Organizational Adaptation in Volatile Environments

Kathleen M. Carley

Dept. of Social and Decision Sciences

&

H. John Heinz III School of Public Policy and Management

Carnegie Mellon University

September 1997

This work was supported in part by the Office of Naval Research (ONR), United States Navy Grant No. N00014-97-1-0037 and by the National Science Foundation NSF IRI9633 662. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Office of Naval Research or the U.S. government.

Citation: Kathleen M. Carley, 2000, "Organizational Adaptation in Volatile Environments" Ch. 11 in C. L. Hulin & D. R. Ilgen (Eds.) *Computational Modeling in Organizational Behavior: The Third Scientific Discipline*. Washington, DC: American Psychological Association.

Organizational Adaptation in Volatile Environments

Today, many organizations are continuously engaged in redesigning and reengineering themselves. The purpose of such change is, ostensibly, to locate that organizational design that is better suited to the task or environment facing the organization. Indeed, we can think of organizations as engaged in a process of searching through the space of possible designs for that design which will optimize their performance in this environment. Over time, through a series of incremental changes in their design, organizations have the potential to tune themselves to fit with their environment. However, for many organizations the environment is not stable. The organizational design that works well today may not work well tomorrow. In other words, the search for the optimal organizational design may be complicated by environmental volatility. This paper uses a computational model of organizational learning to theorize about the nature of organizational adaptation in a volatile environment. Computational theorizing makes it possible to move beyond these general claims to more specific claims about organizational adaptation. Analysis suggests that there is no single process that results in an adaptive organization. What organizations emerge as adaptive depends on the environment, the change strategies employed in the organization, and the organization's initial condition. Change per se' is not adaptive. Rather, adaptation is a function of being on the right path, and continually changing the process of change to find that combination of design and experience that yields high performance.

Organizational Adaptation in Volatile Environments

Theories of organizational adaptation address the value of exploration, flexibility, and change. However, exploration, flexibility and change are not without cost. Thus, theories of adaptation also address the need to balance these forces for change against the costs of change and the benefits of exploiting current expertise. Arguments for change assume that organizational adaptation is the result of balance at a single level (strategic or operational) or with respect to a single type of learning (individual or group). In contrast, what I suggest is that this balance is not at a single level but is the result of maneuvering within an ecology of learning in which change is occurring at many levels — individual, organizational, and environmental. Thus, adaptation is not so much a matter of balance as it is of finding the right evolutionary path and of trading change in one dimension for stability in another.

This way of thinking about organizational adaptation as resulting from strategic maneuvering in a learning ecology will be illustrated using a computational model, ORGAHEAD. Computational theorizing is an important tool for developing our understanding of the complex systems we refer to as organizations; it enables the researcher to think through (in a computer assisted fashion) the interactions of dynamic non-linear processes. Using ORGAHEAD the types of maneuvers that result in adaptation and maladaptation for organizations faced with environments that range from stable to highly volatile will be explored. This computational model can be thought of as the embodiment of a theory about how organizations change. Analysis of the results generates a series of theoretical predictions about organizational adaptation.

This computational model, ORGAHEAD, can be thought of as an operationalized grounded theory. The basis for ORGAHEAD is the body of research, both empirical and theoretical, on organizational learning and organizational design. The model has built into it several theories of different aspects of organizational behavior. From the information processing tradition comes a view of organizations as information processors composed of collections of intelligent individuals each of whom is boundedly rational and constrained in actions, access to information by the current organizational design (rules, procedures, authority structure, communication infrastructure, etc.) and by his or her own cognitive capabilities. Organizations are seen as capable of changing their design (Stinchcombe 1965a; DiMaggio & Powell 1983; Romanelli 1991), and as needing to change, if they are to adapt to changes in the environment or the available technology (Finne 1991). Different organizational designs are seen as better suited to some environments or tasks than are others (Lawrence & Lorsch 1967; Hannan & Freeman 1977). Aspects of the model have been tuned to reflect the findings of various empirical studies related to these theories. The set of theories that are unified into a single computational theory of organizational behavior interact in complex fashions to determine the overall level of organizational performance. The aspect of organizational performance of concern here is accuracy; i.e., whether or not the organization given some set of tasks makes the correct decision for each task. A number of theories speak to the impact of learning on organizational performance or accuracy. Herein, the concern is with two such theories — one focused on experiential learning and the other on strategic learning.

The Experiential Learning Perspective: According to this perspective organizations learn as the individuals within them perform tasks and accumulate experience. Individuals gain experience as they face a series of similar or identical tasks over time and receive feedback on how well they are doing (e.g., by being told what the correct answer is). For individuals, experiential learning results in large initial gains in performance, but with decreasing gains over

time as the individuals gain familiarity with the task (the S shaped experience curve). Such experiential learning on the part of personnel can be beneficial to the organization. That is, organizations with more highly trained, more experienced personnel, are expected to perform better; unless, the personnel are over-trained, become bored, or the task changes. Individual experience, however, does not always translate directly into organizational experience (Darr, Argote and Epple, 1995). Rather, the organization's design serves as a procedure for combining or aggregating the experience of individuals (Carley 1992; Cohen, March & Olsen 1972; Hastie 1986). Thus, if we could clone personnel, two organizations with different designs but the same personnel would exhibit different performance. For example, the same group of people will, given the same information, often generate different decisions when they simply vote and when they report their opinions to a CEO or team leader who then makes the decision.

The Strategic Learning Perspective: According to this perspective organizations learn by exploring new ways of doing business; but, organizations also become more staid and do less exploration as they age. In this case, organizational learning results from a process in which strategic planning by the CEO or some central unit about how to position the firm to achieve high future performance leads to changes in procedures, re-assignments of personnel, and re-engineering of the workforce.¹ Over time, institutional norms, sunk costs, traditions, emergence of stable social networks, and so forth gradually increase the cost and decrease the value of, and likelihood of, exploration. Such strategic learning, which can be observed as changes in the organization's design, can be beneficial to the organization. That is, organizations which adopt designs that are forecasted to improve their performance are expected to perform better; unless, the forecasts are wrong. To be sure it is an individual, such as a CEO, or a group of individuals, such as an executive committee, that is engaged in the strategic planning exercise. However, the CEO or executive committees actions are realized as embedded changes in the organization's design (just as the interactions among friends become embedded in underlying social network). The organization's design and the myths about changing it (such as stories about what worked previously) become in effect, a long term social memory, external to any one individual but that can be more or less relied on when contemplating future decisions. In this social memory or distributed cognition sense, change to the organization's design is an act of organizational learning.²

¹ Clearly change to the organization's design is not always strategic. That is, sometimes forces beyond managerial control, such as turnover, result in alternative designs. Herein, the focus is on strategic change as such changes are, for many organizations, more sweeping (as when entire divisions are cut) and possibly more likely. It is important to recognize that such non-strategic design change, like the strategic changes, are also likely to interact with experiential learning, individual performance, and possibly organizational performance (Krackhardt, 1992). Thus, the observations herein about change at individual level and at the structural level are likely to have some relevance even to cases of non-strategic design change.

² Underlying both experiential and strategic learning is the idea that all types of learning result in changes in the connections among different pieces or sets of information. For the individual, these connections occur within the brain as new ideas become interconnected to old. Artificial systems, such as organizations, also have the same capacity to learn by creating connections among information. For organizations, the pieces of information exist as knowledge bases held in a variety of ways such as within the minds of individuals, in file systems, in job descriptions, in data bases, or in artificial agents such as webbots.

