

**STRUCTURAL CHANGE AND LEARNING WITHIN
ORGANIZATIONS**

Kathleen M. Carley
Carnegie Mellon University

Vanessa Hill
Winthrop University

Direct all correspondence to:

Kathleen M. Carley
Dept. of Social and Decision Sciences
Carnegie Mellon University
Pittsburgh, PA 15213

email: kathleen.carley@andrew.cmu.edu

tel: 1-412-268-3225

fax: 1-412-268-6938

Citation: Kathleen M. Carley & Vanessa Hill, 2001, "Structural Change and Learning Within Organizations". In *Dynamics of organizational societies : Models, theories and methods*. Edited by Alessandro Lomi, MIT Press/AAAI Press/Live Oak.

Conventional wisdom often tells us that when it comes to “getting things done” it is “who” you know rather than “what” you know that matters. Contrarily, this same wisdom advises that “knowledge is power”. The conflict between these perspectives is reflected in research that separates these two factors, focusing on how structure influences behavior devoid of personal attributes, or on attributes excluding the relationships that contextualize the individual and their attributes. This research has shaped our thinking about behaviors in organizations, often focusing on either the relationship or the attribute unable to reconcile the interrelationship between these two factors. Today, however, a new view of organizations as inherently complex, computational and adaptive systems is emerging. This new perspective urges a reconsideration of the roles of relations and attributes. These two mechanisms are symbiotic, simultaneously impacting the behavior in organizations. Organizations are composed of intelligent adaptive agents constrained and enabled by their positions in networks linking agents and knowledge. Consequently, organizations are themselves synthetic agents in which knowledge and learning reside in the minds of the participant agents and in the connections among them.

This paper presents a knowledge level approach to organizational learning. This approach defines learning, culture, structure and behavior at the individual and organizational level in terms of both “who one knows” and “what one knows”. Relationships among individuals are important as they facilitate individual access to knowledge and serve as a form of organizational knowledge. Learning is conceptualized as the acquisition or loss of nodes and relationships, as with each node or relationship gained or lost knowledge is likewise acquired or forgotten. This is true whether the nodes are pieces of information or agents. This is true whether the relationships are among individuals, among information, or between individuals and information. Thus learning and memory exist at the individual and organizational level. When organizations, as synthetic agents learn important organizational behaviors emerge. Such behaviors reflect the emergent structuration of the organization’s culture and structure through learning at the individual and structural level.

Herein, this view is explicated and given precise form by defining the primary constructs in terms of the meta-network linking people and knowledge and the processes for evoking change in this meta-network. Essentially a knowledge level perspective is used to extend the traditional approach to social networks to include both people and ideas. Social networks are seen as affecting a wide range of behaviors ranging from power to consensus to adaptability. According to the common formulation such networks are in terms of ties among personnel. However, networks are more ubiquitous than this conception implies. Networks of ties link not just people, but people, knowledge, resources, tasks etc. (Krackhardt and Carley, 1998). We explicate this idea, and describe a computational model that incorporates this perspective. This computational model is

then used to examine the relation between structural stability and flexibility, performance, diffusion of new information, and consensus within organizations.

KNOWLEDGE LEVEL PERSPECTIVE

Increasingly organizational theorists are concerned with issues of organizational learning and cognition (Cohen and Sproull, 1996, Argote, 1999). While traditional learning studies pointed to changes in market share, productivity or performance as indicators of learning (Argote, Beckman, and Epple, 1990); more recent work looks deeper at the link between individuals, knowledge, and organizational outcomes at both the micro and macro level (Argote, 1996; Wegner, 1987, Argote, 1999).

Learning within organizations is ultimately tied to culture. This might seem an odd perspective given the ways in which culture is traditionally described. For example, the overall culture of a group is often operationalized as norms (Cooke and Rousseau, 1988, Rousseau, 1990), values (Posner, et al. 1985, O'Reilly et al, 1991 Chatman, 1991), rituals (Trice and Beyer, 1984, 1987), or the types of stories told (Martin et al, 1983 Gundry and Rousseau,1994), or using vague empathic terms describing the overall ambience such as "this place is very energetic" (Deal and Kennedy, 1982; Peters and Waterman, 1982). Each operationalization is an attempt to tangibly represent culture which is defined as "...pattern of basic assumptions that the group learned as it solved its problems of external adaptation and internal integration that has worked well enough to be considered valid and therefore, to be taught to new members as the correct way to perceive, think and feel in relation to those problems"(Schein,1983:13). In this sense, culture is the way in which the group responds to change in external and internal environments as well as a framework that guides the way individuals relate to each other. Thus, relationships as they provide the context for communicating basic assumptions play an integral role in culture. "The who" is an important aspect of learning (Argote, 1999, ch. 4). "Who" one is connected to influences the communication and adoption of culture among organizational members (Hill, 1999).

