

# LEARNING WITHIN AND AMONG ORGANIZATIONS

Kathleen M. Carley

Social and Decision Sciences

and

H. J. Heinz III School of Policy and Management

Carnegie Mellon University

Running Head: Learning Within and Among Organizations

Direct all correspondence to:

Prof. Kathleen M. Carley

Dept. of Social and Decision Sciences

Carnegie Mellon University

Pittsburgh, PA 15213

email: [carley+@andrew.cmu.edu](mailto:carley+@andrew.cmu.edu)

tel: 1-412-268-3225

fax: 1-412-268-6938

## ABSTRACT

Change is readily seen both within organizations and within populations of organizations. Such change has been characterized as organization or population level learning or evolution. Underlying such change, is change at the individual human and social network level. Herein, it is asked, how does the way in which individuals learn and the way in which networks evolve reflect itself in organizational and population level learning. What changes should emerge at the organization and population level due to learning, information diffusion, and network change at the individual level? Herein it is argued that organization and population level phenomena, such as performance improvements, mis-learning, and shakeouts emerge from the on-going processes of change at the individual level. Looking at learning and information diffusion enables the organizational theorist to link micro and macro level organizational phenomena.

## LEARNING WITHIN AND AMONG ORGANIZATIONS

Few organizations can characterize themselves as unchanging. Although longevity is a sign of success, inability to, unwillingness to, or simple lack of change is not considered a sign of success. Indeed, organizations and their environments are in a continually state of flux. Change within and among organizations can have dramatic organizational consequences. Although the degree of flux, and the extent to which it has major consequences on the organization's performance and the population of organizations, varies both by industry and over time. Drawing on metaphors from human psychology and biology such change is often characterized by organizational scholars as organizational learning or evolution. The presence of this continual change, this learning, this evolution does not imply unpredictability. If we are to understand, and possibly even predict, the behavior of organizations and populations of organizations, then we will need to understand the mechanics, which bring about the observed change. We will need to understand the mechanics by which change occurs within and among organizations.

An initial basis for such mechanics lies in the underlying cognitive and network processes. That is, whether or not learning and evolution occur at the organizational level, it is irrefutably the case that individual humans do learn, exchange information, and alter their networks (who talks to, works with, reports to, whom). What changes should emerge at the organization and population level due to learning, information diffusion, and network change at the individual level?

Looking at organizations from an individual learning and network perspective enables the micro and macro levels to be linked. This linkage comes about for

several reasons. Two are particularly relevant. First, the dual focus on learning and networks leads to a set of base principles for reasoning about organizations. Second, since individual learning is a dynamic process co-occurring for all individuals in a group, group level phenomena emerge automatically.

Herein, a set of learning and network based principles are first described. These principles collectively define a cognitive-network way of thinking about organizations. Then a computational model that is consistent with these principles is briefly described. This model is used to generate a series of predictions about organization and population level phenomena. Implications from this model are then discussed.

### COGNITIVE-NETWORK PRINCIPLES

To develop our understanding of organizations and populations we need to develop a cognitive-network mechanics. In this paper, some of the basic principles of such a theory are put forward. These principles rest on research in a number of disciplines and are related to a variety of theoretical conceptions. Collectively, however, these principles provide a basis for understanding organizational change from the ground up.

The principles to be discussed here are: primacy of learning, an ecology of learning, synthetic adaptation, an ecology of networks and constraint based action. Each will be described in turn. In describing these principles the goal is not to provide an exhaustive list of all the necessary underlying principles. For example, to understand organizations in greater detail and to relate their behavior to other types of social groups, additional principles relating to issues such the nature of action, agency and knowledge, are needed (see Carley, forthcoming). Herein the

goal is more modest. It is simply to provide a discussion of a few core ideas that differentiate this perspective from other approaches to organizational learning and that provide guidance in thinking through how change occurs at the organization and population level.

### Primacy of Learning

Individual humans learn. This learning is ubiquitous - any time any place and in many ways, individuals learn. Learning is a concurrent activity in which all individuals take part. Such learning does two things. First, learning alters what individual's know. As individuals learn their mental models evolve. Second, learning alters how they relate to others (humans or artificial agents or animals) and other objects (resources, tasks, technology). From a network perspective, individual learning results in the construction of nodes (ideas) and relations (connections among ideas) within and among individuals. Moreover, as individuals learn who they share what information with changes. This can result in changes in the underlying social network (Carley 1991; Kaufer and Carley, 1993). Learning results in change not only in the network or ideas, but the network of interactions that connect individuals.

### Ecology of learning

Not all learning is of the same type. We can think of an ecology of learning types, some of which build on or rely on others and many of which can come into competition with each other or clash. There are knowledge networks (networks connecting agents to knowledge) within and between agents, networks formed of interactions, and networks formed by joint decisions. As agents learn these

networks change. There are at least two dimensions on which types of learning can be characterized — mode and activity.

For example, some of the modes in which individuals learn are observation, experience, expectation, and by being told. These different types of learning, or "learning mechanisms" differ in the feedback that is available to the individual.<sup>1</sup> Observational learning relies on a variety of types of information as feedback, with the key being that the individual learner selects which information to attend to and so use as feedback at the time of the observation. Experiential learning has its basis in task repetition by the individual and the provision of feedback by an external source. Expectation based learning occurs when individuals plan and think ahead about the future, and then use these expectations, rather than or in addition to experience, as a basis for future reasoning. Communication based learning, learning by being told, occurs when the individual learns something simply by listening without also observing, doing or planning. Although these types of learning may rely on common cognitive processes and may be related at the cognitive level, from an organizational perspective we can think of them as being relatively distinct. The important factor from an organizational point of view is that these different types of learning can be going on simultaneously. Thus, even though all individuals within the organization may be learning at the same time, what they are learning may vary in part because they are utilizing different learning mechanisms.

By activity, I mean that learning can take place in a passive or pro-active fashion. Individuals can sit back and act as information receivers, such as when they watch television or listen to a lecture. In this case, the information they learn

may or may not be new. Or they can be pro-active and go out and seek information. In this case, the intent is to seek out information that is not known; although, the individual may not be successful in acquiring exclusively novel information. Learning can occur during passive exchanges, such as when individuals garner new ideas from the mass-media. Many actions and interactions are accidental. Even such non-purposive behavior can result in stability at the organization and population level. Goal directed, or purposive, behavior can be characterized in terms of optimization; i.e., in an unchanging environment stability is reached when some function is optimized. Such purposive behavior can also result in at the organization and population level.