The interaction between such experiential learning and strategic learning is complex and non-linear. Strategic changes in the organization's design may mitigate or enhance the value of experience. Thus, it is difficult to predict the joint impact of both types of learning on organizational performance. In the computational model, ORGAHEAD, both experiential and strategic learning are simultaneously active. We can use this computational model to theorize about how the interactions among these two types of learning influence overall organizational performance (measured as accuracy). ORGAHEAD is a procedural theory of organizational performance. As such, the relationship between performance and learning is determined by a series of processes such as how decisions are made, how the organization's structure is changed, and how individuals learn. These processes will be described later. For each organization, at any point in time its performance is uniquely related to the task, the organization's design, the experience of all personnel, the impact of previous designs on the experience garnered by personnel, and the processes by which personnel learn and the design changes. Since we can characterize various aspects of the organization's design and various aspects of individuals as a set of factors we can at any point in time explore the relationship between these factors and the organization's performance. In fact we can think of the organization as simultaneously operating in two spaces — design and experience. An important point to keep in mind is that there is no small set of factors that uniquely captures all aspects of design or experience (this is true both for the model and the real world). Thus a high number of dimensions may be needed to capture all salient aspects of the design space and the experience space. Another important point to keep in mind is that, in a process model, these factors are not simple inputs, rather they may be shifting their values dynamically over time and all factors are inter-related in often complex ways.

We can think of the relation of performance to these various factors as the performance surface. This is graphically illustrated in Figure 1 — where the expected performance of an organization given two elements of design, size and workload, is shown. The resulting surface can be thought of as a landscape that all organizations are, over time, trying to search to find the optimal form. In the experience space, search is carried out, in part, by personnel doing tasks and gaining experience. In the design space, search is carried out, in part, by the CEO (or some executive committee) sequentially altering the organization's design. Of course, for real organizations and for ORGAHEAD simulated organizations, design (and, for that matter, experience) is a function of more than two variables and the performance surface may be even "messier" than that illustrated.

There are several important points that must be made about the performance surface. First, the processes which result in organizations exhibiting a certain level of performance at a certain point in time also result in a systematic relationship between performance and a large number of design and individual factors such as size, density, span of control, workload, average experience, and retention. This relationship, though systematic, is rarely linear in any dimension. For example, in Figure 1 an organization that tries to climb this surface by either just increasing size or just increasing workload will often suffer drops rather than increases in performance, even though to be a top performer the organization will need to be high in both size and workload. Second, the large number of factors means that it is not possible to graphically display the performance surface relative to all factors (for either or both of design and experience) at once for either real organizations or those simulated in this paper; i.e., at most two factors can be shown at a time. Third, the performance surface is the space of possibilities; i.e., it shows the expected relationship between performance and design (or experience) for all possible designs (or experiences). In a world where there are only a few organizations (real or simulated) there may not be an organization at each point in the space. Over time, organizations may come to cluster in certain regions of this landscape. For example, over time, one would

expect adaptive organizations to migrate toward the highest peaks; whereas, maladaptive organizations might migrate toward lower peaks on which they become trapped. Fourth, the over time trajectory of any one organization is unlikely to be linear. For example, an organization might start as small with a low workload and over time first increase in size, then decrease, then increase in workload then increase in size, increase in size, increase in workload, decrease in size, decrease in size, increase in workload, etc. One of the implications of this is that although maladaptive organizations may be lower/higher on some dimension than adaptive organizations, if the maladaptive organizations change in the direction of the adaptive organizations they may not see a gain in improvement.

Figure 1. Illustrative Performance Landscape

For both real organizations and those simulated in this paper the performance surface is not necessarily stable. The environments faced by organizations can change for a number of reasons; e.g., technology, legislation, change in cultural norms, and change in availability of resources. Changes in the environment are often realized as changes in either the task, the way in which organization's can redesign themselves, or changes in the way personnel learn and make decisions. Consequently, such environmental shifts can fundamentally alter the performance surface, particularly as perceived by the organization. Consequently, the relation between performance and design may change. This is graphically illustrated in Figure 2 where the change over time in the performance surface as a function of just two of the possible dimensions of organizational design is shown. The more volatile the environment, the more often one would expect to see shifts in these performance landscapes.

Figure 2. Illustration of Changing Performance Landscapes

At issue is “when is change good?”. To answer this it is important to distinguish change and adaptation. Change is defined as any difference between an organization at one time and the next. These changes may or may not represent learning. Adaptation is defined as those changes that enable the organization to maintain or improve its performance. Since different designs are better in different circumstances, adaptation can be thought of as the successful result of a process of searching for the optimal, or at least a better, design for the environment confronting the organization given that the organization itself is altering its behavior as the individuals within the organization garner experience. Adaptive organizations then are those that end up situated on the higher peaks.

Change, or exploration, by the organization is seen as a force enabling, but not guaranteeing, adaptation. That is, change in design enables the organization to walk around a performance surface and move toward a peak. There are at least two potential drawbacks. If the performance surface is volatile then the organization runs the risk in moving to a new point in design space that while a “peak” at time 1 is actually a “valley” at time 2. Further, for most organizations, performance improvements are often garnered by staying at the same point in the design space but exploiting current talents, knowledge, and skills of the personnel. Indeed, March (1996) discusses continuity as being critical for true adaptation. Levinthal and March (1993) suggest that organizational learning, because it involves balancing exploration and exploitation, can result in various types of learning myopias. For example, experiential learning can result in a tendency to exploit knowledge (keep doing what one does well) and so overlook or discount the future, discount what others are doing, and discount contradictory information or failures. They conclude by arguing that despite these myopias, learning is still advantageous to the firm.

The question is why? In this paper a partial answer is provided: since organizational learning occurs at many levels, myopia at one level may be countered by, and possibly even encourage, a broader view at another level. Moreover, without learning, the chance of finding a better design is very low. The chance of finding a better design is non-zero for several reasons. Essentially, the organization can benefit from fortuitous accidents. At the strategic level, selective retention or random actions on the part of the CEO or the executive committee can result in the organization moving to a better design. At the operational level, selective retention or random actions on the part of personnel can result in the organization with a particular design appearing to be better or worse than it is when all personnel are acting perfectly. This would engender a misperception of the relationship between that design and the organization's performance which may then discourage or encourage the CEO's or executive committee's search for a new design.

Learning is advantageous to the firm precisely because the organization is engaged simultaneously in multiple types of learning. As previously noted, the concern in this paper is with two types of learning — experiential and strategic. Experiential learning, which resides largely in the individual, is seen as resulting from the process of trial and error and learning from feedback. Strategic learning, which resides largely in the linkages or connections among people, is seen as resulting from organizations going through processes of planning, re-organization, and re-engineering. Such change is strategic as such re-organization is generally done because it is expected (for example by the CEO) to improve performance (Kilman & Covin 1988; Butler 1993). The value of these two types of learning may depend on the environment. For example, in a stable environment experiential learning should be quite effective as the lessons of experience should remain valid. In a volatile environment hard won experience may be less valuable as the jobs of tomorrow may not be the jobs of today. In a stable environment there may be little to strategize about and the organization may quickly lock onto the right structure; whereas, in a volatile environment being able to anticipate the future may be critical. But is this the case?

We can begin to address questions such as these, questions about the dynamics of organizational life, using computational theorizing. In this paper, such questions will be addressed using ORGAHEAD.