While relationships provide the mechanism by which culture is communicated and adopted, the content of culture is the result of *what* individuals know. The content of culture, the pattern of basic assumptions which individuals use as a framework to interpret events and subsequently guide behavior motivates the consideration of "the what" in the learning process. Knowledge exists within and between individuals, and thus within and between any group that contains individuals. As individuals learn, they alter the distribution of information, i.e., the cognitive content, the group's ability to learn, and culture. In essence then, what we want to suggest, is that when culture is viewed from a knowledge level perspective, all of these other characterizations of culture, norms, values, stories, goals, and ambience, are artifacts that emerge from the changing pattern of knowledge and interaction in the organization. Indeed, all social, cultural and individual behavior emerges out of the

ongoing interactions among intelligent adaptive agents (Carley, forthcoming; Padgett, 1997; Epstein and Axtell, 1997; Kauffman, 1995).

Culture as a pattern of knowledge and interaction is itself a form of distributed cognition (Hutchins, 1991, 1995) and so determines general behavior as well as changes to specific responses. Consequently, dramatically different behavior (at both the individual and the organizational level) can result from seemingly innocuous changes in the underlying social knowledge networks, and so changes in the underlying structure and culture of the organization. Such changes are ubiquitous. However, continual change does not imply that we cannot predict the behavior of the organizational system. If we are to understand and predict organizational behavior then we will need to understand “the who”- social relationships – and “the what”- knowledge – which result in learning. If we are to understand and predict organizational behavior then we will need to understand the socio-cognitive mechanics which bring about the observed change in the meta-network linking the who and the what (Carley, forthcoming). The knowledge level approach considers the who and the what in conceptualizing learning at the individual and organizational levels.

Role of “The Who” in the Knowledge Level Approach

Learning within organizations is ultimately tied to structure. A common conception of structure is in terms of the set of linkages among personnel; e.g., the authority network, the communication network, and advice network, the friendship network are all part and parcel of the structure of the organization. This conception of organizational structure is based on recognizing that the individuals in the organization are not independent. Social networks, the connections among these individuals, influence individual and group behavior (McPherson, 1983) and serve to constrain and facilitate change (Granovetter, 1985). Thus individual agency emerges from, is constrained by, and is enabled by this structure.

Any agent that can reposition itself in this interaction-knowledge space has agency. This view of agency draws from the familiar information processing approach (March and Simon, 1958; Simon, 1944; Galbraith, 1973). However, it extends that notion by incorporating it in a network framework (see also discussion in Carley, forthcoming). This provides grounding for talking about the information that the agents have as including not just the “what”, but also their perceptions of who knows who knows who (the cognitive social structure – Krackhraad, 1987) and who knows who knows what (the transactive memory Wegner, 1987, 1995; Moreland et al., 1996; Moreland, in press; or cognitive knowledge structure Monge and Contractor, forthcoming). Action results from opportunity, and the inter-relation among knowledge and capability. An aspect of capability is the passiveness or activeness with which the agent seeks new information.

Whom individuals interact with defines and is defined by their position in the social network. Therefore, in order to understand structural learning, it is particularly important to incorporate a knowledge level approach into our conception of networks within organizations. That is, let us

extend the network conception to include two types of nodes “the who” and “the what”. For the sake of simplicity, we can think of the “who” as the organizational member or more formally, as the information processing agent. The who is capable of knowing at least some of “the what” and is capable of taking action. Such actions might include storing, retrieving, manipulating, combining, creating, communicating the information or taking other actions based on the information known. Again, for the sake of simplicity we can think of the what as knowledge. More formally, the what is the information¹ .

There are linkages at all levels. Such linkages and the resulting networks and some of their characteristics are summarized in Table 1. Collectively, personnel, knowledge and the links among them form a single “meta-network”. The characteristics for each sub-network that are identified are the basic ties, the general phenomenon that this network represents, the core processual outcome that results when patterns in this structure become stable, and the type of learning. Organizational theorists are very familiar with the network by linking people to people- the social network. The communication network, the authority network, and the friendship network, are popular examples of variations on this theme. The point here, is that all such networks link people to people. The relationship to other people provides access and exposure to knowledge, which in turn impacts the individual who then updates his or her knowledge with that absorbed during interaction with another person.