### Synthetic Adaptation

Composite agents can be formed through the synthesis of other agents. Humans are the quintessential agent. They are intelligent, adaptive and computational. Composite agents include teams, groups, organizations, institutions, societies, and so forth that are composed of humans. Synthetic adaptation refers to the idea that the synthesis process which creates composite agents also endows these composite agents with knowledge based properties akin to those of the underlying agents (see Carley, forthcoming). Specifically, any agent composed of intelligent, adaptive, and computational agents is also an intelligent, adaptive, and computational agent. Since humans are intelligent, adaptive and computational all teams, groups, organizations, institutions, societies, and so forth that are composed of humans are also intelligent, adaptive and computational agents.

Through synthetic adaptation composite agents can form which can interact with and perform the same tasks as non-composite agents. In response to natural or technological disasters, there are different composite and non-composite agents all of which act as response units — single companies, consortiums of companies, network organizations, and groups of individuals acting collectively as an institutional unit all play a similar role (Dynes and Quarantelli, 1968; Topper and Carley, 1997). Synthetic adaptation enables the composite agent to take action distinct from, and not predicted on, an aggregation of individual actions. Thus the organization, in and of itself, is an intelligent, adaptive and computational entity (Carley and Gasser, forthcoming) whose capabilities result from the detailed, ongoing, interactions among, decisions of and behavior of the member agents. Group behavior emerges from complex interactions and concurrent activity and not through simply aggregation. This conception is consistent with the work on distributed cognition (Hutchins, 1991, 1995) and transactive memory (Wegner, 1987, 1995; Moreland et al., 1996; Moreland, in press).

One result of synthetic adaptation is that learning can and does occur at multiple levels — e.g., individual, organization, and population. It is particularly useful to distinguish between individual and structural learning (Carley and Svoboda, 1996; Carley and Lee, 1998). Individual learning occurs within the agent. As agents alter their mental models by adding or dropping either ideas and/or relations among ideas we say that individual learning has occurred. Such changes may precipitate changes in interaction among agents (Carley, 1991; Kaufer and Carley, 1993). In this way individual learning mediates structural learning. When the agents are human we think of such changes as changed in the agent's



mental model. When the agents are composites, such as teams, organizations, or populations we think of such individual learning in terms of the shared mental model, group knowledge, or culture. Structural learning occurs among agents and within composite agents (such as groups, organizations, or institutions). Structural learning occurs as the composite agent adds or drops member agents (individual or composite) and/or the relations among member agents. Changes in interaction that result from such structural learning can influence others knowledge and so attitudes and beliefs (Krackhardt and Porter, 1985). Individual and structural learning can occur at the organization and population level.

### Ecology of Networks

Within organization theory the term network is often used to refer to the set of social or work related connections among individuals within an organization or institutional or economic arrangements among organizations. However there are many other networks. Indeed, we can think of networks as existing within an ecology of networks. For example, the social network denoting who talks to whom is intertwined with each individual's cognitive network (the way in which each individual links ideas, i.e., the individual's mental model) and the transactive knowledge network (each individual's perception of the network linking people to their ideas). Within organizations, the authority or reporting network (who reports to whom) is connected to many other networks including the task structure (which tasks are connected to which), the task access structure (who is assigned to what task), and the capabilities networks (who has what capabilities or access to what resources).

Change in any network results in a cascade of complex repercussions for the other networks in this ecology. Consider the space defined by the interaction networks and the knowledge network. One of the most common bases for interaction is similarity. One basis for similarity is shared knowledge. Change in who knows what (the knowledge network) alters the distribution of knowledge and who is similar to whom. This results in changes in the interaction network. Thus, learning results in individuals moving about in the interaction-knowledge space. As individuals within organizations learn, the organization's knowledge also changes. Thus, even as learning results in people, organizations, and populations of organizations moving through the interaction-knowledge space.

#### Constraint Based Action

Cognition and networks mediates agent action. Interaction is a function of external constraints (opportunity), cognitive and knowledge constraints (such as agent's perception of their relative similarity to others). The actions of all agents, synthetic or not, are constrained and enabled by the learning mechanisms they employ and the networks in which they are embedded. If we imagine a hierarchy of agents then higher-order composite agents constrain the actions of their member agents. The networks in which agents are embedded influence and constrain individual and group behavior (McPherson, 1983) and serve to constrain and facilitate change (Granovetter, 1985). Constraints reduce the set of potential actions to the set of acceptable actions. Constraints can be so severe that they define all activity; e.g., certain assembly line technologies define a network ordering of tasks that severely constrains which agent does what when. This notion of constraint based action appears in the information processing (March and Simon, 1958;

Galbraith, 1973), social information processing (Salancik and Pfeffer, 1978) and resource dependency perspectives (Pfeffer and Salancik, 1978). However, from an applied or modeling perspective, specifying constraints requires more than a recognition that constraints exist. To really evaluate the impact of constraints, and to make concrete predictions, the precise set of tasks, networks, institutions, resources, knowledge, agents and technology that affect the flow of information need to be specified (Carley and Prietula, 1994).

### MODEL AND METHODOLOGICAL APPROACH

Computational models consistent with these principles can be, and have been, used to illustrate, explain and theorize about organizational behavior. Using such a computational model, ORGAHEAD<sup>2</sup>, some of the implications of this theoretical approach are described and illustrated for organizations. ORGAHEAD has been previously described in detail in the literature (Carley and Svoboda, 1996; Carley and Lee, 1998; Carley, 1998). Thus, only a cursory description is provided below, with a focus on the learning aspects.