ORGANIZATIONAL ADAPTATION MODEL

ORGAHEAD is a computational theory of strategic organizational behavior as search through the problem space of potential organizational designs. ORGAHEAD was designed to reflect basic realities of organizational life. That is, in any organization there is a task or set of tasks being done, new tasks arrive over time, a set of personnel, each of whom occupies a particular role in the organization, reports to others, does the tasks and gains experience, and a strategic or management function that tries to anticipate the future, assigns personnel to tasks and determines who reports to whom. These fundamental features of organizations are captured in Figure 3 which serves as a top-level view of the ORGAHEAD model. Thus, from a learning perspective, within ORGAHEAD organizations are characterized at two levels — operational and strategic. At the operational level, the organization is characterized as a collection of intelligent adaptive ACTS agents (Carley & Prietula 1994); i.e., each agent is cognitively capable, socially situated and engaged in working on a specific task. At the strategic level, the organization is characterized as a purposive actor; i.e., there is a CEO or executive committee which tries to forecast the future and decides how to change the organization to meet that need. Each of the three primary components within ORGAHEAD will be described in turn — the task model, the operational model, and the strategic model.

Figure 3. Top Level View of ORGAHEAD

Task and Task Environment

In ORGAHEAD, the organization is faced with a sequence of tasks. Each task has associated with it a “correct decision”. This correct decision is not known, a priori, by any one within the organization. For each task, the CEO must make the organization’s decision. This decision may or may not match the correct decision. Organizational performance is a function of how often this decision matches the correct decision; i.e., of the organization’s accuracy.

There are many tasks in which organizations can engage. A common one is the classification task. Classification tasks have the following properties: There is an item that needs to be classified as being of a certain type. The item can be characterized as having a set of attributes, and based on the values of these attributes, the item is classified. Examples of classification tasks include: diagnosis, situation assessment, and multiple choice tests. A simple version of the classification task is the binary choice task. In the binary choice task the objective is to determine for a given binary string whether that string is of type A or B (a string cannot be both A and B). For each string there is a correct decision; i.e., each string is either of type A or B. Of course, the decision maker may not correctly classify the string; i.e., the decision maker may misclassify an A string as a B string or vice-versa. The specific task used in the ensuing analysis is a nine bit binary classification task. Classification tasks in general, and the binary choice task

in particular, have been extensively studied by researchers interested in team and organizational performance (see for example, Tang, Pattipati, & Kleinman 1992; Pete, Pattipati, & Kleinman 1995; Carley 1992; Lin & Carley forthcoming; Hollenbeck, Ilgen, Sego, Hedlund, Major, & Phillips 1995; Hollenbeck, Ilgen, Tuttle, & Sego 1995).

Environments can be thought of as task environments; i.e., the set of tasks faced by an organization. For example, when the task is the binary choice task then the environment might be the set of binary strings of length N that must be classified by the organization. Environments have various characteristics such as bias and stability/volatility. Environments may be biased (another interpretation of this is that they may be a niche environment). One way of thinking of bias is as the prevalence of a specific. For example, in the binary choice task, the environment is unbiased if all outcomes (whether the string is of type A or B) are equally likely given a set of strings. The more likely one outcome is than the other, given a set of strings, the higher the bias of that environment. For the binary choice task, it is unbiased if a string is of type A if there are more 1's than 0's in the set of task features and B otherwise, all strings of length N are possible, and the probability that any task features is a 1 is .5. Under these conditions, neither outcome (A or B) is more likely to occur.

Environments can also vary in their stability. Hannan and Freeman (1977, p. 952) characterize environmental instability as an oscillating sequence of “patches”. Seasonal markets — such as lawn care — have this oscillatory nature. Environmental instability is a function of both the frequency of shift (Hannan and Freeman’s “grain size”) and the degree of the shift. We can operationalize this as oscillations between one class of tasks and another; e.g., between an unbiased and a biased task.³ The degree of shift is operationalized as the difference in the bias between the two classes of tasks, the greater the difference in bias, the greater the degree of the shift. Frequency of shift is the number of tasks of a specific class the organization faces before it shifts to a different class of tasks. In a stable environment, all tasks are chosen from the same distribution. In an unstable environment the tasks are chosen from different distributions, both unbiased and biased, in an oscillatory fashion.

The specific task being done by the organization can affect the optimal organizational design and may constrain the level of performance achievable (Carley 1992). The task largely determines the specific optimal structure. Thus, unless one is interested in finding the organizational design for a specific task, the issue is not what design is found, but what factors affect differences in what designs are found or the rate at which the optimal design is found. In the ensuing analysis, the focus is not on what structures are found, but in what causes different structures to be found and whether different structures are found in stable and volatile environments.

Operational Level

In ORGAHEAD, the organization does this task in a distributed fashion. That is, at the operational level, each individual in the organization works on at most part of the task (no individual sees the whole string). The operational part of ORGAHEAD is built out of the CORP model (Carley 1992; Carley & Lin 1995, forthcoming). CORP is an information processing model of organizational performance in which organizational performance results from aggregating the performance of individuals. All individuals are boundedly rational in terms of organizational access to information and cognitive ability to process information (Simon 1955, 1956; Carley & Newell 1994). Individuals are limited cognitively; they cannot do the task by

³ For an alternate operationalization using formal logic see Bruggeman (forthcoming).

themselves (can only handle a maximum of seven pieces of information), forget, have primacy and recency biases (they remember only the first 500 tasks and the most recent 500 tasks they have seen), remember general trends rather than particular tasks, and are effectively overconfident in their decisions. Individuals are also limited structurally; i.e., they cannot access all information, what information they can access depends on their role in the organization, and they often must act not on the basis of actual task information but on the interpretation of that information as forwarded by others in the organization. The individual's role is a function of the authority structure (who reports to whom), S , and the resource access structure (who has access to what resources or information), R . Individuals have the capability of learning from their experience over time. Specifically, in each time period in which they receive information, they report whether they think (on the basis of that information) the overall task is of type A or B, and then they find out whether or not the task really was of type A or B. Individuals then learn by keeping track of the likelihood that when they see pattern x the task will be of type A/B. Each individual, regardless of his or her position in the organization, acts as an experiential learner, acquires information, classifies the pattern of information that he or she sees, recommends the choice that was most often correct in the past for that pattern to his or her superior, receives feedback as to whether the task was really of type A or B, and increments his or her memory. If the individual has no previous experience he or she simply guesses. The pattern of information seen by an individual can be either, or both, raw information on this specific task or the recommendations of other individuals for this task.

The process of pattern classification, recommendation, and reporting up through the authority structure is the organizational process by which individual level decisions are aggregated into a single organizational decision — the decision made by the CEO. See Figure 3. At the apex of the organization is a single superior or CEO. Below the CEO there is a network that connects the individuals into a single organizational structure. We can think of this as a directed graph, S , showing who reports to whom. Overlaying this authority structure is a second network that connects individuals to raw task information. We can think of this as a directed graph, R , showing which information on the task is observed or monitored or handled by which individual. This second graph is the resource access structure. Each piece of information is a resource for the organization. The resource access structure determines which part of the task is observed by which individual. Each time period there is a single task for which the organization collectively makes a classification which is the decision made by the CEO. Each time period, the flow of information (and opinions) is from the task to the CEO as filtered by these two networks — the authority structure, S , and the resource access structure, R . The organization's design can be formally characterized by the joint structure — design (S,R) .

Organizational performance is measured as accuracy, the percentage of problems in a window of opportunity that the organization correctly classifies. The organization is faced with a sequence of tasks drawn with replacement at random from the population of tasks. Each time period the organization, as represented by the CEO, must make a decision for the current task. Of the $N(t)$ individuals who are collectively processing $I(t)$ pieces of information some individuals are processing raw information on the task and reporting their opinions to other individuals; whereas, other individuals are processing these opinions, and still others may be processing both raw task information and opinions. Who is doing what, and who reports to whom, depends on the organizational design. An organization may be composed of anywhere between 2 and 45 individuals organized in one to three tiers below the CEO with a maximum of 15 individuals per tier. There is an additional or first tier composed of a single CEO. As noted, the organizational decision is the CEO's decision. By convention the CEO is referred to as Level 1 and the tier furthest removed from the CEO as Level 4.

