cycle begins (Figure 2): general operation, evaluate performance, suggest design, evaluate design, decide whether to accept design, change design. After training, the organization commences on its life cycle. First, given a sequence of tasks (general operation), performance is calculated over a sequence of tasks (evaluate performance). Then, an executive suggests a new design (a strategy from the move set), looks ahead for 100 tasks and tries to imagine how the proposed new design will impact performance (evaluate design), and then decides whether to accept the change according to the Metropolis criterion. The limited look ahead is simulated by creating a hypothetical organization with the proposed new design and simulating its performance on a sequence of tasks. If the design change is accepted, the organization puts the

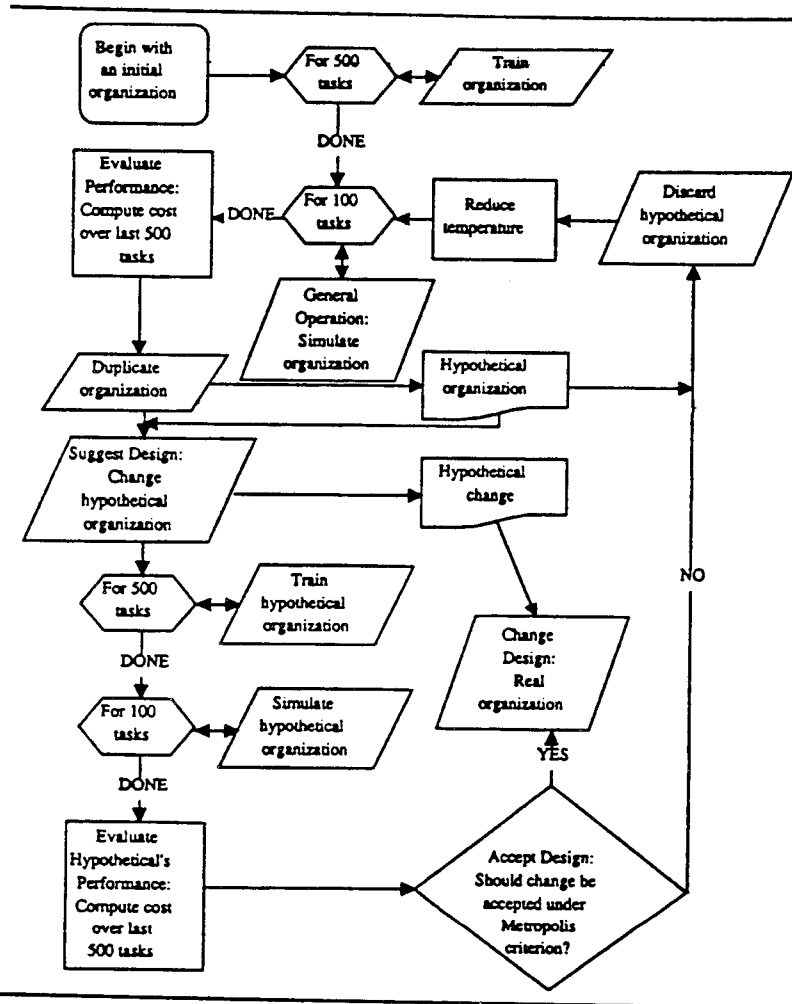


Figure 2: Organizational Life Cycle

change in place and then proceeds to process another sequence of tasks, at which point another design is suggested. If a design change is not accepted, the organization goes on as it is for another sequence of 100 tasks. The rate of change is set by the temperature cooling schedule. Temperature will drop every 100 tasks (time periods) as $T(t + 1) = \alpha * T(t)$, where α is the rate at which the organization becomes risk averse and t is time.

The probability of accepting a new design (a strategy from the move set) is based on the Metropolis criterion. According to this criterion, the organization always accepts the change if the resulting hypothetical organization is a better performer than the current organization. Otherwise, the change is accepted with a probability given by $e^{-\Delta cost(t)/T}$ such that $cost(t) = 0 - performance(t)$. *Performance(t)* is the percentage of tasks in the past 500 that the organization correctly calculated.

The organization is doing a binary classification task—given a set of independent task features (each of which can take on two values, true or false), determine whether the object is of type A or B. This task in either its binary or its trinary version has been widely studied by organizational researchers (Hollenbeck, Ilgen, Sego, Hedlund, et al. 1955; Hollenbeck, Ilgen, Tuttle, and Sego 1995; Tang, Pattipati, and Kleinman 1992, 1993; Pete, Pattipati, and Kleinman 1993, 1994) and is analogous to a wide number of tasks in actual organizations (Carley 1992). The specific task is an unbiased decomposable task in which the true decision is an A if there are more 1s than 0s in the set of task features, and B otherwise. Organizational performance is measured by accuracy, the percentage of problems in a window of opportunity that the organization correctly classifies as an A or B.

The agents in the organization are experiential learners (Carley 1992). They classify the pattern of incoming information that they see and then select as their decision the choice that was most often correct in the past for that pattern. If an agent has no previous experience with this pattern, the agent simply guesses. For an agent, the incoming information can be either or both raw information on this specific task or the decisions of other agents for this task. If an agent sees x pieces of information, then the number of patterns the agent tracks is 2^x . After each task, the agent receives feedback on what was the correct choice for the organization for that task. This information is used to increment the agent's memory. If the organizational design is changed, then the agent may see fewer or more pieces of information and adapts accordingly.