ORGAHEAD is a model of organizational performance and dynamics. ORGAHEAD has been informed by empirical studies on human learning and adaptation within human organizations. Within ORGAHEAD both individual and structural learning are present. Organizational action proceeds at both the operational and strategic level. Operational personnel learn as they work on a series of quasi-repeated tasks, which can vary in complexity. Both experiential and communication mechanisms are employed by these agents. The position of agents in the organization (as defined by the authority network, communication network, capabilities network, etc.) constrains what information is available to whom and

who communicates with whom. This in turn affects what is learned by whom. At the strategic level, the CEO (or team manager) employs both experiential and expectation mechanisms. The CEO makes a decision about the task for the organization as a whole. CEOs can alter the organizational structure and so alter access to information. All agents have mental models which contain a task model, past experiences, expectations, and a knowledge network for subordinates (which subordinates know what). The CEO's mental model includes information on organizational performance, who knows what, previous structural changes, and expectations about alternative structures. ORGAHEAD has been used to look at a variety of organizational issues including the impact of different change strategies, training scenarios, constraints on organizational re-design, in both stable and changing environments.

ORGAHEAD is a dual-level model of organizational adaptation in which the organization can change at both the strategic and the operational level. At the operational level the organization is modeled as a collection of adaptive agents, each of whom occupies a particular organizational position, can learn over time, and gain task based experience. Individual learning is operationalized using a variant of a Bush and Mosteller (1955) stochastic learning model with additional limits on attention, memory, and information processing which bounds the agent's rationality more than in the original stochastic models. At the strategic level, the organization adapts strategically in response to changes in its performance through structural learning (which results in hiring, downsizing, expansion, and re-engineering). Structural learning is operationalized as a simulated annealing process.

Why Use a Computational Model?

Recent advances in computational analysis and distributed artificial intelligence (DAI) suggest that multi-agent models can be usefully employed for theory creation in the domain of organizational dynamics. In these models, core organizational processes are abstracted so as to lay bare the relationships among the various key components of organizational design and adaptation. Despite such abstraction, complex non-linear processes are a central feature of these models. Computational analysis is one of the few techniques that enables the theorist to think through the possible ramifications of non-linear processes and to develop a series of consistent predictions. Within organization theory, computational models can be, and have been, used in a normative (and sometimes a prescriptive) fashion to generate hypotheses that can then be tested in other empirical settings. Running a virtual experiment<sup>i</sup> on the model and then statistically analyzing the result generates hypotheses.

Adaptive agent models, such as genetic algorithms and neural networks, are being used to answer questions about the evolution of industries and the sets of organizations within a market (Axelrod 1987; Axelrod and Dion 1988; Crowston 1994; Holland 1975; Holland and Miller 1991; Padgett, 1997). Experiential and symbolic learning models are being used to answer questions about turnover, organizational learning, CEO activity, agent activity (Carley 1992; Lin and Carley forthcoming; Verhagan and Masuch 1994; Lant 1994), learning ecologies (Carley and Svoboda, 1996; Kim 1993). Issues examined include coordination and communication (Levitt et al, 1994; Durfee and Montgomery 1991; Ishida, Gasser

---

and Yokoo 1992), planning (Gasser and Majchrzak 1994), monitoring (Elofson and Konsynski 1993) and socially shared cognition (Hutchins 1990, 1991).

### Virtual Experiments

A virtual experiment is an experiment conducted using a computer model. They are computational analogues of human laboratory experiments. The agents are computational not human beings or animals. The term simulation has been used in many ways. For example, it has been used to refer to both the computational model and the result of running that computational model at least once. Modern computational models define a space of options that is sufficiently large that virtual experiments become a feasible technique for defining the response space. One difference here from a laboratory experiment is that many more data points can be collected. Using ORGAHEAD two different virtual experiments were done to look at the impact of individual and structural learning on change at the organization and population level.

Experiment 1 was done to explore the over time behavior for a small population of organizations for an extremely long period of time. In this case, the behavior of a set of 100 organizations varying in structure and faced with a sequence of 40,000 tasks (one per time period) was simulated. Experiment 2 was done to explore the types of organizations and change patterns that emerged over time within a large population. In this case, the behavior of a set of 1000 organizations varying in structure and faced with a sequence of 20,000 tasks (one per time period) was simulated. Analysis shows that there is no fundamental difference in the way the organizations behaved in these two experiments.

In both experiments the organizations were defined such that their initial size and authority/communication structure was selected at random from all possible organizations with between 2 and 45 personnel and 1 to 4 levels in the hierarchy. In each organization the individual agents, or people, could learn and the organization as a whole engaged in structural learning. Agents were treated as being boundedly rational in terms of how many pieces of the task they could process (at most 7), how many others they can communicate with (at most 7) in one time period, how much they can remember. The CEO could alter the organization by hiring or firing personnel, changing who reported to whom, and changing who did what. The number of such changes in one time period could vary between 1 and 15. Initially, all such changes were equally likely. All organizations faced a set of 2000 tasks, one per time period. This enables sufficient time for noticeable patterns of organization and population level change to emerge. The over-arching task was a classification choice task. Variables measured include final performance (percentage of final 500 tasks for which the organization made the correct decision, number of structural changes, the size, density, number of isolated personnel, number of task factors ignored, and redundancy in the organization initially and at the end.

### IMPLICATIONS FOR CHANGE

Organization and population level change emerges out of the ongoing interactions among intelligent adaptive agents. Not only is interaction the fundamental social act, it is also the basis for individual behavior within organizations and the structuring of ties among organizations (Carley, 1991; Carley and Newell, 1994). For example, the structure of the networks and the ability to

learn via communication can alter the rate at which the team learns and affect the diffusion of erroneous beliefs (Prietula and Carley, 1994). Further, sets of intelligent agents are self-organizing due to capability and knowledge constraints (Padgett, 1997; Epstein and Axtell, 1997; Kauffman, 1995). This self organization leads to the structuration of the organizational field. Thus, regularities in behavior, organizational structure, networks, and culture across, among, and over time for agents and composite agents emerge as agents interact. Consequently, the specifics of who interacts with whom when determines which norms, regulations, what type of culture, what type of organizational structure, and so on emerges (White, 1992).

These ideas are supported by the results from the ORGAHEAD model. Moreover, using ORGAHEAD, we gain a more precise notion of what it means for interactions among intelligent agents to generate a structured field of organizations. Structuration leads to the following three phenomena. First, there will ultimately be a quasi-stable pattern of behaviors; e.g., a ranking among firms will appear. Second, in the process of moving toward this stability, more or less similar firms will differentiate themselves thus resulting in a stratified industry. Third, firms will lock into strategies of change and these strategies will differ across industrial strata. However, these change strategies will not result in industrial homogeneity.