CORP and related models have been extensively described in other venues and their implications for organizational performance characterized (Carley 1992; Carley & Lin 1995, forthcoming; Ouksel & Mihavics 1995; Mihavics & Ouksel 1996). CORP has been shown to be a reasonable model of organizational performance for organizations that are not collegial voting or democratic structures when compared with both experimental lab studies (Carley 1996) and archival data on actual organizations (Lin 1994; Carley & Lin 1995).

Strategic Level

At the strategic level, organizational performance is affected by the ability of the CEO or central unit to anticipate the future and take the appropriate strategic actions to alter the organization's design in response to environmental cues. Change at this level is determined by the organization's change strategy.⁴ The change strategy⁴ has three parts: what changes to the current design are allowable, conditions under which changes are made, and the likelihood of the various changes.

Recall that the organization's design is characterized in terms of the authority structure and the resource access structure. Change in design involves changing these structures. There are four types of changes possible regardless of the change strategy: *downsizing* — drop n individuals (such that $1 \leq n \leq N_O(t)$), *upsizing* — add n individuals (such that $1 \leq n \leq N_{max} - N_O(t)$), *redesigning* — delete the tie between individual i and j (i reports to j) and reassign individual i to report to individual k , and *re-engineering* — delete the tie between individual i and piece of information s and add a tie between individual j and piece of information s . Exactly how many changes of a particular types occur at a time is given by a Poisson distribution.⁵ Changing connections through re-design and re-engineering moves connections and will not lead to an absolute increase or decrease in personnel. Changes in personnel, through upsizing and downsizing, cause changes in all of the ties associated with those personnel.

The organization begins with a particular design ($S(0), R(0)$) and proceeds to process 500 tasks. After this, regardless of the change strategy, the organization continuously cycle through the following actions: make decisions for a sequence of tasks, evaluate performance, propose a change to the structure, change the structure. The four strategies vary in terms of when they move from the decision phase to the structural change phase and in the way in which they decide on a possible new design. The four strategies are stable (no change), random, procedural, and strategic.

⁴ A strategy is a set of moves and the statistical distribution for each move that determines the probability with which that move will be chosen. Different strategies can be defined either by altering what moves are possible, or by altering the underlying distributions. Conceptually, however, some reserachers might be more comfortable thinking of changes in the underlying distribution, particularly small changes, as perturbations of a single strategy. For example, imagine the two following strategies. In strategy 1, the only moves possible are hires and fires and the probability of each is determined with a Poisson distribution with a mean of .2. In strategy 2, the only moves possible are hires and fires and the probability of each is determined with a Poisson distribution with a mean of .3. For some researchers, this second strategy will represent a new strategy, for others it will represent only a perturbation on an existing strategy.

⁵ ORGAHEAD is a discreet event simulator. Thus, the number of personnel or connections that can be added or dropped each period is a positive integer. Poisson is a standard distribution used when x (in this case the number of changes) can take on only positive integer values.

Organizations with a stable change strategy do not alter their design over time. This can be thought of as a baseline case against which to evaluate the extent to which change is adaptive. Organizations with a random change strategy, randomly alter their design every time the organization's performance changes by 5% or more. The level at which change occurs in the organizational structure and the type of change is random. After a change occurs the organization was required to remain 'dormant' long enough for new people (and old) to get used to the new organizational design. Change was equally likely at all levels (25%) and there was a 25% probability that no change would occur. Organizations with a procedural change strategy hired personnel when things were going well (5% or more improvement in performance) and fired personnel when things were going poorly (5 % or more drop in performance). Under this strategy, change was more likely at lower levels in the organization. Specifically, the probability of change at level 4 is 50%, 30% at level 3, and 5% at level 2. Again, there was a 25% probability that no hiring or firing would occur when the efficiency warrants it. Under the 'strategic' strategy, the organization used strategic planning to determine whether or not to make a change and whether a change was made depended on expectations. Over time, the organization learned what types of changes were effective and became increasingly adverse to making changes that were not expected to improve performance. This was implemented via simulated annealing.

Simulated annealing is a reasonable computational analog of organizational strategic adaptation (Carley & Svoboda, forthcoming). Simulated annealing is a heuristic approach to optimization (Kirkpatrick, Gelatt, & Vecchi 1983; or for an overview see Rutenbar 1989). In an annealer, the system tries to optimize some function by moving through states sequentially, and the path through those states is constrained by the set of ways of altering the current state. Further, in an annealer, a move is chosen from the set of possibilities and then its potential impact is evaluated before it is accepted (or not) according to the Metropolis criteria. According to the Metropolis criteria non-risky moves are always accepted and risky moves are sometimes accepted but with a decreasing likelihood over time. This rate of decrease is controlled in part by the temperature of the system. The temperature of the system determines its excitation; thus, as the temperature drops, the system becomes more rigid and less capable of change. The result is a heuristic based optimization process that tends to locate better states but is not guaranteed to locate the best state. Similarly, the organization moves through a series of organizational designs, one at a time. Over time, the organization attempts to optimize its design given some cost function (such as minimizing salary, maximizing the number of widgets produced, or maximizing decision accuracy). The CEO or central unit has a set of possible strategies (move set) that dictates which designs are possible given the current design. The CEO does not compare all strategies, but simply evaluates a strategy by trying to anticipate the future (Allison 1971; Cohen & March 1974; Axelrod 1976). Strategic adaptation requires the CEO to have knowledge about which individuals in the organization have what information and which capabilities. Consequently, the organization's optimization process is imperfect. Strategic change moves the organization closer to the goal but may not achieve it (Simon 1944; March & Simon 1958). Organization's gamble on redesigns that might possibly "increase costs", and they are much more prone to this kind of risky behavior when they are new (Stinchcombe 1965b). With maturation the organization becomes staid and trapped by its own competency (Levitt & March 1988).

The goal of the strategic change process as implemented herein is to find the organizational design that maximizes performance. In other words, the goal is to alter the organization's design sequentially so as to locate the tallest peak shown in Figure 1. For the strategic approach, first the organization's performance for 100 tasks is calculated. Then the CEO chooses one change

from the set of possible changes, thereby suggesting a new organizational design. This design is evaluated (using a limited lookahead of 100 tasks). Then the forecasted performance of the proposed design is compared with the previous performance of the current organization, and a strategic decision is made as to whether or not to accept the change. Finally, if the change is accepted the organization's design is altered and the process begins again. If the change is not accepted the process begins again with the unchanged organization. Performance at time, t , for the current organization is the percentage of the most recent 500 tasks that the organization correctly classified prior to time t .

The probability of accepting the new design is determined via the Metropolis criteria. Specifically, the change is always accepted if the forecasted performance for the hypothetical organization is better than the known performance of the current organization. Further, when the forecast is poorer the change may still be accepted. In fact, we can think of the probability of accepting the "bad" design as resulting from the impact of the organization's risk aversiveness on its decision. This probability is calculated, using the Boltzman equation, as $P = P_O e^{(-2 \text{ cost}(t)/T)}$, such that $\text{cost}(t) = 0 - \text{performance}(t)$ and P_O is the initial probability of accepting a "bad" design. This probability decreases as the "temperature" decreases. We can think of temperature as the organizations current level of risk aversion. Herein, temperature drops every 100 tasks (time periods) as $T(t+1) = \alpha * T(t)$ where α can be thought of as the rate at which the organization becomes risk averse.

Summary

ORGAHEAD predicts organizational performance (measured as accuracy in doing a task) from the task, operations needed to perform the tasks (individual actions and experience), and strategic changes to alter the organization to do the task better (refer back to Figure 3). All aspects of ORGAHEAD are dynamic. That is, each time period the specific task faced by individuals (and the organization) is different. Over time, the environment of tasks may be changing. At the operational level, individuals learn through experience as they do the task, but what they do and what they learn is a function of their position in the organization's design. At the strategic level, attempts at anticipating the future often result in changes in this design.