All agents are boundedly rational. First, they do not remember all tasks they have ever seen. Rather, they suffer from both a primacy and a recency effect and thus remember only the first 500 tasks and the most recent 500 tasks they have seen. Second, agents can handle only a maximum of seven pieces of information. Third, agents do not

remember exactly what happened on a particular task; rather, they remember general trends. The result of these limitations is that no one agent can do the task him- or herself; each agent is effectively overconfident in his/her decisions, and the organization's memory (number of tasks it bases its decision to adapt on) is greater than the memory of any one individual (number of tasks each individual remembers). Further, the agents are bound by their positions in the organization. That is, the position determines to what information the agent has access.

The organization's decision is the decision made by agent or agents in the top tier in the organization. If there is only one agent at the top, then the organization's decision for a specific task is that agent's decision. If there are multiple agents at the top, then the organization's decision is the result of an equally weighted vote by the agents at the top. If both options receive an equal number of votes, then the organization's decision is chosen randomly from the two options.

VIRTUAL EXPERIMENT

Theoretical computational models, like mathematical models, are normative. They show how the society or organization would behave in the ideal case—that is, when the conditions specified in the model are met. Computational techniques are turned to when the problem is not tractable—for example, when it is NP-complete. Typically, with computational models, it is not possible to characterize all of the model's behavior. Rather, a general sensitivity analysis given a set of parameters is generally done. This examination can take the form of a virtual experiment. From these results, predictions for organizational behavior can be induced.

For this article, three virtual experiments were run using this computational model—only individual learning, only structural learning, and dual-mode learning. In the only-individual-learning condition, organizations could not change their organizational design over time; however, individuals could and did learn. Only the stochastic learning model for boundedly rational agents is used. Organizations are given a randomly determined initial design and did not restructure. In the only-structural-learning condition, organizations tried to opti-

mize their organizational design over time in response to environmental feedback; however, individuals could not learn and instead acted as majority classifiers² (which is what the majority of them would have acted like had they been able to learn and were given sufficient time to learn). Only the simulated annealing model of organizational restructuring is used. In the dual-mode condition, organizations tried to optimize their organizational design over time in response to environmental feedback, and individuals learned on the basis of experience. For each of the organizational learning conditions, the following constraints were set.

Organizations were limited to having at most three levels of personnel with at most 15 individuals on each level and nine pieces of information on the task. The organization's, $O_{(t)}$'s, organizational structure $S(t)$ was represented as a 45×45 matrix such that $S_{ij}(t) = 1$ if i reported to j (by convention $S_{ii}(t) = 0$). The organization cannot have fewer than one agent. Also, by convention, for the three levels (A, B and C), agents cannot report to agents at the same level, agents in Level C can report to agents in either Level A or Level B, and agents in Level B can report to agents in Level A. No other types of reporting are allowed. These constraints enable hierarchies, teams, and matrix structures with dual-reporting options to form. Similarly, the organization's task structure, $R(t)$, was represented as a 45×9 matrix such that $R_{is}(t) = 1$ if i used resource s in making i 's decision. These constraints mean that the maximum possible number of organizations is 21,080. In total, 1,000 different organizations were simulated, for each condition, such that their initial designs (these matrices) were randomly generated from the set of all possible organizational designs.

A task is characterized by a nine-bit binary string such that each bit can be thought of as a different resource. There is a population of 36 distinct tasks. Each organization is simulated for 20,000 tasks (time periods) or until it reaches quiescence, whichever comes first. An organization is said to reach quiescence if the probability of accepting a new design drops to approximately 55 percent (this is a freezing temperature of 0.0345). The initial temperature is set to 0.433, which means that initially approximately 90 percent of the changes are accepted.³ The rate at which the organization learns to be risk averse is set such that $\alpha = .975$.

The timing of the various stages in the organization's life cycle is controlled by a number of windows. Each window is defined by a number of tasks. The training window is set to 500 tasks. These occur prior to the 20,000 tasks over which the organization is simulated. The 20,000 tasks are divided into a series of 200 cooling windows, each of length 100 tasks. Temperature is dropped each cooling window. Organizational performance is calculated over the last 500 tasks. The proposed design is simulated for 100 tasks, and its performance is calculated over this look-ahead window.

In the individual and dual-mode conditions, individuals learn from experience. Individuals update their experience after each task. Individual memory is based on the first 500 tasks the agent sees and the last 500 tasks. This gives a slight bias to recent information. This slight tendency on the part of individual agents to attend more to recent information enables the individual agents to learn to respond to changes in the organizational structure. These values mean that the individual agents have a shorter time horizon than does the organization—that is, they base their decision on only 350 tasks, whereas the organization bases its decision on 500 tasks. Because the organization can change every 100 tasks, potentially, the organization could be moving personnel about and interfering with their memory. In the structural and dual-mode conditions, organizations adapt by hiring, firing, retasking, or redesigning. All four types of moves are initially equally likely.

Organizational performance and design is characterized by the last performance window (final 500 tasks of the 20,000 examined). The following measures of organizational design, change, and performance are captured: size, span of control, redundancy, isolates, ignorance, hires, fires, retaskings, reassignments, and performance. A variety of measures are captured for both the initial 1,000 organizations and the final 1,000 organizations under each condition. Both the mean and variance are reported. Those measures that are used in the following illustrative results are as follows:

Size: Average number of personnel in the organization, N .

Span of control: Average number of personnel supervised by a manager, $S(t)$.