### Quasi-Stability

This trend toward stability is a projected eventuality. In most cases, disruptions will occur long before stability is reached and may prevent stability from ever being achieved. The issue then is not — "what is the stable configuration?", but "what happens along the way to stability?". The results here suggest that along the way to stability, stratification and shakeouts occur. The basic



cognitive-network mechanisms and the learning processes that cause short term oscillations in interaction and performance also cause organizations to enter into strategic choice sets through which they become increasingly differentiated. Many oscillations in performance and the resultant shakeouts are a result of collisions and collusions between the types of learning.

At both the organization and population level we see a general trend to stability. At the population level this trend is seen in that, in a stable environment, initially more or less equivalent firms fall into a rough rank order in performance over time. That is, clear top and bottom performers emerge. This is only a quasi-stable situation however. Due to the on-going learning the exact ranking will continue to oscillate. However, the degree of oscillation will decrease. An illustration of this appears in Figure 1.

#### INSERT FIGURE 1 ABOUT HERE

The resultant behavior for the three organizations that ended up with the highest and lowest performance is shown in Figure 1. This figure illustrates that despite individual and structural learning, over time the variation in performance decreases for each organization even though for the population the performance bifurcates. A familiar example of such ranking shifts is the ranking of universities by national magazines. Over time, however, these organizations lock into a structure, a pattern for how to change it, and the individuals into patterned ways of responding to the environment. This trend toward quasi-stability, in a stable environment can be beneficial for those organizations that locate satisfactory structures and patterns of change. However, for those organizations that lock into

poor change strategies, simply being able to learn will not alter their long-term trajectories.

In the short term, relative to the long term, there is greater oscillation in organizational behavior and firms move in their relative rankings quite dramatically. Notice in Figure 1 that all organizations are initially undifferentiated. As learning occurs both at the individual and structural level, the population of organizations stratifies. Despite appearances, the behavior of organizations and the industry is not randomly generated. Rather, this pattern of change is for each organization a reflection of the historic path it has followed in altering its structure, responding to expected changes in the environment or its own performance, and the specific nature of what the agents within the organization have learned. Change in relative ranking across organizations are caused by different responses to the same task. These different responses are the result of the specific path of adaptation chosen by the organization; i.e. history matters and there is path dependency. The specific networks within the organization, the learning procedures within agents, the specific sequence of tasks, and who learned what information when lead organizations, even those with identical structures, to respond to the same task differently. Moreover, organizations adapt their structure given the same performance differently. There are many consequences of this differential response. For example, in some organizations the individual learning that results as the task changes clashes with the structural learning thus degrading performance; whereas in other organizations the two types of learning at that point may work synergistically to dramatically improve performance. An example of a clash is when downsizing results in laying off core talent for doing a particular task;

whereas, a synergy would occur when downsizing results in laying off the key individuals who were generating erroneous decisions. Thus, the process that generates what appear to be random fluctuations, is actually non-random.

### Industrial Stratification

This trend toward ultimate quasi-stability does not imply a trend toward a single outcome. In fact, in most situations there are multiple endpoints and stability means that industrial stratification has occurred. For example, the ORGAHEAD simulations (experiment 1 and 2) demonstrate that over time similar organizations will diverge into high and low performers. From a complexity standpoint we can say that there are multiple attractors to which the organizations can gravitate. Once the division has taken place most organizations will stay in their performance strata.

As performance stratification happens the internal structure of high and low performance organizations will come to be different (Carley, 1998). As noted in Table 1, based on experiment 2, in the long run the top performers will not ignore any of the information needed to do the task, will be less dense (in who talks to whom) and more redundant (in who does what) than will be low performers. They will also be slightly (though not significantly) larger.

### INSERT TABLE 1 ABOUT HERE

Even though organizations begin very similarly they will tend to take different paths despite using the same learning mechanisms. Organizations will get locked into patterns of change where the internal knowledge and interaction networks that develop lock them into a particular pattern of performance. Divergence occurs because different organizations learn different things which affects whether and when the various learning mechanisms clash or act synergistically. What this means

for an organization is that learning is not guaranteed to benefit the organization. Organizational differences result from the content of what is learned, from local choices that move the organization through the interaction-knowledge space

#### INSERT FIGURE 2 ABOUT HERE

This divergence is seen in a separation in the structure of the high and low performers (see Figure 2). In Figure 2 we see that over time, the size of the top performers increases and the size of the bottom performers decrease. In Figure 2, the mid point of the circle is the mean size of the organizations in that quartile at that time, the size of the circle is the standard deviation. The overall initial and final distribution of sizes is also shown. Over time, the population of organizations is becoming more heterogeneous in their structure.

Learning clashes result in diversity at the population level. Thus, there is not a shared experience at the population level in terms of what structure is the best. An organization that exhibits high performance has located a structure and a change strategy that is well suited for the current task environment. Poor performance organizations have not located such an ideal. However, structures and change strategies that do well in one task environment may not do well in other environments. Thus, it may be that leaning clashes, although detrimental to a particular organization, may result in sufficient population diversity that in an environmental shift there may be a pre-existing set of organizations with the appropriate structure and change strategy.

In addition to diversification, learning clashes also can result in dysfunctionality. For example, in most organizations there are a few people who are isolated. These are people who, by choice or happenstance, are loners and

rarely interact with or share information with others. In some sense their knowledge and expertise is lost to the rest of the organization. Due to the potential for lost opportunities, we might think of the organization as being more dysfunctional if there are more isolates. In experiment 2 we find that over time, for most organizations, the quest for a more high performing structure leads organizations to adopt forms in which there are fewer isolates (see Figure 3). The quest for high organizational performance results in more common individual experience. This occurs without these organizations necessarily adopting a strategy of firing these isolates. Despite this trend, low performing organizations end up having more isolates than do high performing organizations. In Figure 3, the mid point of the circle is the mean number of isolates in the organizations in that quartile at that time, the size of the circle is the standard deviation. The overall initial and final distribution of the number of isolates is also shown. In the case of isolates, organizations are, as a population, becoming more homogeneous. The population of organizations has collectively learned, without mimicry, that isolates detract from performance. However, high performing organizations have learned the lesson better and will still have fewer isolates and exhibit less variance among themselves than will low performers. The learning clashes between individual and structural learning result in low performing organizations, despite appropriate changes, retaining their dysfunctionality and coming to appear more dysfunctional than their high performing counterparts.