Virtual Experiment

Using ORGAHEAD a series of virtual experiments were run to examine the impact of change strategies on performance and adaptability in a changing environment. In each experiment two or more initial conditions were considered and then for each condition, using Monte Carlo techniques, the over time behavior of a large number of organizations were simulated.

As with human experiments it is not possible to explore all possible variations possible given the proposed model. Thus, hypotheses in the literature are used to give guidance as to what factors are important to look at and when we are most likely to observe differences in organizational performance. The particular issue of interest is strategic adaptation to a dynamic environment. Thus, the experiment is design to look at variations in environment and variations in strategy. This concern follows in part from the work of Hannan and Freeman (1977) who suggest that environmental instability may require a different type of organizational design than environmental stability. The more unstable the environment the more likely we are to see performance differences. In an unstable or oscillatory environment the strain on the organization's design should be higher as the difference between the two types of tasks increase; i.e., the greater the difference in bias between the two classes of tasks, the greater the strain on the organizational design. This line of reasoning suggests that an examination of organizational

adaptation in a changing environment can be made by examining the impact of switching between unbiased and highly biased tasks. One might further expect that the more rapidly the organization switches between the two classes of tasks the greater the strain on the design. This suggests that an examination be made of non-oscillating, slowly oscillating, and quickly oscillating environments.

In the first experiment two environments were examined: a stable environment under an unbiased task condition and a volatile environment oscillating between an unbiased and a highly biased task every 1000 tasks. For the highly biased task, the outcome was an A if there were 3 or more 1's in the string and a B otherwise. In the second virtual experiment three environments were examined: stable (no shifts), low volatility (shifts every 5000 tasks), and high volatility (shifts every 1000 tasks) for organizations that changed strategically. It is reasonable to expect that the strategy the CEO follows in deciding which change to make to the organization's design should impact the organization's ability to adapt to the different environments. Thus, in both experiments, four change strategies are examined: stable (no change), random, procedural, and strategic.

In both experiments, for each initial condition, 1000 organizations were generated randomly in a Monte Carlo fashion. The initial size of the organization (2 to 45 individuals), the number of levels below the CEO (1 to 3), the initial reporting structure ($S(0)$), and the initial resource access structure ($R(0)$), were all chosen randomly with replacement from the set of possibilities. Each organization was simulated for 2000 time periods (after the initial 500).

The goal of the analysis is to gain insight into strategic adaptation to a dynamic environment. One way of examining this is to see what characterizes the organizations that are adaptive (end up on the peaks in the long run) or maladaptive (end up in the valleys in the long run) in the various environments. As previously noted, the performance surface is multi-dimensional even in the design space. To contain the analysis, only three aspects of design will be examined: size, density, and workload. Size, the number of individuals in the organization, is one of the most widely studied of all structural design variables in the empirical literature. Size is important for a number of reasons; e.g., size is expected to impact performance because as the size of the organization is increased more aspects of the task can be analyzed but the coordination problems increase. Density is the fraction of possible connections in the reporting structure that actually exist. Although density per se is less well studied than size, factors that are highly related to density such as the communication, worker autonomy, and decentralization have all been shown to contribute to performance. Essentially, density is expected to impact performance because the higher the density the higher the management workload but the greater the communication and so potential for noticing errors. Workload is the average number of items of information that each employee needs to handle to make a decision. Workload is expected to affect performance because the higher the workload the greater the information processing requirements on the individuals and thus the slower they will learn. Since the concern is with adaptation, only the performance of the organizations in the long run is considered. Since organizations begin with individuals who have no experience initial organizational performance is approximately 50%; i.e., initially, each CEO only correctly classifies 50% of the problems. The adaptiveness of the organization can thus be seen by measuring its performance at a later time period, such as the last 500 tasks.⁶ Since strategic

⁶ An interesting question is why not examine the trajectories of organizations over time. There are two basic answers. First, empirical complexity. That is, to examine the relation between just these 3 design factors and performance for each point in time (2000 time points) for

learning is relatively slow compared to experiential learning, looking at organizational performance during later (rather than earlier) time periods provides a more complete understanding of the impacts of strategic change on the organization.

DETERMINANTS OF ADAPTIVITY

A series of questions about adaptability were raised. For example: Under what conditions is change beneficial? Why is learning advantageous? Does strategic learning provide greater advantage in a volatile environment and experiential in a stable environment? The ensuing analysis examines these and other questions about organizational adaptation. To make clear the determinants of adaptation, the focus will be on a comparison of the relatively adaptive and maladaptive organizations. Adaptive organizations are defined as those 50 organizations (out of the 1000 simulated) that exhibit the highest level of performance at the end of the 2000 trials. Maladaptive organizations are defined as those 50 organizations (out of the 1000 simulated) that exhibit the lowest performance at the end of the 2000 trials. By examining differences and similarities in the those organizations which, in comparison to each other, are relatively more adaptive or maladaptive, a series of lessons about organizational adaptation can be garnered.

Lesson 1: Change is Not Adaptive

In Table 1, the average amount of change for several different types of change for adaptive and maladaptive organizations is shown. As can be seen, change is not adaptive. In a stable environment, change is typically maladaptive. In general, maladaptive organizations end up changing more than adaptive organizations (in particular they downsize more). Downsizing is maladaptive if it gets rid of the organizational knowledge residing in the minds of the personnel — both knowledge about the task and knowledge about people. Even when the environment is changing, some of that knowledge is useful and transferable. Further, in an oscillatory environment, the need for that type of knowledge returns periodically. In contrast, the value of upsizing depends on the environment. Consequently, adaptive firms downsize less. Finally, in a stable environment organizations employing the strategic approach to change tend to change more; however, their final performance is comparable to those which change either randomly or procedurally.

Table 1. Average number of changes made by adaptive and maladaptive organizations

Adaptive				Maladaptive			
Redesign	Re-engineer	Upsize	Downsize	Redesign	Re-engineer	Upsize	Downsize

each of the 1000 organizations for each of the 20 different initial conditions would require analyzing 160 million pieces of data. The second reason has to do with the length of time it takes to run a simulation. Essentially, IO (reading and writing to files) is a major component of the time it takes to run these simulations. Thus, to speed the programs to the point where data could be extracted with less than 3 weeks of run time it was necessary to restrict the amount of data that was gathered and when it was gathered. The option taken was to gather more types of data but at fewer points in time as it was felt that this afforded a more comprehensive view of system performance. For a small number of organizations, complete information was gathered to see if this approach was reasonable. The third reason has to do with the analysis of trajectories. That is, it is very difficult to analyze n dimensional trajectories. Future work involves exploring the use of Fourier analysis for analyzing these organizational trajectories.

Stable Environment								
Random	2.38	0.94	1.22	0.52	5.98	2.72	1.62	1.48
Procedural	1.70	1.00	1.02	0.40	6.02	3.24	1.38	1.68
Strategic	54.98	29.92	25.28	14.32	63.30	31.80	15.98	23.44
Volatile Environment								
Random	31.38	12.79	12.48	6.48	30.32	15.40	10.58	7.44
Procedural	30.04	16.24	12.70	6.84	29.94	18.96	10.66	7.18
Strategic	56.78	31.90	11.12	7.60	64.38	33.20	10.10	14.38

In general, in high volatility environments, more change is needed than in stable environments to achieve a reasonable level of performance. Interestingly, on average, the adaptive organizations in a volatile environment tend to perform as well or better than the adaptive organization in a stable environment. Essentially, in a stable environment, despite being able to change, organizations typically get locked into competency traps not in what they do, but in how they do it (i.e., in structure)

In a previous study, Carley (forthcoming) found that adaptive organizations changed more than did maladaptive organizations. This result would seem to contradict those reported above. After aligning the models⁷ it was found that those results were the product of having averaged the effect for different change strategies (only personnel change, only connections change, and both) for the cumulative effect of all four types of change. In particular, if the change strategy is personnel change (upsizing) then adaptive organizations that act strategically do tend to change more. In contrast, in this study, both personnel and connection changes are allowed and the number of each of the four types of changes is shown separately.