Redundancy: Average number of ties in the resource access structure, $R(t)$.
 Isolates: Average number of agents not connected to any other agent in the organizational structure.

Ignorance: Average number of pieces of information not connected to any agent.

Hires: Average number of agents hired over the last 20,000 tasks.

Fires: Average number of agents fired over the last 20,000 tasks.

Retaskings: Average number of ties between agents and resources that were changed over the last 20,000 tasks.

Reassignments: Average number of ties between agents that were changed over the last 20,000 tasks.

Performance: Average percentage of the 500 tasks that the organizations correctly classify during the final performance window.

DESIGN AND PERFORMANCE

THE PERFORMANCE SURFACE

The first question addressed is simply, What is the optimal organizational design given the cognitive and structural limitations previously identified? If the individuals and the structure were not limited, then the optimal organizational form would be a single individual with perfect recall. However, the limitations exist. To locate the optimal form, a sampling approach was taken. A random sample of 1,000 organizations each varying in terms of N , $S(t)$, and $R(t)$ were selected. Each was then simulated for 1,000 tasks, and their performance was calculated over the 1,000 tasks. No design changes were allowed. All agents have equilibrated and learned as much as they could learn. The resultant performance of these organizations by size and span of control is shown in Figure 3.

If all agents in the organization were simply guessing, then the organization would make the correct decision 50 percent of the time. Note that the better, more optimized designs are rare and are surrounded by a plateau of designs that are approximately equally good and all of which exhibit performance over that achievable via guessing (50 percent). Figure 3 was generated under the individual-learning condition; however, the general features of this figure are shared by the other learning conditions. These features are few top performers, all top performers surrounded by a plateau of equally good second-best

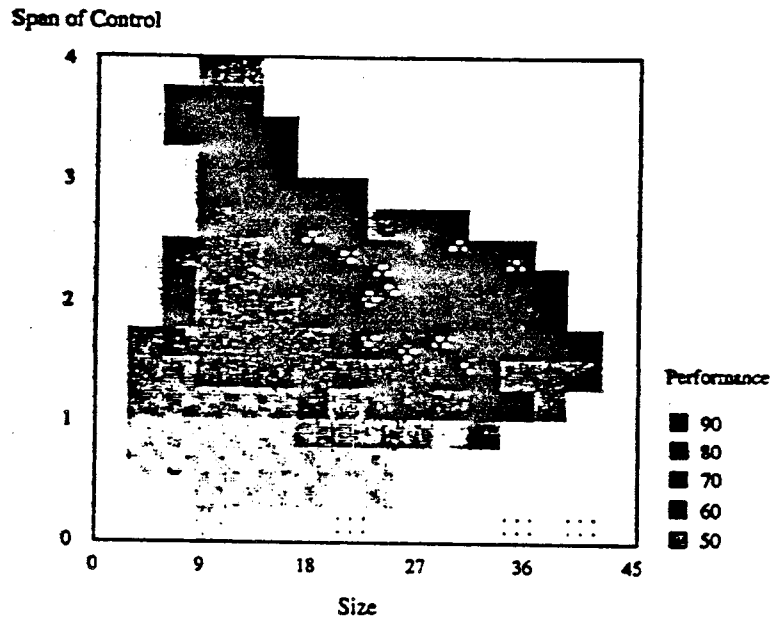


Figure 3: Contour Plot of Performance by Span of Control and Size for Individual-Learning Condition


structures, and a general increase in performance with size. Not one of the generated designs, regardless of the type of learning, exhibits perfect performance. It is theoretically possible for certain designs to exhibit perfect performance. One such design is the team of nine individuals, each handling a different piece of information and together making their decision by voting. Such a structure just happens to never occur in any of the runs completed.

Under structural and dual-mode learning, the performance surfaces are similar to that in Figure 3. In Figure 4, the performance surfaces for these two other learning conditions are shown. Note that the pattern of the surface is somewhat different but that the basic features of equi-performance plateaus and an increase in performance with size hold.

MODE OF LEARNING AND ORGANIZATIONAL PERFORMANCE

For the condition in which only individuals learn and all organizations begin in a state in which the individual agents are simply

Span of control



Dual-Mode Learning

Span of control




Figure 4: Contour Plot of Performance by Span of Control and Size, for Structural and Dual-Mode Learning Conditions

guessing, performance is 50 percent. The final performance, on average, for organizations in which all changes are due to individual learning is 76.14 percent. When final performance is calculated, the personnel have reached equilibrium—that is, they have learned all that they can learn. Thus the average increase in performance attributable to individual learning (over guessing) is 26.14 percent.

No design exhibits perfect performance. The designs exhibiting the best performance vary in the number of individuals at each level and the exact pattern by which individual agents are connected. An important point here is that in an organization in which individuals can learn, traditional measures of structure do not control performance. To be sure, structural features of the organization impact performance; however, they do not account for all of the variance in performance. This point will be returned to in a later section. Much of the variation in organizational performance, and there is a great deal, is due to differences in what individuals learn conditioned upon their position in the organization and the variation in the specific problems faced by the organization. As to the first point, small variations in the amount of information an organizational agent attends to and in the quality or accuracy of the reports to an agent by subordinates can have major organizational ramifications. Structures differ in performance not just because they alter the flow of information and so affect specific decisions but also because they alter the flow of information and so affect what the agents within the organization learn.