INSERT FIGURE 3 ABOUT HERE

Stratification Strategies and Complexity

For many tasks, and for the task used in these experiments, these attractors to which the organizations gravitate do not take the form of "the right organizational structure for this task." That is, there are structures that are known to perform optimally, and the simulated organizations do not all gravitate to these forms despite attempts to improve performance. For high performance organizations, the attractors are patterns of change. Over time, the way in which these organizations change also diverges (see Table 2). In Table 2 the percentage difference in final and initial conditions for each of the structural variables is shown. There are two points to notice. First, top and bottom performers do not always change in the same direction. As noted earlier, learning clashes within organizations can result in diversification and stratification at the population level. Second, regardless of whether or not low and high performers change in the same direction, the variance in change is much less for the high performers than the low performers. In other words, high performing organizations may not move to the same organizational form, but they tend to change at the same rate. This is reflected in the fact that the variance in the distribution of the various structural indicators may or may not be lower for high than low performers; however, the variance in the percentage change in these indicators over time is always much lower for high performers than low performers. In this sense, the experience that is shared by high performers is not experience in operating within a particular form, but experience in process and degree of change.

INSERT TABLE 2 ABOUT HERE

## DISCUSSION

The argument here is quite simple, a unified theory of organizational change should be possible if it has at its basis a processual approach to learning. Principles underlying such an approach are: the primacy of learning, a learning ecology, synthetic adaptation, a network ecology, and constraint based action. Complex behavior within and among organizations result from simple learning mechanisms operating within a system of constraints and an ecology of learning mechanisms and changing networks. Importantly, these networks include both knowledge and interaction networks all of which change dynamically as agents learn.

There are many ramifications of saying that interaction and learning at the individual level result in organizational movement through the interaction-knowledge space. One implication is that structure and culture emerge at the organization and population level as agents interact. This means that there is no ontological imperative that gives one agent (such as the CEO, corporate founder, or funding agency) more apriori import on an organization's or industry's ultimate future. This is not to say that specific individuals do not play an important role, merely to say that the behavior of all individuals collectively is what is key.

Organizations that begin quite similar can end up very different. The exact same learning mechanisms result in divergence in form, strategy and performance for organizations. Learning in and of itself will result in a stratification of the organizational landscape and a potential shakeout in terms of firm survival. Firms do not fail because they are not learning, but because they are learning the wrong things, or learning the right things in the wrong order and at the wrong time. Path, or history, matters. What is interesting is despite the importance of history, and hence the exact networks within and among organizations, the organizational

structures and patterns of change that emerge at the population level are not random. For individuals, experiential learning results in improved performance, which can be mitigated by changes in the structure of the organization (structural learning). For an organization, individual learning results in performance improvements, but the structure of the organization places a cap on these improvements. Structural learning can remove these caps but can result in clashes or synergies with individual or experiential learning. These clashes or synergies results at the population level in industrial stratification and the appearance of mimicry in process among high performers.

In general, this work illustrates how complex dynamics at the organization and population level result from simple, but non-linear processes, at the individual level. Consequently, thinking through the implications of adaptation processes is non-trivial even given simple starting conditions. For example, herein, all the behavior that emerged did so under the following conditions. All organizations began with different, but randomly chosen structures. All individuals learned in the same way. All organizations were faced with the same sequence of tasks and receive the same feedback. Individual, organizational, and population level changes that emerged did so because of the specific networks, the position of individuals and organizations in the interaction-knowledge space, and their specific history of response to events. Complexity emerged simply do to learning occurring within a system of cognitive and network based constraints.

Taking a population ecology perspective, we recognize that organizations can, unlike those in the foregoing virtual experiments, cease to exist. A cause of such organizational death is low performance. In the foregoing simulations



performance was measured as accuracy — the percentage of tasks in a decision window that were correctly classified by the organization. If we expect returns on accuracy, then we can think of this as a proxy for profits. Low accuracy would indicate low profits. Low profits, or lack of profits, over an extended period could lead to the demise of the organization. Such determination of the economic well being of an organization would occur at a slower rate than the task decision-making rate. That is, it is reasonable to assume that the organization does not fail after misclassifying a single task. Rather, sustained poor performance is more likely to result in organizational demise. Based on these assumptions we can ask what would be the impact of organizational death?

If in fact, organizations only fail if their performance goes under a particular threshold for an extended period of time then what should happen, in the absence of mimicry, is that there will be an initial shakeout in most industries. This shakeout will occur as organizations get locked into change strategies that push them to the lower strata. The severity of the shakeout and when it occurs will vary by industry depending on initial conditions and the complexity of the task. A consequence of the shakeout will be increased homogeneity in the population; although this homogeneity will be far from perfect. Since learning does not stop, at any level, organizations will continue to change after the initial shakeout. However, to the extent that they have locked into high performance change strategies then fewer organizations are likely to fail at later dates than in the initial shakeout (barring major changes in technology or environment). Additionally, it is likely that it the organizations that are in the middle strata that will be most at risk. Future work using these computational models in conjunction with population level learning

models are needed to make predictions about the distribution and timing of these shakeouts.

At the population level another type of learning exists that has not been considered — mimicry. Mimicry, may alter the severity and timing of the shakeout, as well. We see among high performers behavior that looks like mimicry of process. However, there is no process in the model for individuals to transfer among organizations or to learn via communication from those in other organizations. The apparent mimicry among high performers occurs for two reasons. First, to the extent that structures have become more similar through structural learning will result in the same individual learning which means similarity in the expectation based learning among the CEOs which results in them taking similar actions. In other words, they have learned to parse incoming information on the task in the same or similar ways. Second, these organizations have locked into similar change strategies. This means that they have learned to parse information on how their own organization operates and responds to the environment in the same or similar ways.