Table 1. The Meta- Network formed of Sub- networks among people and knowledge.		
	People	Knowledge
People Tie Phenomenon Patterning process Learning	Social Network <i>Who knows who</i> Structure Structuration Structural learning	Knowledge Network <i>Who knows what</i> Culture Socialization, acculturation Individual learning
Knowledge Tie Phenomenon Patterning process Learning		Information Network <i>What informs what</i> Intellectual corpus Erudition Discovery

¹ The term information is being used in its broadest sense to include that anything that is codable, storable, etc. It can include beliefs, values, goals, graphs, pictures, words, numbers, information on the task, on what other’s know and whom they interact with, on one’s own and other’s demographic characteristics, and so forth. Since knowledge is hierarchical there will be macro indicators that are basically indicative of a wealth of background information that the individual is more or less privy to. For example, associated with demographic characteristics, degree titles, job titles, etc. is a set of knowledge that individuals with those characteristics are more likely to know than not. Thus, at the macro level the heterogeneity in these characteristics is likely to be correlated with overall heterogeneity in knowledge.

As previously noted, learning is ubiquitous. From a socio-knowledge network perspective, learning results in the construction of nodes and relations. Learning can, and does, occur in any of the sub-networks. Learning can, and does occur at both the individual and the structural level (Carley and Svoboda, 1996). Individual learning occurs within the individual through change in information nodes and links among them. Individual learning, including learning by being told and learning from experience, is reflected in changes in who knows what, i.e., in changes in the knowledge network. Experiential learning results in the individual changing his or her mental model by adding or dropping information and the relationships among those pieces of information that enable the individual to recall and use certain pieces of knowledge. These knowledge-level changes may precipitate changes in interaction if they alter the individual's relative similarity to other individuals.

Structural learning occurs when changes occur in the social network. Structural learning results in the adding or dropping of agents (individual representatives of the organization or the organization as a single entity) or choosing to continue or discontinue relations with those agents. An interesting aspect of structural learning is that it is often based on expectations about the future, and not just on direct historical experience (Carley and Svoboda, 1996; Carley, 1998; Carley and Lee, 1998). For example, an organization may choose to establish a relationship (selling, vending, acquisition, etc.) based on the forecasted profitability of the firm in question. These changes result in changes in the social network. Consequently, structural learning can make the organization appear adaptive, even though there is no intent of being adaptive.

A third time of learning is discovery. Discovery entails the creation or elimination of links among pieces of information for the first time. After discovery, the information is knowable.

This view of organizations and the consequent view of structural learning, result in a recognition that such learning induces structuration. Whereas individual learning, given that the individuals are constrained and enabled by their position in the social network, induces socialization and acculturation. Discovery leads to erudition as innovations illuminate previous unsuspected connections among information or defines away other relations. This conceptualization of learning makes it possible to distinguish individual versus organizational level learning. Within organizations, knowledge resides in the minds of the individuals in the organization, and it is also captured and stored in databases, procedural routines and organizational structure. The process of acquiring this knowledge, or learning, can be conceptualized, at least in part, as a structural phenomenon.

Role of “the what” in the Knowledge Level Approach

Learning on either the individual or organizational level describes a process by which knowledge is created, adopted, or dismissed. Through social networks, knowledge is acquired, thus

motivating the importance of “the who” in conceptualizing learning in organizations. Nevertheless, the discussion of the knowledge level approach is incomplete without an analysis of the role of “the what”. The knowledge network is the set of linkages between individuals and information, between “the who” and “the what”. Logically, there is another network of importance - the information network. The information network is the set of linkages among information. This information network has been described and analyzed in many fields and an abundance of approaches exist for conceptualizing this network including mental models, frames, schemas and many others (see Carley, 1997 for a discussion). The important point here, is that even as the “whos” can be related, so can the “whats”.