In the structural learning mode, organizations are essentially acting as procedural organizations in which all individuals follow a series of procedures that lead them to basically act as majority classifiers. Thus, when new personnel are hired and when individuals are retasked or reassigned, there are no start-up or retraining costs. An organization composed of these particular procedural agents is acting like an organization composed of individuals who have learned all that they can learn at all points in time. Under structural learning, the average final performance is 79.69 percent. Recall that the average final performance under the condition of only individual learning was 76.14 percent. This suggests that the increase in performance uniquely attributable to structural learning is 3.55 percent. On average, individual learning is thus a much greater contributor to organizational

learning than is structural learning. Finally the landscape of the performance space, peak performance areas surrounded by plateaus with approximately equally successful organizational forms, is not conducive to organizations learning structurally. That is, the organization can search through the space trying a variety of forms and yet not happen on to the specific form with high performance.

Under dual-mode learning, new personnel are hired, and when individuals are retasked or reassigned, there are start-up or retraining costs. New personnel have no prior experience and so begin by just guessing. Thus it is possible that structural learning might interfere with individual learning (and possibly vice versa). Under dual-mode learning, the average final performance is 79.46 percent. Contrast this with performance in the last two conditions—individual learning at 76.14 percent and structural learning at 79.69 percent. What this suggests is that structural learning might interfere with individual learning; however, the effect may not be noticeable.⁴ Organizations can learn even when they restructure.

Other performance characteristics of the space of organizations by mode of learning can be considered (see Table 1). These results suggest that although structural learning may have little impact on overall performance, it does improve the organization by reducing the number of isolates and the amount of ignorance. Structural learning impedes the ability of some organizations to reap the full benefits of individual experience, as it results in fewer top-performing organizations (organizations with a performance levels greater than or equal to 90 percent). Further, structural learning lets the field of organizations improve. That is, although the average performance has decreased, the probability that any organization will witness at least a 60 percent improvement in performance is higher in the structural and dual-mode conditions. Thus, although individual learning has the biggest initial impact on organizational learning, structural learning enables the organization to fine-tune its behavior. Another way of saying this is that managers can always restructure the organization. Such changes may have little or no long-term effect, but they will increase the appearance of improved performance by increasing efficiency (decreased isolates and pockets of ignorance).

TABLE 1: Mode of Learning and Performance

<i>Measures of Performance</i>	<i>Mode of Learning</i>		
	<i>Individual</i>	<i>Structural</i>	<i>Dual</i>
Average performance	76.14%	79.69%	79.46%
Pr(80% + improve)	.0210	.0100	.0060
Pr(60% + improve)	.4080	.5600	.5280
Number of top performers	21	10	6
Average isolates	1.87	1.14	1.25
Average ignorance	0.42	0.21	0.24

NOTE: Pr(80% + improve) = the probability that any organization will witness at least an 80 percent improvement in performance; Pr(60% + improve) = the probability that any organization will witness at least a 60 percent improvement in performance.

ORGANIZATIONAL DESIGN AND ORGANIZATIONAL PERFORMANCE

The simple rules for change that have been investigated lead to a wide range of emergent organizational structures. These organizations vary in size, redundancy, and span of control (see Table 2). As can be seen in Table 2, when organizations can engage in structural learning, the emergent structures are larger, are less redundant, and have a lower span of control than do the initial structures. The initial structures have the same properties as the organizations in which only individual learning is allowed, but over time, the structures that emerged are more finely tuned. Importantly, the range of structures increases (the standard deviations are higher) when the organization changes structurally. Structural learning can increase the variance in organizational form even when all organizations are responding to the same environment and are all trying to optimize their organizational form for high performance.

All agents in all organizations are receiving identical feedback. Performance differentials at the organizational level, given a particular mode of learning, are due to information absorption as decisions flow up the organizational hierarchy from Level C to B to A, to differences in the complexity of the specific subtask faced by the different agents (as they handle more resources), and to the extent to which decisions are made by "averaging" individual agent's decisions. One question

TABLE 2: Mode of Learning and Variance in Design

<i>Design Characteristics</i>	<i>Mode of Learning</i>		
	<i>Individual</i>	<i>Structural</i>	<i>Dual</i>
Size	20.95 (7.32)	26.60 (9.22)	27.00 (9.69)
Redundancy	1.60 (0.41)	1.49 (0.48)	1.45 (0.50)
Span of control	1.39 (0.50)	1.11 (0.52)	1.09 (0.53)

NOTE: $N = 1,000$. Standard deviations are in parentheses.

is whether there is an interaction between what designs are optimal and the mode of learning.

In Table 3, the results of a regression showing the impact of various aspects of design on performance (using standardized beta coefficients) for each mode of learning are shown. When organizations do not engage in structural learning, the impact of design is what might be expected. Specifically, performance is higher the greater the organizational size, the task density, and the span of control, and the fewer the number of isolates and pockets of ignorance. Of these factors, size and isolates have the most effect. In other words, organizations can most improve performance by getting rid of deadwood and increasing in size. This advice is fine, however, only if the organization is in the position of not laying off personnel and allowing personnel to act on the basis of their experience.

Now consider what happens when the organization engages in structural (or dual) learning. In this case, the impact of design changes. Size remains the dominant determinant of performance. Importantly, the impact of span of control reverses. That is, under only individual learning, the larger the span of control, the better the organizational performance. However, when the organization can alter its structure, the lower the organizational span of control, the higher the performance. Organizations that can adapt structurally can embed learning in the connections and need not rely on managerial experience in sorting through agents' differing perceptions. The final important difference is that when organizations can restructure, the impact of ignorance is more devastating.