Learning clashes for the population, result in higher diversity in form on some dimensions, and higher diversity in form for all low performers on all dimensions. In this sense, there are many ways to fail and only a few ways to succeed. But these few ways to succeed are not simple algorithms for the structure of the organization. Rather, they are a complex historical pattern of changes, networks, and cognitions all connected through simple but ubiquitous learning mechanisms.

## ACKNOWLEDGMENT

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 under Grant No. 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, the National Science Foundation, or the U.S. government.

---

<sup>1</sup> A large number of individual learning models exist (for a review see Pew and Mavor, 1998).

Some researchers argue that a single learning mechanism may be sufficient to account for all of the apparent different types of learning. Whether or not this is the case, from the vantage of organizations, learning types that take different feedback or occur at different rate have different organizational implications.

<sup>2</sup> ORGAHEAD is written in a combination of C, PERLSSCRIPT, and C++ and runs only on UNIX platforms. Interested readers should contact the author to determine the most feasible mode of access.

## REFERENCES

- Axelrod, Robert and Doug Dion. Dec 9, 1988. "The Further Evolution of Cooperation." *Science* 242(4884): 1385-1390.
- 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.
- Bush, Robert R. and Frederick Mosteller, 1955. *Stochastic Models for Learning*. New York, NY: Wiley.
- Carley, Kathleen 1992, "Organizational Learning and Personnel Turnover." *Organization Science*, 3(1): 20-46.
- Carley, Kathleen and Allen Newell. 1994. "The Nature of the Social Agent." *Journal of Mathematical Sociology* 19(4): 221-262.
- Carley, Kathleen and Michael Prietula. 1994. "ACTS Theory: Extending the Model of Bounded Rationality" in *Computational Organization Theory*, edited by Kathleen Carley and Michael Prietula. Hillsdale, NJ: Lawrence Earlbaum Associates.
- Carley, Kathleen M. forthcoming. "On the Evolution of Social and Organizational Networks." In David Knoke and Steve Andrews (Eds.) special issue of Research on the Sociology of Organizations. on *Networks In and Around Organizations*." JAI Press, Inc.
- Carley, Kathleen M. 1991. "A Theory of Group Stability." *American Sociological Review* 56(3): 331-354.

- Carley, Kathleen M. 1998, "Organizational Adaptation." *Annals of Operations Research*. 75: 25-47.
- Carley, Kathleen M. and David M. Svoboda. 1996. "Modeling Organizational Adaptation as a Simulated Annealing Process." *Sociological Methods and Research* 25(1): 138-168.
- Carley, Kathleen M. and Ju-Sung Lee. 1998, Dynamic Organizations: Organizational Adaptation in a Changing Environment. Ch. 15 (pp. 269-297) in Joel Baum (Ed.) *Advances in Strategic Management, Vol 15, Disciplinary Roots of Strategic Management Research*. JAI Press.
- Carley, Kathleen M. and Les Gasser. forthcoming. "Computational Organization Theory." Ch. 7 in *Distributed Artificial Intelligence* edited by Gerhard Weiss. Cambridge, MA: MIT Press.
- Crowston, Kevin. 1994. "Evolving Novel Organizational Forms," in K.M. Carley and M.J. Prietula (Eds.) *Computational Organization Theory*, Hillsdale, NJ: Lawrence Erlbaum Associates.
- 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.
- Dynes, Russell R. and E. L Quarantelli. 1968. "Group behavior under stress: A required convergence of organizational and collective behavior perspectives." *Sociology and Social Research* 52:416-429
- Elofson, Gregg S and Benn R. Konsynski. 1993. "Performing Organizational Learning with Machine Apprentices." *Decision Support Systems*, 10(2): 109-119.

- Epstein, Joshua and Rob Axtell. 1997. *Growing Artificial Societies*. Boston, MA: MIT Press.
- Galbraith, Jay 1973. *Designing Complex Organizations*. Reading, MA: Addison-Wesley.
- Gasser, L. and A. Majchrzak. 1994. "ACTION Integrates Manufacturing Strategy, Design, and Planning," in P. Kidd and W. Karwowski (Eds) *Ergonomics of Hybrid Automated Systems IV*, IOS Press, Netherlands.
- Granovetter, Mark. 1985. "Economic Action and Social Structure: The Problem of Embeddedness," *American Journal of Sociology* 91: 481-510
- Holland, J.H. and J. Miller. 1991. "Artificial Adaptive Agents in Economic Theory," *American Economic Review, Papers and Proceedings* 81:365--70.
- Holland, John H. 1975. *Adaptation in Natural and Artificial Systems*. Ann Arbor, MI: University of Michigan Press.
- Hutchins, Edwin. 1991. "The Social Organization of Distributed Cognition." In *Perspectives on Socially Shared Cognition*, edited by L. B. Resnick, J. M. Levine and S. D. Teasley. Washington D.C.: American Psychological Association.
- Hutchins, Edwin. 1990. "The Technology of Team Navigation." Pp. 191-220 in *Intellectual Teamwork*, edited by J. Galegher, R. Kraut and C. Egido, Hillsdale, N.J.: LEA.
- Hutchins, Edwin. 1995. *Cognition in the Wild*. Cambridge, MA: MIT Press.
- 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.

- Kaufer, David S. and Kathleen M. Carley. 1993. *Communication at a Distance: The Effect of Print on Socio-Cultural Organization and Change*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Kauffman, Stuart A. 1995. *At home in the universe : the search for laws of self-organization and complexity*. New York, NY: Oxford University Press.
- Kim, D. H. 1993. "The Link Between Individual Learning and Organizational Learning." *Sloan Management Review* Fall issue: 37-50.
- Krackhardt, David and Lyman Porter. 1985. "When Friends Leave: A Structural Analysis of the Relationship between Turnover and Stayers' Attitudes," *Administrative Science Quarterly*, 30(2): 242-261.
- 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.
- 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 Carley. 1993. "Proactive or Reactive: An Analysis of the Effect of Agent Style on Organizational Decision Making Performance." *International Journal of Intelligent Systems in Accounting, Finance and Management* , 2(4): 271-288.