The culture of the organization can be defined in terms of the knowledge network. For example, Carley (1991) defined culture as the distribution of knowledge across people in the population. Culture is popularly defined as a pattern of basic assumptions that provide a framework for organizational members to interpret events and subsequently motivate appropriate responses to these events. New organizational members are taught to refer to this framework as it guides behavior. Inherently, culture has been traditionally conceptualized as a mechanism that engenders consensus, homogenizing the organizational members’ view of the world. However, organizational scholars have pointed out that culture makes differences salient (Martin, 1992, Van Maanen, 1991) so that consensus is experienced only in the context of similarly minded individuals in subgroups. At one extreme, scholars argue that culture is experienced idiosyncratically, thus different experiences result in unique interpretations (Martin, 1992) and consensus is fleeting and issue-specific. Despite disagreement on the extent to which culture is shared, all three perspectives can be seen as characterizing culture as knowledge that is shared, debated, negotiated, or distributed across organizational members. This, common theme is at the heart of the knowledge level conceptualization of culture.

Every individual enters the organization with experiences, beliefs, and information that comprises knowledge that is unique to that individual. This knowledge affects how the individual processes new information and how the individual chooses to develop, avoid, or discontinue relationships with other organizational members. In this sense “the what”, knowledge, influences the learning process by motivating the individual to add or drop particular nodes in his or her network, thereby changing the network of ties in the social network. This process of adding or dropping particular nodes and relationships from the network results in access to knowledge (experiences, beliefs, basic assumptions etc.) possessed by other individuals in their networks. As knowledge diffuses through this network culture is created and maintained. Culture is thus a product of knowledge diffusion. Unlike knowledge, culture is a distributed object. As knowledge is communicated in the context of social relationships among organizational members culture is created and maintained. As culture is a product of knowledge, as knowledge changes with the

addition or dropping of nodes and relations, culture changes. As such the nature of culture is dynamic and constantly in the process of being negotiated. Consensus observed in an organization's culture is simply a tentative answer to a problem that can be reaffirmed or renegotiated as organizational members see fit.

A Caveat on Learning

Changes in the networks, either through adding or dropping nodes or adding or dropping linkages is learning. In this paper we are exclusively concerned with learning vis change in the linkages, thus other types of learning are not described. From a cognitive perspective there are additional ways of discriminating types of learning. Many of these alternate conceptualizations have to do with the construction of the information network. Herein, we take the information network as a given and do not concern ourselves with the discovery or other processes that lead to the creation of bits of knowledge and connections among them.

Conceptualizing the organization in terms of this meta-network makes it clear that change in any of the sub-networks may ultimately affect all other sub-networks and the behavior of the entire system. Additionally, this meta-network effects the rate of information diffusion among individuals and within organizations, the ability of individuals to acquire and use information, and the speed, quality, and accuracy of organizational decisions. Consequently, sub-network change can cascade to effect dramatic organizational consequences.

The rate of change of these networks is, in part, a function of the way in which individuals use these networks. For example, sometimes people actively seek out information and at other times people learn new information passively, simply because they happen to be interacting with someone who tells them something new. The prevalence of active to passive information seeking behavior is likely to effect organizational behavior. In particular it should alter the rate of information diffusion and who gets what information when. This in turn might alter the degree of consensus and the performance of the group.

Although the social network and the knowledge network co-evolve they are different sub-systems (Carley, 1995;1999b). Consequently, they may change at somewhat different rates or reach stability at different times. The pattern of behavior observed within the organization will be a function of the relative rate of change of these networks and where they are at in the patterning process. Change in the social network, the structuration of the organization, occurs as groups (such as triads) form and re-form. Change in the knowledge network, the socialization and acculturation of the people in the group affects the degree of agreement. The rate of information diffusion will be affected by changes in both networks. Change in the information network, will interact with changes in the other networks to affect overall organizational performance.

SIMULATING STRUCTURAL CHANGE AND ORGANIZATIONAL LEARNING

Organizations, as synthetic agents, are difficult to theorize about. One reason for this is that the principles of combination that generate organizational behavior are more complex than simple aggregation. The level of complexity, number of inter-linked factors, and non-linear relations is such that processing aids, such as simulation, are needed. Herein, we use a computational model as an aid in theory development. Essentially, we are using the model as a hypothesis generator (Carley, 1999a). That is, the theory is realized as a computational model and then using formal procedures hypotheses are deduced.

Computational analysis and modeling, including but not limited to simulation, is fundamentally reshaping the way we do science. Through judicious use of reasonable programming and modeling practices, validation analysis, experimental design, and computational analysis — ideas can be examined, propositions compared, and theories can be developed and honed. Computational models can embody the fundamental nature of human information processing behavior (Simon 1981). Hence, computational models provide a logical and reasonable basis for reasoning about humans and organizations. Compared with human experiments and field studies, computational models are easier to control, more flexible, more replicable, more objective, and less prone to errors due to random noise. There are a large number of arguments about when computational analysis is useful and why (see for example, Levinthal 1999 this volume). Ostrom (1988) argues that computer simulation is a symbol system which “offers a substantial advantage” to researchers “attempting to develop formal theories of complex and interdependent social phenomena.” Two of these reasons are particularly relevant to the study of organizations in general and organizational learning in particular — scale and adaptation.