TABLE 3: Regression of Design Features on Organizational Performance by Mode of Learning

Design Feature	Learning Mode		
	Individual	Structural	Dual
Size	.370***	.390***	.490***
Redundancy	.087**	.284***	.300***
Span of control	.191***	-.102***	-.189***
Isolates	-.308***	-.026	-.086***
Ignorance	-.052*	-.225***	-.189***
R ²	.363	.265	.315

NOTE: Values shown are standardized betas. $N = 1,000$.
 * $p < .05$; ** $p < .01$; *** $p < .005$.

These results suggest that the relationship between organizational design and performance, despite the simplicity of the rules for change, may be chaotic. Organizations with almost equivalent organizational forms can exhibit very different performance. Organizations as complex systems have the characteristic that the organizational design features are not systematically related to performance; rather, small differences in organizations' initial design and their response to the environment can result in dramatically different emergent structures and final performance. Thus, although organizations can locate good designs regardless of the intelligence of the agents within them, what designs emerge depends on the interaction between structural and individual adaptability.

AMOUNT OF CHANGE AND ORGANIZATIONAL PERFORMANCE

As both the number of individuals the organization hires or fires and the number of retaskings and reassignments increase, the rate of change in the organization effectively increases. Under both structural and dual-mode learning, the amount of change tends to be related to the final level of performance the organization achieves (see Figure 5). Except for reassignments, less organizational change generally ultimately results in higher organizational performance. With respect to reassignment, the more the organization changes, the higher its performance. This suggests that rather than structural learning impeding individual learning (as is the case for hiring and firing), individual

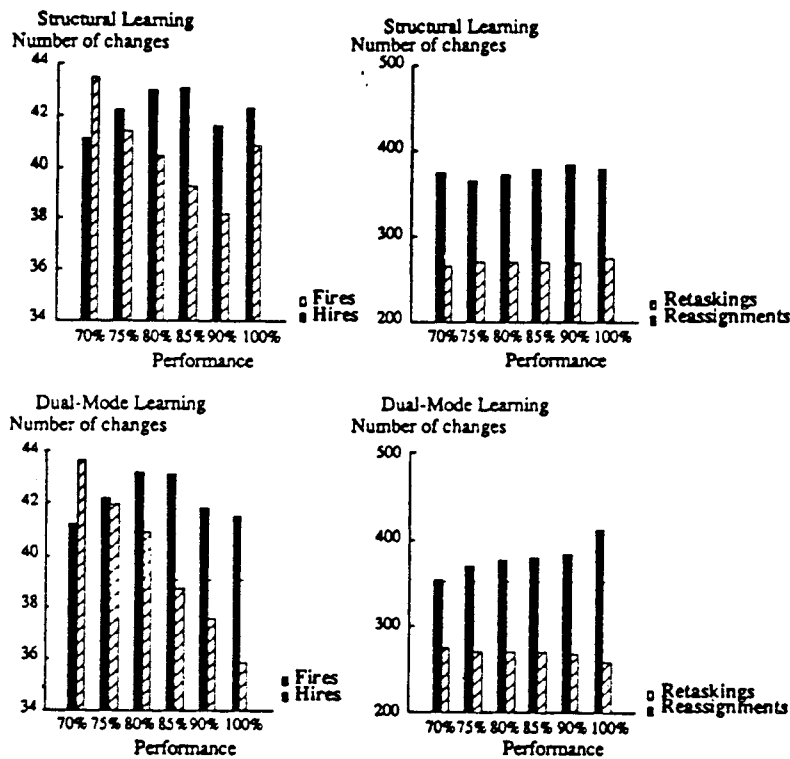


Figure 5: Amount of Organizational Change and Final Performance

learning facilitates the organization's ability to improve performance through redesign.

Under both modes of learning, upsizing the organization may reduce the likelihood that the organization will achieve high levels of performance. Upsizing the organization is not guaranteed to improve the performance and may be as likely to reduce performance as enhance it, unless the individuals in the organization are just following procedures (structural learning case). Regardless of whether individuals can learn, downsizing typically leads to improved performance; however, when individuals are simply following procedures, the top organizations do not exhibit the greatest level of downsizing.

Retasking has less of an effect on performance than does reassignment. However, lower levels of retasking, when individuals can learn, tend to correspond to higher performance. In contrast, higher reassignments correspond to higher final performance. In general, there is less of a relationship between rate of change and performance when agents are acting procedurally than when they can learn. Organizations composed of nonadaptive agents, therefore, have greater latitude in changing their structure. Structural changes have less of an impact when individuals cannot themselves adapt.

THE DESIGN OF TOP PERFORMERS

Finally, let us consider the organizational design of those organizations exhibiting the highest performance (greater than 90 percent performance) under the different modes of learning. Table 4 shows the number of agents at each level and the average number of links from each level to the one above. In the structural and dual-mode conditions, as opposed to the individual-learning condition, the degree of hierarchy (in terms of division of labor) is higher, and the amount of work that the individuals at the lowest level do (Level C) is higher. We see this increased hierarchy in the fact that when structural learning is admitted, agents at Level C report only to Level B. In addition, under the dual-mode condition, neither does Level C report to Level A nor is basic information assigned to Level A. When individual learning or dual-mode learning is allowed, the organizations are on average smaller. This characterization for a stylized three-person organization is shown in Figure 6.