- March, James G. and Herbert A. Simon. 1958. *Organizations*. New York, NY: Wiley.
- McPherson, J. Miller. 1983. "An Ecology of Affiliation," *American Sociological Review* 48:519-532.
- Moreland, Richard L. forthcoming. "Transactive memory in work groups and organizations." In *Shared knowledge in organizations*, edited by L. Thompson, D. Messick, and J. Levine. Mahwah, N.J: Lawrence Erlbaum Associates
- Moreland, Richard L., Linda Argote and R. Krishnan. 1996. "Socially shared cognition at work: Transactive memory and group performance." Pp. 57-84 in *What's social about social cognition? Research on socially shared cognition in small groups*, edited by J.L. Nye and A.M. Brower. Newbury Park, CA: Sage.
- Padgett, John F. 1997. "The Emergence of Simple Ecologies of Skill." Pp. 199-222 in *The Economy as an Evolving Complex System II* Vol. XXVII, edited by B. Arthur, S. Durlauf and D. Lane. Reading MA: Addison-Wesley.
- Pfeffer, Jeffrey, and Gerald R. Salancik. 1978. *The External Control of Organizations: A Resource Dependence Perspective*. NY: Harper and Row.
- Prietula, Michael J. and Kathleen M. Carley. 1994. "Computational Organization Theory: Autonomous Agents and Emergent Behavior," *Journal of Organizational Computing* 41(1): 41-83.
- Salancik, Gerald R. and Jeffrey Pfeffer. 1978. "A social information processing approach to job attitudes and task design," *Administrative Science Quarterly* 23: 224-253.

- Sanil, Ashish, David Banks and Kathleen Carley. 1995. "Models for Evolving Fixed Node Networks: Model Fitting and Model Testing," *Social Networks* 17(1): 65-81.
- Stryker, Sheldon. 1980. *Symbolic Interactionism*. Menlo Park: Benjamin Cummings.
- Topper, Curtis and Kathleen M. Carley. 1997. "A Structural Perspective on the Emergence of Network Organizations," Working Paper, H. John Heinz III School of Public Policy and Management, Carnegie Mellon University, Pittsburgh, PA.
- Verhagen, H. and M. Masuch. 1994. "TASCCS: a synthesis of Double-AISS and Plural-SOAR," in K.M. Carley and M.J. Prietula (Eds.) *Computational Organization Theory*, Hillsdale, NJ: Lawrence Erlbaum Associates.
- Wegner, Daniel M. 1987. "Transactive memory: A contemporary analysis of the group mind." Pp. 185-208 in *Theories of group behavior*, edited by B. Mullen and G. R. Goethals. New York: Springer-Verlag.
- Wegner, Daniel M. 1995. "A computer network model of human transactive memory," *Social Cognition* 13(3): 319-339.
- White, Harrison C. 1992. *Identity and Control: A Structural Theory of Social Action*. Princeton, NJ: Princeton University Press.

Figure 1. Over time changes in performance for lowest and highest performing organizations

Figure 2. Change in size of organizations over time.

Figure 3. Change in the number of isolates in organizations over time.

---

Table 1. Predicting performance from structural features.

---

<u>Variable</u>	<u>Coefficient</u>	<u>Std Coefficient</u>
Constant	74.89 ***	0.00 ***
Final Size	0.05	0.03
Final Number of Isolates	-0.10	-0.02
Final Amount of Ignored Information	-2.23 ***	-0.28 ***
Final Density	-14.64 ***	-0.09 ***
Final Redundancy	1.152 ***	0.28 ***

R2 = 0.321, \* < 0.05 \*\* < 0.01 \*\*\* < 0.005

---

---

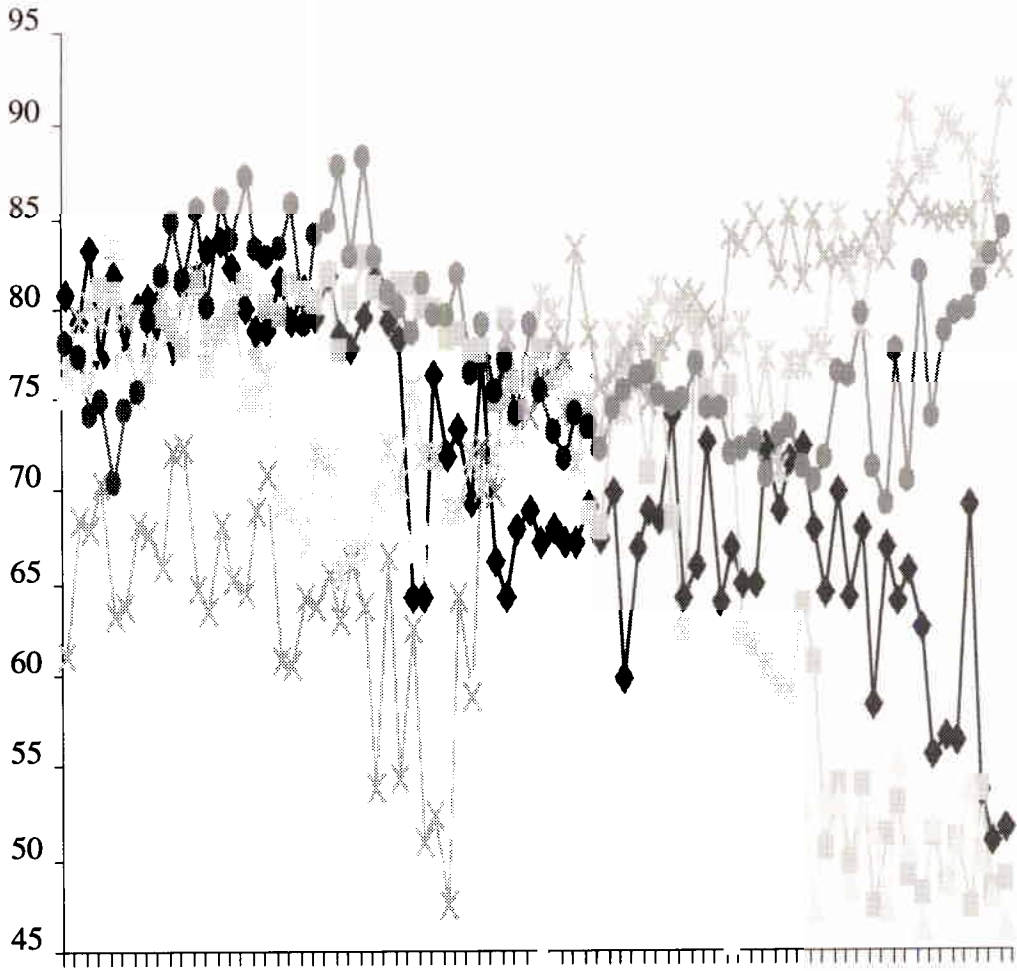
Table 2. Diversification in change strategies.