Lesson 2: History Matters

As was just seen, each of the change strategies leads to a different level of change but results in structures that are comparable in terms of efficiency. Do the same organizational structures emerge regardless of which strategy is used? If so, then the strategy that causes less change could be considered the more efficient one. The point here is not to find the optimal structure for this task. The optimal structure is idiosyncratic to the task and, as such, is not particularly interesting from the standpoint of general principles. The point is to see whether the same structure emerges regardless of which change strategy is employed. In Figure 4, the final size, density, and workload for the organizations by change strategy and environment is shown. As can be seen in Figure 4, the design of the adaptive organizations is different by environment and by change strategy. For example, in stable environments, the adaptive organizations are somewhat smaller and less dense than are the adaptive organizations in a volatile environment. There is an interesting interaction between change strategy and environment. Specifically, adaptive organizations that have changed randomly tend to be smaller, more dense, with lower workload than their strategic counterparts in the stable environment, in the volatile environment the random changes are smaller, more dense, but have a higher workload than their strategic counterparts. As you move from a stable to a volatile environment, the organizational designs

⁷ This was done using a process called docking, an analysis technique for the comparison of computational models, described by Axtell, Axelrod, Epstein & Cohen (1996).

that emerge as being adaptive when a strategic change strategy is used tend to be similar. This is not the case for either the random or the procedural change strategy.

Stable Environment

Volatile Environment

Figure 4. Impact of change strategy and environment on design of adaptive organizations

In the foregoing analysis we examined only the adaptive organizations and saw that the change strategy affected the shape of the adaptive organizations. But does it also affect the shape of the maladaptive organizations? In Figure 5 we see that the change strategy leads to different organizations emerging in both the maladaptive and the adaptive arena. In particular, when the change strategy is random or procedural the organizations that emerge as maladaptive tend to be larger, less dense, and possess lower workloads than those emerging as adaptive. Whereas, under a strategy of strategic change maladaptive organizations are smaller, more dense and with lower workloads than their adaptive counterparts.

Figure 5. Impact of change strategy on relative design of maladaptive and adaptive organizations in a volatile environment.

Regardless of the change strategy, two organizations with the same design given the same set of tasks will exhibit the same performance if the employees within the organization have the same knowledge. The results displayed in figures 4 and 5 however, suggest that two organizations with the same design and different change strategies are actually making different decisions and thereby achieving different levels of performance. How can this be? There are two reasons. First, the structures that emerge under the different change strategies even when they have the same size, density and workload are somewhat different in configuration. Size, density and workload do not capture all of the nuances of the design.

In other words, design matters — the exact configuration of the organization affects its performance. Second, as was noted, two organizations with identical configurations will perform the same only if the employees are also identical. But, the change strategy affects what the employees know. The change strategy determines what employees learn what when and the context of that learning. Thus two employees who 'grow up' in organizations with different change strategies, even when those organizations are faced with the same set of tasks, will learn different things. Collectively this results in the organizations with different change strategies learning different things, and with two organizations with the same design and different change strategies exhibiting different performance. In other words, history matters — the path by which an organization reaches a particular design determines the effectiveness of that design.

Lesson 3: Avoid Locking in to a Change Strategy: Overshooting, Undershooting, and the Downward Spiral

Consider those organizations that in the end are either adaptive or maladaptive. Organizational change results in increased variation in organizational design (see Figure 6). To be sure, on average there is some difference initially. However, and importantly, this difference grows over time. Maladaptive organizations using a procedural or random strategy tend to overshoot in size, and undershoot in density and workload. Whereas, maladaptive organizations acting strategically tend to undershoot in size and workload and not to decrease far enough in density. The same change strategy can move the adaptive organizations in one direction and the maladaptive organizations in another.

Figure 6. Over time change in design

A detailed analysis of a few organizations that ended up as adaptive or maladaptive suggests that organizations get locked into certain ways of changing. For example, one highly adaptive strategic organization began by initial high levels of hiring and then over time, when it changed, only changed connections (who was doing what and who reported to whom). As another example, a highly maladaptive strategic organization began small, downsized even more, and then got locked into a cyclical pattern of hiring and firing. Thus, it appears that organizations can get locked into a spirals of behavior where over time they come be locked not just into poor/good performance, but into organizational designs and approaches to changing that design that further serve to degrade/improve performance more. Over time this pattern of behavior causes initial difference between the organizations to grow. In this sense, being adaptive is locking into the right pattern of change given the chosen change strategy. Maladaptive organizations appear to be ones that lock into a way of changing that causes them to overshoot or undershoot their objectives, or to re-enforce their current negative tendencies. Future work on the over time trajectories and patterns of change is needed.

COMPUTATIONAL THEORIZING

The dynamics of change result from simple, but non-linear processes. Certainly, the interactions among the different types of learning — individual experiential learning and structural expectation based learning — are non-linear. Thinking through the implications of change processes and trying to predict adaptation is thus non-trivial. Consider the following learning interaction. When the environment is stable, the lessons of experience are valuable, and performance can be improved simply by relying on that experience. When the environment is volatile, the lessons of experience may be less valuable but certainly not irrelevant. In this case, re-engineering the organization to meet the demands of the changing environment may improve performance as the design developed may more closely match that needed for the task, or it may degrade performance further by making the lessons of experience held by old personnel even less relevant.

Computational models are valuable tools for theory building and for examining issues of social and organizational dynamics because they enable the researcher to generate a set of consistent theoretical propositions from basic principles even when there are complex interactions among the relevant factors. Models, can be thought of as abstractions of reality that strip away many of the complications existent in actual organizations and so lay bare the relationships among the remaining components of organizational design and adaptation. In the organizational domain, however, most organizational action is the result of interactions among adaptive systems, of emergent structuration (DiMaggio and Powell, 1983) in response to non-linear processes, of detailed interactions among hundreds of factors. Complex, non-linear and adaptive systems are generally not mathematically tractable and must be examined using computational techniques. Computational analysis is one of the few techniques that enable the theorist to explore the possible ramifications of such processes and to develop a series of consistent predictions. Additionally, computational models allow researchers to show proof of concept and to demonstrate that certain factors, which can be completely modeled, can or cannot generate certain phenomena. In this way, computational models can be used to show the potential legitimacy of various theoretical claims. Finally, many computational models are grounded operational theories. As such, the traditional distinction between normative and descriptive often becomes blurred in these models. Many models have an element of both. For example, the models may be descriptive at the individual level — describing individuals as boundedly rational, with various built in biases — but normative at the structural level — finding the best organizational design subject to a set of constraints. Interpretation of the results depends on which aspect of the model is being considered.