DISCUSSION

Computational models are one of the few ways of developing a coherent and consistent theory of organizational change. The results from these models can be thought of as predictions that can then be studied in other settings. Field studies of organizational learning have difficulty finding organizations that are sufficiently similar and that are not undergoing periods of adaptation. Lab studies cannot consider organizations of sufficient size. Further, in human organizations,

TABLE 4: Design of Top Performers

	<i>Learning Mode</i>		
	<i>Individual</i>	<i>Structural</i>	<i>Dual</i>
Size	27	28	25
Number of agents at Level C	10	8	8
Number of agents at Level B	7	12	7
Number of agents at Level A	9	9	10
Average number of links between task and Level C	2	3	4
Average number of links between task and Level B	1	1	2
Average number of links between task and Level A	1	1	0
Average number of links between Levels C and B	2	1	1
Average number of links between Levels C and A	1	0	0
Average number of links between Levels B and A	1	1	1

performance indicators are typically difficult to collect. Computational techniques obviate these problems and allow the researcher to build and explore a theoretical model of change. Future empirical studies can then test the predictions that are generated from these computational models.

In this article, the cost function used was simply in terms of optimizing performance. Alternate cost functions, such as simultaneously optimizing performance and minimizing personnel, could be considered. Work on computational models suggests that changing the cost function can affect the particularities of the results. In this case, it would affect the number of isolates, the level of ignorance, and so forth. However, changing the cost functions is unlikely to change the basic results: that individual learning had a greater impact than structural learning, that the type of structural change influences behavior, that there are equi-performance plateaus, that there are few high performers, and that the small changes in organizational form can result in dramatic changes in performance. Nevertheless, future studies should investigate alternate cost functions.

This article also links into the area of complexity theory, chaos, and complex adaptive agents. Herein, organizations are modeled as complex adaptive agents composed of complex adaptive agents. As in any complexity model, stable structures and patterns of behavior emerge from the application of simple rules on the part of the agents, and agents coevolve. Rate of learning influences what patterns emerge.

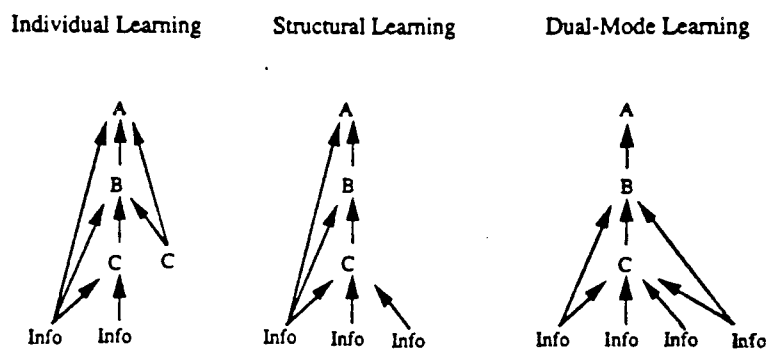


Figure 6: Stylized Organizational Designs for Top Performers With a Minimum of Three Agents

Rather than focus on how this model exhibited standard chaos and complexity results, we focused on those results that are particularly relevant to organizational theory. However, it is worth noting that such standard results are generatable with this model.

CONCLUSION

In most organizational studies, the R^2 for models exploring performance tends to be quite small. This is often attributed to measurement error or missing variables. An alternate explanation suggested by these results is that organizations are complex systems composed of complex adaptive agents. In such systems, linear models cannot capture the complexities, and the use of such models enables only low levels of prediction. Moreover, behavior, even that which is process driven, as in this simulation, often appears chaotic, without a simple relation between input/independent and output/dependent variables. Despite the simplicity of the process, results can vary dramatically given only small changes in parameters or initial conditions. In this study, it was seen that small changes in organizational design can result in dramatically different performance levels.

The foregoing results indicate that there is a relation between organizational design and mode of organizational learning. In this study, organizations had difficulty locating the optimal form, where

optimal is defined in terms of accuracy in performance. However, most organizations developed a form with vaguely hierarchical characteristics. Most of the organizations that emerged had multiple levels, specialization of function by level, and so forth. As many studies have shown, hierarchical designs, though not optimal from an accuracy perspective, do have many advantages from the organization's perspective, such as high resiliency in the face of personnel loss, stress, and information errors. Contingency theory suggests that the reason most organizations have somewhat different forms and that performance is not perfectly related to structure is that there must be a match between structure and task or structure and environment for high performance to result. This analysis suggests another explanation. Organizations vary in form but only slightly because they are searching for the optimal form in essentially a performance plateau in which chance more than strategy will lead them to the particular small change in design that improves their performance.

In this study, many organizations with highly similar designs exhibit almost equal performance. These results are consistent with an interpretation that mimicry may be necessary to achieve better organizational designs. Because there are many designs that are equally good, locating the best design will occur largely by chance. Consequently, mimicry across organizations can increase the likelihood of any one organization finding the optimal structure—assuming, of course, that organizations can accurately assess the performance of other organizations. This result also suggests that organizations can redesign with relative impunity. After all, in the market it is not necessary to be optimal—just as good or better than your competitors. Executives can achieve credit for doing “something,” for restructuring their organization. Moreover, these restructurings, which may not be lead to performance improvements, may be advantageous to the organization in other ways, such as making the organization more efficient by lowering isolation, ignorance, and redundancy.