For example, consider the issue of scale. Computational analysis makes it possible to economically, and in a reasonable amount of time, examine more groups, larger groups, more tasks, larger tasks, more factors, more interactions, more dynamics (learning, adaptation, evolution), and harsher conditions than is practical, possible, or ethical to consider in either laboratory or field settings. This is particularly critical for studies of organizations. Two of the most common modes of data collection in organizational science have been laboratory experiments and case studies. Both types of studies bring to the fore issues of generalizability. To wit, it is not clear under what conditions laboratory experiments on small groups generalize to the larger organizational level nor is it clear to what extent single in-depth case studies on a particular firm generalize to the larger population of organizations. Since computational analysis obviates the problem of scale it provides a means by which the organizational researcher can ask: do the lessons learned in the experimental or field setting generalize.

As another example, consider the issue of complex adaptive systems. There is little doubt that organizations are complex adaptive systems. Organizational behavior is contingent on a huge number of factors which interact in complex and non-linear ways. Organizations are dynamic in a

variety of ways - they learn, adapt and evolve. Human beings, because they are cognitively limited (Simon, 1955; Carley and Newell, 1994) have difficulty simultaneously reasoning in multi-dimensional space and about a large number of factors, particularly if those factors interact in complex non-linear ways and or in a time varying fashion. In other words, human beings are not good at reasoning about complex adaptive systems. Computational analysis is a tool perfectly suited to assist in reasoning about complexity and/or dynamic systems at all levels - cognitive, ecological, historical. Computational agents have the ability to learn, adapt, and evolve and they can be constructed in ways that meet the Social Turing test (Carley and Newell, 1994) for matching complex adaptive systems to social behavior. Moreover, organizations and computer models are similar in kind. That is, they can both be complex adaptive systems. Moreover, organizations are inherently computational — they have a need to scan and observe their environment, store facts and programs, communicate among members and with their environment, and transform information by human or automated decision making (Carley and Gasser, 1999). As noted by Jim March in the introduction "Simulation represents an approach that appears both to match the phenomena of interest and to provide some analytical power."

In this paper, we develop our theoretical understanding of the impact of organizational dynamics on behavior and the evolution of the organization's culture and structure using a computational model. This model incorporates the previously described knowledge level extension of the standard network perspective. We examine the consequences of organizational change and learning for the stability and performance of small organizational units. This is done by conducting a series of virtual experiments using this computational model to explore the relative consequences of organizational learning on organizational performance and stability. The specific model used is CONSTRUCT-O. CONSTRUCT-O is a simulation engine in which the organizations are modeled in terms of agents, their knowledge, and their interactions and where performance is task based. Consequently, the models and theory embodied in the engine draws on recent work in cognitive science, social networks, sociology and organization theory. Aspects of the simulation engine have been validated (Carley, 1990, 1991, Carley and Krackhardt, 1996).

CONSTRUCT-O is based on the constructal model. The original constructal model, CONSTRUCT (Carley, 1990, 1991, 1995; Kaufer and Carley, 1993) is a computational model of social and individual change in response to the diffusion of information among individuals as they interact, communicate and learn. CONSTRUCT is essentially a multi-agent model in which each agent has a position in the social network, a mental model or knowledge (including knowledge about what, possibly defective knowledge about who knows what and who knows who). When individuals interact they communicate a piece of their knowledge (chosen at random). Individuals learn as they interact (through communication) and thus changing who knows what. Information diffusion can be measured for any specific piece of information as the number of people who

knows that piece at that point in time. Whom an individual interacts with is a function of their relative similarity in knowledge. Thus changes in who knows what ultimately lead to changes in who knows who and vice-versa. CONSTRUCT has been used to look at changes in workgroups, friendship networks, communication networks, and the impact of a group or organization's social structure or culture on the diffusion of information and the production of consensus. CONSTRUCT-O extends CONSTRUCT in three important ways – new interaction style, task analysis and triad analysis.