Computational theorizing has a long history in the organizational sciences. One of the earliest and best known works is the behavioral theory of the firm (Cyert & March 1963). In that work simple information processing models were used to address issues of design and performance. Today, computational organizational theorists are drawing on work from a variety of computational disciplines including: research in distributed artificial intelligence (DAI) (Levitt et al, 1994; Gasser & Majchrzak 1994; Durfee & Montgomery 1991; Ishida, Gasser, & Yokoo 1992) and that on multi-agent systems (Tambe 1996), research on adaptive agents (Holland & Miller 1991; Crowston 1994), and research on organizational design and learning (Lant 1994; Prietula & Carley 1994; Levinthal & March 1981). Some organizational models use complex adaptive agent models, such as genetic algorithms and neural networks, to answer questions about organizational change. In this case the focus is often on the evolution of industries and the sets of organizations within a market, rather than adaptation within a single organization (Axelrod 1987; Axelrod & Dion 1988; Crowston 1994, 1996; Holland 1975;

Padgett forthcoming). Other models focus on learning, particularly experiential learning by the individuals within the organization and the impact of this experience on organizational performance (Carley 1992; Lin & Carley 1997; Verhagan & Masuch 1994). Still others focus on aggregate organizational learning or search (March 1996; Lant & Mezias 1992; Lant 1994; Levinthal & March 1981). Despite discussions of the importance of considering learning at different levels (Huber 1996; Kim 1993) there has to date been little research on the possible interactions among the multiple levels of learning (for exceptions see, Carley & Svoboda 1996). This paper is a step at filling this gap in our understanding of organizational behavior.

THE VALUE OF LEARNING

A conception of organizations as information processors and search engines was presented. From this theoretical conception of organizations it follows that ultimate performance is a function of where you start and the path you follow, that learning although necessary for adaptation is not sufficient, that change is not adaptation, and that the path to success depends on factors both internal to and external to the organization, thus are many paths to success. Organizations were also characterized as being capable of learning at many levels. Learning at one level can interfere with or support learning at another level. Thus organizations with identical designs but different histories will perform differently as the structural learning affects what is learned experientially and the experiential learning affects what changes are made in the organization's structure.

This work is a step in looking at organizational adaptation in different environments. But, it is only a step. The environments examined herein are only volatile in a simple sense—the same class of task is not always being done. But what of environments where the organization is periodically faced with a truly novel task, where over time the organization sees many different tasks, rather than just cycling between two known tasks? Would such novelty affect the answers gleaned herein? In such situations, we might expect that the negative effects of downsizing would be less pronounced, and in fact that downsizing might even be advantageous. In a related vein, in this analysis change was seen as costless apart from information processing costs. If additional costs were added for change then again the value of change might be less. Clearly other potentially relevant factors have also not been considered in this analysis. Such factors, if included would alter the specific results; e.g., that large low density organizations are most adaptive in stable environments for organizations that change strategically. However, they are not likely to change the critical results — that what designs are adaptive depends on the environment and the change strategy, that change is not necessarily adaptive but become more so in a volatile environment, that history matters, and so on. Finally, one of the values of formal theorizing is that it allows the theorist to incrementally add on additional factors and work through the implications of these factors in a rigorous and consistent manner. In this way, insights about design, performance and change build up incrementally.

Hannan and Freeman's (1984) proposed a structural theory of organizational inertia suggesting that though uncommon, change was related to organizational survival. However, in an empirical study Kelly and Amburgey (1991) found that organizational change was not necessarily related to either environmental change or to survival. Their argument was that inertia was not lack of change but change in the old ways. The theoretical conception portrayed in this paper can be viewed as a refinement of both arguments. In keeping with structural inertia theory this work suggests that while change is necessary for survival it is not sufficient. In keeping with Kelly and Amburgey it is argued that for organizations, continuing to change as they have in the past, extrapolation, even in the face of environmental shifts may be dysfunctional. However, unlike Kelly and Amburgey I would suggest that extrapolation can be

beneficial. The point is, whether or not extrapolation is beneficial depends on whether the organization is moving along the correct path. The key to survival/failure is neither change nor extrapolation but determining whether or not you are on the right path. In a sense, this is a much harder search problem. Future research should be directed at determining how to determine which path is the correct one and how to jump between paths.

A possible mechanism for jumping between paths is changing how you change; i.e., strategic change. Consider how the organizations are changing in the foregoing analysis. Over time, the organizations that pursued a random or procedural change strategy did not change how they changed. Only organizations that changed strategically changed how they changed. A closer examination of the strategically changing organizations suggests that for these organizations adaptation was a matter of fine tuning. The organizations appeared to begin by getting the right number and set of people (initial judicious upsizing and downsizing). However, over time the amount of personnel changeover dropped precipitously and the amount of re-assignment and re-tasking increased. This suggests that the organizations were first growing the right set of people and then tuning themselves by moving individuals about so that they might alter whom they were working with or what they were working on.

Eccles and Crane (1988) argue that over time organizations do indeed act as annealers. This study suggests that annealing may be a valuable learning mechanism for organizations. If learning indeed occurs within the organization at many levels. And, if learning at one level can interfere with or support learning at another level, then being able to dynamically change how change is made may be critical if organizations are to adapt. The analysis presented in this paper suggests that organizations that anneal and tune have this capability.

References

- Axtell, R., R. Axelrod, J.M. Epstein and M. Cohen. 1996. Aligning Simulation Models: A Case Study and Results. Computational and Mathematical Organization Theory 1(2): 123-142.
- Axelrod, Robert and Doug Dion. Dec 9, 1988. "The Further Evolution of Cooperation." Science 242(4884): 1385-1390.
- Axelrod, Robert M. 1976. Structure of Decision: The Cognitive Maps of Political Elites. Princeton, NJ: Princeton University Press
- Axelrod, Robert M. 1987. "The Evolution of Strategies in the Iterated Prisoner's Dilemma." Pp. 32-41 in Genetic Algorithms and Simulated Annealing, edited by L. Davis. London: Pitman.
- Bruggeman, J. forthcoming. "Niche Width Theory Reappraised." Journal of Mathematical Sociology.
- Butler, R. 1993. "The evolution of the civil service — A progress report." Public Administration 71(3): 395-406.
- Carley, Kathleen and Allen Newell, 1994, "The Nature of the Social Individual." Journal of Mathematical Sociology , 19(4): 221-262
- Carley, Kathleen M. and D. Svoboda. 1996. "Modeling Organizational Adaptation as a Simulated Annealing Process." Sociological Methods and Research, 25(1): 138-168
- Carley, Kathleen and Michael Prietula, 1994, "ACTS Theory: Extending the Model of Bounded Rationality," in Kathleen Carley & Michael Prietula (Eds.), Computational Organization Theory . Hillsdale, NJ: Lawrence Earlbaum Associates.
- Carley, Kathleen M. and Zhiang Lin. forthcoming. "A Theoretical Study of Organizational Performance under Information Distortion." Management Science.
- Carley, Kathleen M. and Zhiang Lin. 1995. "Organizational Designs Suited to High Performance Under Stress." IEEE - Systems Man and Cybernetics, 25(1): 221-230.
- Carley, Kathleen M. forthcoming. "Organizational Adaptation." Annals of Operations Research, special issue on Artificial Intelligence and Management Science
- Carley, Kathleen M. 1996. "A Comparison of Artificial and Human Organizations." Journal of Economic Behavior and Organization.896: 1-17.
- Carley, Kathleen. 1992. "Organizational Learning and Personnel Turnover." Organization Science 3(1): 20-46.
- Cohen, M.D., J.G. March and J.P. Olsen, 1972. "A Garbage Can Model of Organizational Choice," Administrative Sciences Quarterly, 17(1): 1-25.