---

<u>Change in</u>	<u>Low Performers</u>		<u>High Performers</u>	
	mean	Std.dev.	mean	Std.dev.
Size	22.54	134.08	31.35	1.48
Isolates	23.32	151.58	-17.49	86.56
Overlooked Information	45.85	137.30	-13.80	45.47
Density	-11.00	73.31	-28.92	50.74
Redundancy	-4.49	59.80	33.61	57.54

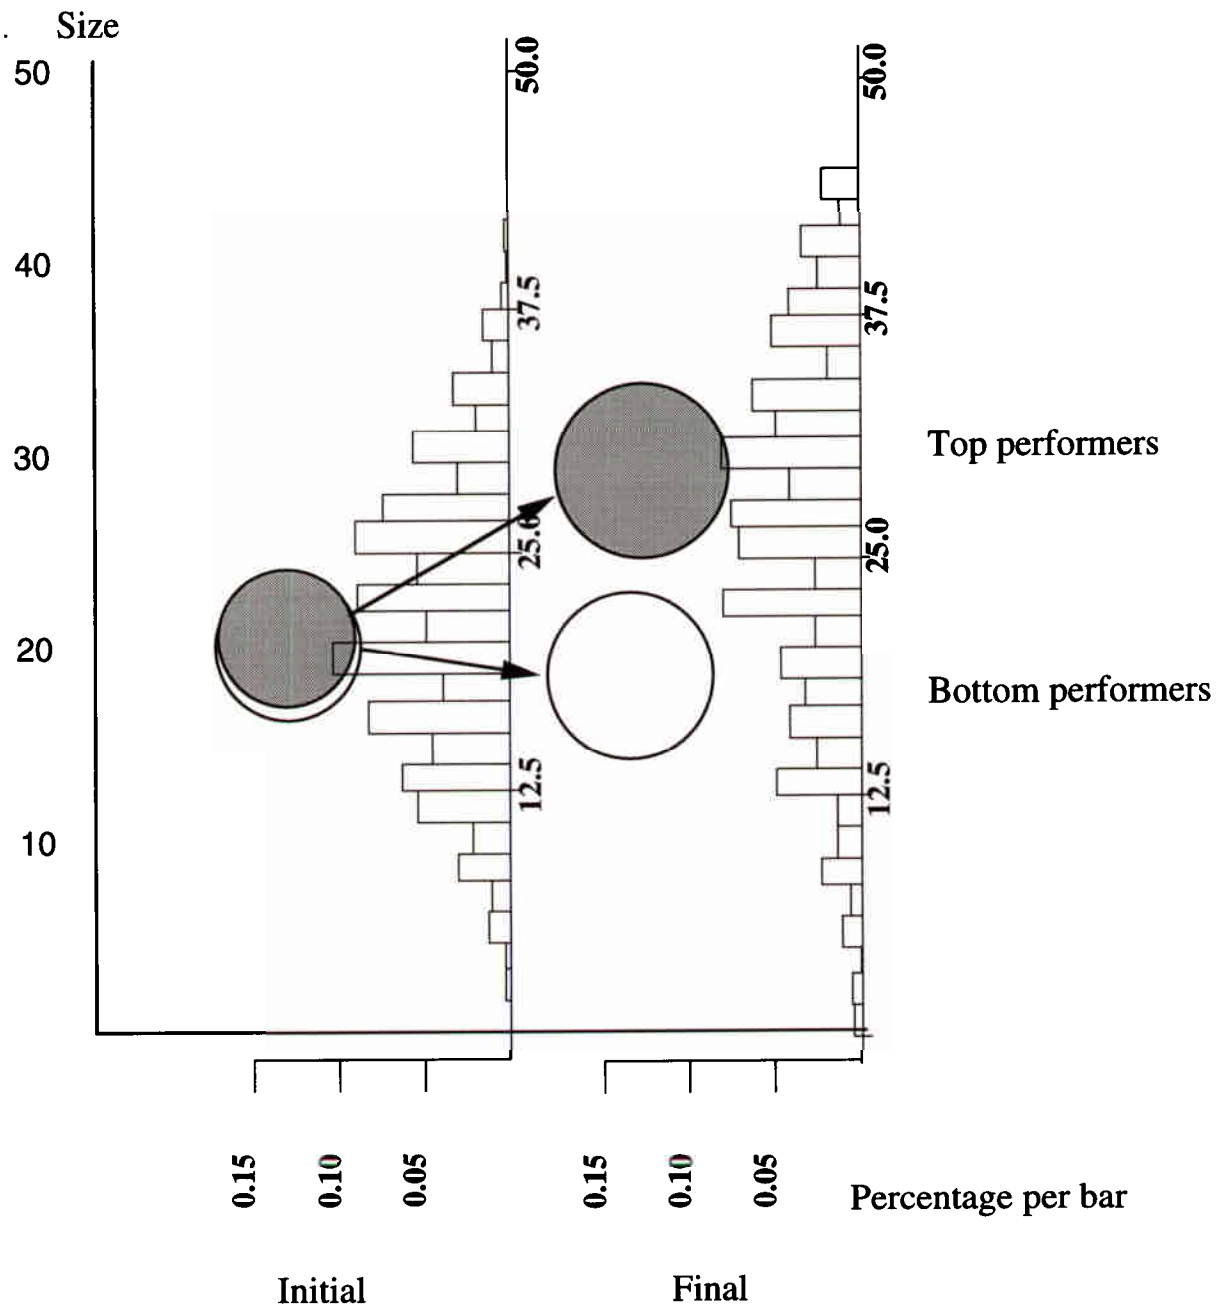
---

Performance



250 2250 4250 6250 8250 10250 12250 14250 16250 18250 20250 22250 24250 26250 28250 30250 32250 34250 36250 38250

Tasks/ Time



Number of Isolates