Interaction Style

In the original CONSTRUCT, the likelihood that one individual was to attempt to interact with another was purely of function of homophily in knowledge. Assume a matrix K such that if individual i knows fact k then $K_{ik}=1$ else $K_{ik}=0$; then, the likelihood of interaction is defined as:

$$\text{Passive Interaction}_{ij} = (\sum_{k=1} K_{ik} * K_{jk}) / (\sum_{j=1} \sum_{k=1} K_{ik} * K_{jk})$$

In CONSTRUCT-O the option was added that interaction could also be determined by having the goal of getting new information. In that case, the likelihood that the individual chooses another to interact with is a function of the relative potential of learning something new from the other. Let $NEW_{ijk}=1$ be the event that individual i does not know the information k , but j does; i.e.,

$$\text{IF } K_{ik}=0 \text{ \& } K_{jk}=1 \text{ THEN } NEW_{ijk}=1.$$

In this case, the likelihood of interaction is defined as:

$$\text{Active Interaction}_{ij} = (\sum_{k=1} NEW_{ijk}) / (\sum_{j=1} \sum_{k=1} NEW_{ijk})$$

Mixed strategies are also allowed for by proportionally averaging these; e.g., 25% seeking behavior occurs by 75% of the time the individual choosing an interaction partner based on the passive strategy and 25% of the time choosing an interaction partner based on the active strategy.

Regardless of the interaction style, when two individuals interact they each communicate a piece of information to the other, and each will learn the other's information if it is new to them.

Task Model

In the original CONSTRUCT the individuals interact and communicate. CONSTRUCT-O sets these individuals more explicitly in an organization, and assigns each agent a role and a task. Essentially, CONSTRUCT-O combines the interaction, communication, individual learning model from CONSTRUCT with the binary-choice task model used in numerous organizational performance models such as ELM (Carley, 1992) and ORGAHEAD (Carley and Svoboda, 1996). The binary choice model is based on the classic classification-choice task. In this task there is a sequence of problems, such that a problem is to decide for a binary string whether there are more 1's or 0's in the string. The task is distributed such that no individual, at least initially, can see the entire string. Strings are chosen with replacement from the set of all possible strings. Performance is calculated as organizational accuracy, the fraction of tasks for which the organization correctly

classifies the task. This task is added to CONSTRUCT-O as follows, each time period, agents interact, communicate, learn, solve problems, and then organizational accuracy is calculated. For the analyses herein, the size of the problem is the same as the amount of knowledge in the society. The number of problems looked at each time period is 25.

What parts of the task an individual sees depends on what pieces of information the individual knows. For example, if the individual knows information bits 3,5 and 7 then the individual can see the values of the task string for bits 3,5 and 7. Since individuals learn, if there are no barriers to learning eventually all individuals will know all bits of information and so will be able to see all parts of the task. Each individual looks at those task bits about which they have information, then the individual acts like a majority classifier. That is, each individual looks at those task bits about which they have information and then if in that subset of the task there are more 1s than 0s (or equal) the person responds with a 1, else the person responds with a 0.

For the organization it is assumed that all individuals opinion about the value of the problem are equally weighted. This is as though we are dealing with a collegial team where decisions are made by voting. Thus the set of individual opinions are averaged across the group, and if more individuals thought the answer was a 1 then that is the organization's decision, else it is a 0. It is this aggregated decision that is compared with the true answer to determine organizational performance. Both organizational performance and group consensus is a function of individual knowledge. Consensus is defined as the size of the group expressing the majority opinion (regardless of whether the opinion is 1 or 0, accurate or inaccurate).

Triads

Discussions of organizational flexibility often center on the ability to respond rapidly to the changing environment. Such notions of flexibility often center on the continual construction and reconstruction of groups. One of the most fundamental types of group is the triad. Change in the number, type, and location of triads is an indicator of the change in the number, type, and membership of groups. Rapid formation and reformation of triads is one key aspect of flexibility.

We measure the number of triads that exist at any point in time as the number of sets of three individuals (a,b,and c) such that for each dyad in the triad at least one of the two exhibits a stronger than average level of interaction likelihood to the other. This is a very weak measure as it does not require the interaction tie to be reciprocated. Nor does it require transitivity in the triad. A tie is said to be stronger than the average if $\text{Interaction}_{ij} \geq (\sum_{i=1} \sum_{j=1, i \neq j} \text{Interaction}_{ij}) / (N * (N-1))$ where N is the number of individuals in the population. Based on this definition of triads, the number of triads at each period is calculated both for the entire organization, and also for the sub-groups.