- Cohen, Michael D. and James G. March. 1974. Leadership and Ambiguity: The American College President. New York: McGraw-Hill.
- 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: LEA.
- Crowston, Kevin. 1996. "An Approach to Evolving Novel Organizational Forms." Computational and Mathematical Organization Theory. 2(1):29-48.
- Cyert, R.M. and J.G. March, 1963. A Behavioral Theory of the Firm. Prentice-Hall: Englewood Cliffs, New Jersey.
- Darr, E., Argote, L., & Epple, D. (1995). The acquisition, transfer, and depreciation of knowledge in service organizations: Productivity in franchises. Management Science, 41: 1750-1762.
- DiMaggio, Paul J. and Walter W. Powell. 1983. "The iron cage revisited: institutional isomorphism and collective rationality in organizational fields." American Sociological Review, 48: 147-160.
- Durfee, Edmund H. and Thomas A. Montgomery. 1991. "Coordination as Distributed Search in a Hierarchical Behavior Space." IEEE Transactions on Systems, Man, and Cybernetics 21(6) : 1363-1378.
- Eccles, Robert G. and Dwight B. Crane. 1988. Doing Deals: Investment Banks at Work. Boston, MA: Harvard Business School Press
- Finne, Hakon. 1991. "Organizational Adaptation to Changing Contingencies", Futures, 23(10):, 1061-1074.
- Gasser, Les and Ann Majchrzak. 1994. "ACTION Integrates Manufacturing Strategy, Design, and Planning." Pp. 133-136 in Ergonomics of Hybrid Automated Systems IV, edited by P. Kidd and W. Karwowski . Netherlands: IOS Press .
- Hannan M.T. and Freeman J. 1984. "Structural Inertia and Organizational Change." American Sociological Review. 49: 149-164.
- Hannan, Michael T. and John Freeman. 1977. "The Population Ecology of Organizations." American Journal of Sociology 82: 929-64.
- Holland, John H. 1975. Adaptation in Natural and Artificial Systems. Ann Arbor, MI: University of Michigan Press.
- Holland, John H. and John Miller. 1991. "Artificial Adaptive Individuals in Economic Theory." American Economic Review, Papers and Proceedings 81: 365-70.
- Hollenbeck, J. R., D. R. Ilgen, D.J. Sego, J. Hedlund, D.A. Major and J. Phillips. 1995. "The multi-level theory of team decision making: Decision performance in teams incorporating distributed expertise." Journal of Applied Psychology, 80: 292-316.

- Hollenbeck, J. R., D.R. Ilgen, D. Tuttle and D.J. Segó. 1995. "Team performance on monitoring tasks: An examination of decision errors in contexts requiring sustained attention." Journal of Applied Psychology, 80: 685-696.
- Hastie, R., 1986. "Experimental Evidence on Group Accuracy," In F.M. Jablin, L.L. Putnam, K.H. Roberts and L.W. Porter (Eds.), Handbook of Organizational Communication: An Interdisciplinary Perspective, Sage, Beverly Hills, CA.
- Huber, G.P. (1996). "Organizational Learning: The Contributing Processes and the Literatures." In M.D. Cohen & L.S. Sproull (Eds.) Organizational Learning. Thousand Oaks, CA: Sage.
- Ishida, Toru, Les Gasser and Makoto Yokoo. 1992. "Organization Self-Design of Distributed Production Systems." IEEE Transactions on Knowledge and Data Engineering 4(2): 123-134.
- Kelly, Dawn and Terry L. Amburgey (1991) "Organizational Inertia and Momentum: A Dynamic Model of Strategic Change" Academy of Management Journal, 34(3): 591-612.
- Kilmann, R.H. and T.J. Covin (Ed). 1988. Corporate transformation: revitalizing organizations for a competitive world. Vol 1. in The Jossey-Bass management series, San Francisco, CA: Jossey-Bass.
- Kim, D. H. 1993. "The Link Between Individual Learning and Organizational Learning." Sloan Management Review Fall issue: 37-50.
- Kirkpatrick, S., C.D. Gelatt and M.P. Vecchi. 1983. "Optimization by Simulated Annealing." Science 220(4598): 671-680.
- Krackhardt, D.1992 . "The Strength of Strong Ties: The Importance of Philos in Organizations." In Nohira, N. and R.G. Eccles (Eds.) Networks and Organizations: Structure, Form, and Action. Boston, MA: Harvard Business School Press.
- 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: LEA.
- Lawrence, Paul R. and Jay W. Lorsch. 1967. Organization and Environment: Managing Differentiation and Integration. Boston: Graduate School of Business Administration, Harvard University.
- Lant, T. K. and S. J. Mezias. 1992. "An organizational learning model of convergence and reorientation," Organization Science, 3, 1, 47-71.
- Levinthal, Daniel A. and James G. March. 1981. "A model of adaptive organizational search," Journal of Economic Behavior and Organization, 2, 307-333.
- Levinthal, Daniel A. and James G. March. 1993. "The Myopia of Learning," Strategic Management Journal, v14: 95-112.

- Levitt, B. and J. March, J. 1988. "Organizational Learning." Annual Review of Sociology, 14: 319-340.
- Levitt, Raymond E., Geoff P. Cohen, John C. Kunz, Chris I. Nass, Torre Christiansen and Yan Jin. 1994. "A Theoretical Evaluation Of Measures Of Organizational Design: Interrelationship And Performance Predictability." Pp. 1-18 in Computational Organization Theory, edited by K.M. Carley and M.J. Prietula . Hillsdale, NJ: Lawrence Erlbaum Associates.
- Lin, Zhiang and Kathleen M. Carley. 1997. "Organizational Response: The Cost Performance Tradeoff." Management Science.43(2): 217-234.
- Lin, Zhiang. 1994. Organizational Performance — Theory and Reality, Unpublished Doctoral Disserrtation, H.J. Heinz III School of Public Policy and Management. Carnegie Mellon University.
- March, James G. and Herbert Simon. 1958. Organizations. New York: John Wiley & Sons, Inc.
- March, James G. 1996. "Continuity and change in theories of organizational action," Administrative Science Quarterly, 41(2): 278-287.
- March, James G. 1996. "Exploration and Exploitation in Organizational Learning." In M.D. Cohen & L.S. Sproull (Eds.) Organizational Learning. Thousand Oaks, CA: Sage
- Mihavics, Ken and Aris Ouksel. 1996. "Learning to Align Organizational Design and Data." Computational and Mathematical Organization Theory. 1(2): 143-155.
- Ouksel, Aris and Ken Mihavics. 1995. "Organizational Structure and Information Processing Costs" Technical Report. IDS-95-1, The University of Illinois at Chicago, College of Business Administration .
- 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 Institute Studies in the Sciences of Complexity.
- Pete, A., K.R. Pattipati and D.L.Kleinman. 1993. "Distributed detection in teams with partial information: A normative descriptive model." IEEE Transactions on Systems, Man, and Cybernetics. 23: 1626-1648.
- Prietula, Michael J. and Kathleen M. Carley. 1994. "Computational Organization Theory: Autonomous Individuals and Emergent Behavior." Journal of Organizational Computing 41(1):41-83.
- Romanelli, E. 1991. "The Evolution of New Organizational Forms." Annual Review of Sociology 17: 79-103.
- 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.
- Simon, Herbert. 1955. "A Behavioral Model of Rational Choice." Quarterly Journal of Economics, 69: 99-118.
- Simon, Herbert. 1956. "Rational Choice and the Structure of the Environment." Psychological Review, 63: 129-138.
- Stinchcombe, A. 1965a. "Organization-creating organizations." Trans-actions, 2: 34-35.
- Stinchcombe, A. 1965b. "Social Structure and Organizations." Pp. 153-193 in Handbook of Organizations edited by James G. March . Chicago: Rand McNally
- Tambe, M. 1996. "Teamwork in real-world, dynamic environments." Proceedings of the International Conference on Multi-agent Systems (ICMAS).
- Tang, Z., K. R. Pattipati, and D. L. Kleinman. 1992. "A Distributed M-ary Hypothesis Testing problem with Correlated Observations." IEEE Transactions on Automatic Control, 37: 1042-1046.
- 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: LEA.