NOTES

1. Different strategic approaches than that explored herein can be characterized by altering the mean value for the Poisson distributions for each of these four types of moves.

2. Each agent sees a binary string of numbers for each task. The agent simply counts up the number of 1s and 0s he or she sees. If there are more 1s than 0s, the agent reports 1, else 0. If there are the same number of 1s and 0s, the agent randomly chooses between the two options.

3. Initial and final temperatures were chosen after the relationship between probability of acceptance and temperature was analyzed for this organizational model. Given a set of 1,000 randomly generated organizations, each organization was simulated for 100 tasks for an initial partition. The results are shown in the following table. As can be seen, when the partition includes 99 percent of the theoretical range for the Δ cost, all 100 changes are accepted. When the partition was $1e-95$ (virtually 0), half of all moves are accepted. This is because the performance space is a large plain and the chances of a move being downhill are 50/50, even for efficient organizations, because of the "flatness" of the plain and the inherent randomness in organizational performance.

<i>Initial Partition</i>	<i>Temperature</i>	<i>Ratio of Accepted to Total Moves</i>
.99	199	100/100
1e-1	8.69e-1	96/100
1e-2	4.34e-1	89/100
1e-3	2.9 e-1	87/100
1e-4	2.17e-1	90/100
1e-7	1.24e-1	74/100
1e-10	8.69e-2	64/100
1e-14	6.68e-2	72/100
1e-16	5.43e-2	68/100
1e-20	4.34e-2	69/100
1e-25	3.45e-2	55/100
1e-30	2.9 e-2	53/100
1e-35	2.48e-2	55/100
1e-95	9.14e-3	50/100

4. In a second virtual experiment, the size of the agent's memory and the number of tasks between potential changes were varied. The results indicate that for this task, the major effect of structural change occurs when the time window between changes is much less than the size of the agent's memory. In fact, there is a wide range of values in which variation in these parameters is largely irrelevant with respect to performance.

REFERENCES

- Allison, Graham. 1971. *Essence of Decision*. Boston, MA: Little, Brown.
- Axelrod, Robert M. 1976. *Structure of Decision: The Cognitive Maps of Political Elites*. Princeton, NJ: Princeton University Press.
- Axelrod, Robert and Doug Dion. 1988. "The Further Evolution of Cooperation." *Science* 242:1385-90.
- Carley, Kathleen. 1991a. "Designing Organizational Structures to Cope With Communication Breakdowns: A Simulation Model." *Industrial Crisis Quarterly* 5:19-57.
- . 1991b. "A Theory of Group Stability." *American Sociological Review* 56:331-54.

- . 1992. "Organizational Learning and Personnel Turnover." *Organization Science* 3:20-46.
- . 1995a. "Communication Technologies and Their Effect on Cultural Homogeneity, Consensus, and the Diffusion of New Ideas." *Sociological Perspectives* 38:547-57.
- . 1995b. "Computational and Mathematical Organization Theory: Perspective and Directions." *Computational and Mathematical Organization Theory* 1:39-56.
- Carley, Kathleen and Allen Newell. 1994. "The Nature of the Social Agent." *Journal of Mathematical Sociology* 19:221-62.
- Carley, Kathleen, David Park, and Michael Prietula. 1993. "Agent Honesty, Cooperation and Benevolence in an Artificial Organization." In *Workshop Notes for the AI and Theories of Groups and Organizations: Conceptual and Empirical Research*, edited by M. Prietula, K. Carley, L. Gasser, and D. King.
- Carley, Kathleen M. and Michael J. Prietula. 1993. "Plural-Soar: Towards the Development of a Cognitively Motivated Theory of Organizations." Symposium conducted for the National Science Foundation at the Proceedings of the 1993 Coordination Theory and Collaboration Technology Workshop, May, Washington, DC.
- Cohen, Michael D. and James G. March. 1974. *Leadership and Ambiguity: The American College President*. New York: McGraw-Hill.
- Cohen, Michael D., James B. March, and Johann P. Olsen. 1972. "A Garbage Can Model of Organizational Choice." *Administrative Science Quarterly* 17:1-25.
- Crowston, Kevin. 1994. "Evolving Novel Organizational Forms." Pp. 19-38 in *Computational Organization Theory*, edited by K. M. Carley and M. J. Prietula. Hillsdale, NJ: Lawrence Erlbaum.
- . Forthcoming. "An Approach to Evolving Novel Organizational Forms." *Computational and Mathematical Organization Theory*.
- Cyert, Richard M. and James G. March. 1963. *Behavioral Theory of the Firm*. Englewood Cliffs, NJ: Prentice Hall.
- DiMaggio, Paul J. and Walter W. Powell. 1983. "The Iron Cage Revisited: Institutional Isomorphism and Collective Rationality in Organizational Fields." *American Sociological Review* 48:147-60.
- Eccles, Robert G. and Dwight B. Crane. 1988. *Doing Deals: Investment Banks at Work*. Boston, MA: Harvard Business School Press.
- Hannan, Michael T. and John Freeman. 1977. "The Population Ecology of Organizations." *American Journal of Sociology* 82:929-64.
- Hastie, Reid. 1986. "Experimental Evidence on Group Accuracy." Pp. PAGES?? in *Handbook of Organizational Communication: An Interdisciplinary Perspective*, edited by F. M. Jablin, L. L. Putnam, K. H. Roberts, and L. W. Porter. Beverly Hills, CA: Sage.
- Holland, John H. and John Miller. 1991. "Artificial Adaptive Agents in Economic Theory." *American Economic Review, Papers and Proceedings* 81:365-70.
- Hollenbeck, John R., Daniel R. Ilgen, Douglas J. Sego, Jennifer Hedlund, Debra A. Major, and Jean M. Phillips. 1995. "Multi-Level Theory of Team Decision Making: Decision Performance in Teams Incorporating Distributed Expertise." *Journal of Applied Psychology* 80:292-316.
- Hollenbeck, John R., Daniel R. Ilgen, Dale B. Tuttle, and Douglas J. Sego. 1995. "Team Performance on Monitoring Tasks: An Examination of Decision Errors in Contexts Requiring Sustained Attention." *Journal of Applied Psychology* 80:685-96.
- Horgan, J. 1994. "Sex, Death and Sugar." *Scientific American* 271 (5): 20-24.
- Huber, George P. 1991. "Organizational Learning: The Contributing Processes and the Literatures." *Organization Science* 2:88-115.