The Virtual Experiment

Given CONSTRUCT-O a virtual experiment was run. A virtual experiment is an experiment run using a computational model where selected parameters are varied over specific ranges. Such virtual experiments, like human laboratory experiments, are designed to address a core set of questions. Virtual experiments are typically done when the computational model is sufficiently complex or time consuming to operate that the entire response surface cannot be calculated. In this case, since behavior is a function of population size (the number of who's), knowledge base size (the number of what's), and the ratio of the two, all three of which can conceivably vary between 0 and ∞ , the entire surface cannot be calculated. Thus, a virtual experiment is called for.

In this virtual experiment the following factors were varied, size (and knowledge), the style of interaction, and whether or not the two divisions in the organization were differentiated or not in terms of size. This is summarized in table 2. Each organization was simulated for 300 time periods. This was sufficient to ensure quiescence in behavior. Each organization was simulated 4 times. However results reported here are for a single run, as this clarifies exposition. Since size and knowledge are locked together, this is a 3 x 5 x 2 design; i.e., 30 different types of organizations were simulated.

Four variables are tracked over time. Diffusion: the number of people who know a the new innovative piece of information at that time. Consensus: the size of the majority group when the decision is made. Performance: the accuracy of the organization. Triads: the number of triads. Based on this data a number of other items can be calculated such as the time to stability in each of these measures, the functional form of the change in these variables, and so forth.

Table 2. Virtual Experiment

Variables	Values	Meaning
Size	Small = 10, Medium = 20, Large = 30	Number of individuals in the organization. (The Number of Who's.)
Knowledge	2xSize	Number of bits of information that can be learned. (The Number of What's.)
Interaction Style	Passive = 0.00, 0.25, Mixed = 0.50, 0.75, Large = 1.00	The fraction of interaction likelihood that is focused on active as opposed to passive information reception.
Structure	Unsegmented = Groupsizes=Size/2, Differentiated = Groupsizes=Size/5 & Size-(Size/5)	The number of individuals in each of the two sub-groups in the organization.

Trends and Emergent Patterns

As individuals interact and learn they eventually come to know everything (assuming there are no structural barriers such as one sub-group with specialized knowledge is completely disconnected from the rest of the organizational members). Hence, over time, if a new idea is discovered, the number of individuals who learn it will approach the size of the organization. Eventually all individuals will come to agree. Thus the size of the group in agreement will approach, over time, the size of the organization. Over time, the level of organizational performance will approach 100%. And over time, interaction will equalize and so all possible triads will exist. Thus, long term patterns, i.e., the equilibrium state, is not at issue. The important factor is how fast the organization reaches this state and what happens along the way.

Thus, we focus on the time at which stability is reached.² When the all simulations are considered, a few general trends stand out. First, regardless of the measure, things take longer the larger the organization's size – see Table 3. Second, the organization's underlying structure has a greater influence on individual behavior than organizational behavior. That is, in a differentiated organization where divisions differ in size then information diffuses more slowly and the triadic structure takes longer to stabilize. However, there are not appreciable average differences in organizational level outcomes such as time to stability in consensus and performance. All else being equal individuals in organizations with a strong minority/majority sub-structure may feel less in the know, more left out, and be more concerned with cliques and coalitions than their counterparts in other organizations.

Size	10	20	30
Diffusion	41.60	73.20	97.70
Performance	61.10	124.60	138.20
Consensus	63.40	129.10	147.10
Triadic	69.00	156.50	248.40

As to the third independent variable, interaction style, the pattern varies. Again looking only at averages we see that as individuals become more information seeking (active) it takes longer for the organization to reach stability in terms of either consensus or performance. More complex curvilinear patterns appear at the individual level for both diffusion and triads. Diffusion is particularly interesting. As individuals become more active in seeking out new information the time

² In non-linear stochastic systems where the data is being sampled, such as is the case here, measuring the exact time at which stability is reached is difficult if not impossible due to measurement noise. A standard measure is the 90% point; i.e., the time at which the measure gets to 90% of its final (theoretical) value.

it takes for a novel idea to diffuse decreases. This is not particularly surprising. What is surprising is that if individuals never interact passively, if they actively seek new information, it actually takes longer for the new idea to diffuse. Essentially, if only a few gatekeepers hold the knowledge that is being sought, then they may be too busy interacting with certain others to interact with everyone thus creating a bottleneck in the flow of information. Further, individuals will spend most of their time not interacting because they think no one can tell them what they want, so they are less likely to information by chance than if they were passive receptors.