- Kim, D. H. 1993. "The Link Between Individual Learning and Organizational Learning." *Sloan Management Review* VOL#?? (fall): 37-50.
- Kirkpatrick, S., C. D. Gelatt, and M. P. Vecchi. 1983. "Optimization by Simulated Annealing." *Science* 220:671-80.
- Lant, Theresa L. 1994. "Computer Simulations of Organizations as Experimental Learning Systems: Implications for Organization Theory." Pp. 195-216 in *Computational Organization Theory*, edited by K. M. Carley and M. J. Prietula. Hillsdale, NJ: Lawrence Erlbaum.
- Lawrence, Paul R. and Jay W. Lorsch. 1967. *Organization and Environment: Managing Differentiation and Integration*. Boston, MA: Graduate School of Business Administration, Harvard University.
- Levitt, Barbara and James March. 1988. "Organizational Learning." *Annual Review of Sociology* 14:319-40.
- Macy, Michael W. 1990. "Learning Theory and the Logic of Critical Mass." *American Sociological Review* 55:809-26.
- . 1991. "Learning to Cooperate: Stochastic and Tacit Collusion in Social Exchange." *American Journal of Sociology* 97:808-43.
- March, James G. 1981. "Footnotes to Organizational Change." *Administrative Science Quarterly* 26:563-77.
- March, James G. and Herbert Simon. 1958. *Organizations*. New York: John Wiley.
- Padgett, John F. Forthcoming. "The Emergence of Simple Ecologies of Skill: A Hypercycle Approach to Economic Organization." In *The Economy as a Complex Evolving System*, edited by B. Arthur, S. Durlauf, and D. Lane. Santa Fe, NM: Santa Fe Institute.
- Pete, Andras, Krishna R. Pattipati, and David L. Kleinman. 1993. "Distributed Detection in Teams With Partial Information: A Normative Descriptive Model." *IEEE Transactions on Systems, Man and Cybernetics* 23:1626-48.
- . 1994. "Optimization of Detection Networks With Multiple Event Structures." *IEEE Transactions on Automatic Control* 39:1702-1707.
- Prietula, Michael J. and Kathleen M. Carley. 1994. "Computational Organization Theory: Autonomous Agents and Emergent Behavior." *Journal of Organizational Computing* 41:41-83.
- Rutenbar, Rob A. 1989. "Simulated Annealing Algorithms: An Overview." *IEEE Circuits and Devices Magazine* 5:12-26.
- Simon, Herbert A. 1944. "Decision-Making and Administrative Organization." *Public Administration Review* 4:16-31.
- . 1955. "A Behavioral Model of Rational Choice." *Quarterly Journal of Economics* 69:99-118.
- . 1956. "Rational Choice and the Structure of the Environment." *Psychological Review* 63:129-38.
- Tang, Zhuang-Bo, Krishna R. Pattipati, and David L. Kleinman. 1992. "On the Role of Confidence of Opinions in Distributed Binary Hypothesis Testing Problems." *IEEE Transactions on Automatic Control* 37:632-36.
- . 1993. "Optimization of Detection Networks: Part II—Generalized Tree Structures." *IEEE Transactions on Systems, Man and Cybernetics* 23:211-21.
- Verhagen, Harko and Michael Masuch. 1994. *TASCCS: A Synthesis of Double-AISS and Plural-SOAR*. Pp. 39-54 in *Computational Organization Theory*, edited by K. M. Carley and M. J. Prietula. Hillsdale, NJ: Lawrence Erlbaum.
- Zald, Meyer N. and Patricia Denton. 1963. "From Evangelism to General Service: The Transformation of the YMCA." *Administrative Science Quarterly* 8:214-34.

Kathleen M. Carley is an associate professor of sociology and organizations at Carnegie Mellon University. She examines how learning, structure, and culture influence social and organizational behavior. She has developed methodological techniques for both textual and network analysis. Her current projects include examining the impact of organizational design, strategies, and member skills on organizational performance in a changing environment; the interplay of social structure, cognition, and technology in affecting information diffusion and group formation; mental models analysis; and statistical analysis of social networks and their evolution over time.

David M. Svoboda received his B.S. degree in applied mathematics/computer science from Carnegie Mellon University in 1990. He has been working as a research programmer for the Electrical & Computer Engineering Department of Carnegie Mellon University since 1991. His interests include group social behavior, social networks, and human-computer interaction.