As a final point, in general stability occurs in diffusion first, performance second, consensus third, and triads last. What this means is that in organizations, even after peak operating performance is reached, there may still be lots of internal activity – people trying to convince each other of their opinions, people changing opinions, coalitions shifting, group structures altering, all with no appreciable effect on performance. It takes longer for consensus to stabilize than it does for information to diffuse because consensus involves the diffusion of multiple pieces of information. Stability in consensus occurs when the distribution of information and the set of problems is such that agreement can be reached even though different individuals have different information. The issue is not how much information is shared, but do individuals regardless of their knowledge reach the same conclusion. Consensus can, and typically does, occur long before everyone knows everything. When information diffuses it can make two people more likely to agree or disagree depending on whether that information makes the overall pattern of information on the task that they see more similar or less similar. These results characterize the typical behavior, with differences in size, structural form, and interaction style averaged out. We now engage in a more detailed analysis where all of these variables are considered simultaneously.

Interaction Style	Passive		Mixed		Active
	0.00	0.25	0.50	0.75	1.00
Diffusion	73.33	73.17	64.17	56.50	90.00
Performance	98.83	106.83	107.17	112.33	114.67
Consensus	103.17	110.17	114.67	115.83	122.17
Triadic	111.67	171.83	165.67	170.17	170.57

COMPETING PROCESSES

Time to triadic stability increases monotonically in size when either individuals search for information actively or when they act as passive receptors of information. Whether individuals strive to meet a goal (acquire new information) or respond by interacting with those with whom they

are most comfortable (passive homophily based interaction) the result is that the larger the organization the longer it takes the structure to stabilize. The effect is at least linear for a passive interaction style and may be exponential in size for an interaction style. Now, the typical assumption is that when two processes, both of which lead to monotonic increases that are at least linear, when they are combined together, the result will also lead to increases that are at least linear. Thus, if individuals spend part of their time passively and part of their time actively, one would expect that as the size of the organization grows the time to triadic stability would grow at least as fast as the slower process. However, this is not the case. As can be seen in Figure 1, when there is some combination of information seeking behaviors the effect of size is mitigated. This type of behavior is typical of a non-linear system. But what that means to the theorist and policy analyst is that reasoning about the impacts of the co-presence of multiple processes on organizational outcomes is suspect to error. Computational models can aid in this regard by illuminating the non-intuitive effects when multiple processes are at work.

Time to Triadic Stability

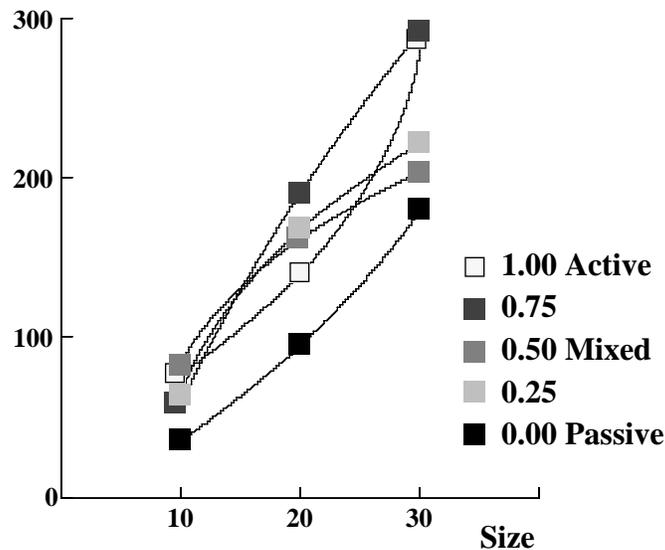


Figure 1. Impact of Interaction Style on Triadic Stability

The effects of these non-linearities in the processes which control learning, communication, interaction and so the evolution of the organization’s culture and structure are even more pronounced for other behaviors. Illustrations of this appear in Figures 2 and 3. With respect to information diffusion (Figure 2), the effect of size on time to diffusion changes radically with the style of interaction. From a passive information reception perspective the typical contagion argument is that information moves more slowly the larger the group, as there are more people who do not know the new idea. From an active information seeking perspective, the argument is that the

larger the group the more likely it is that someone will know the needed information. Thus suggesting that ideas might move faster in larger groups. We see here that neither argument holds, even under pure seeking and reception behaviors, there are still non-linearities with respect to size. Further, when individuals engage in both types of behaviors the effects can be just the opposite of what was observed for either of the pure forms of interaction.

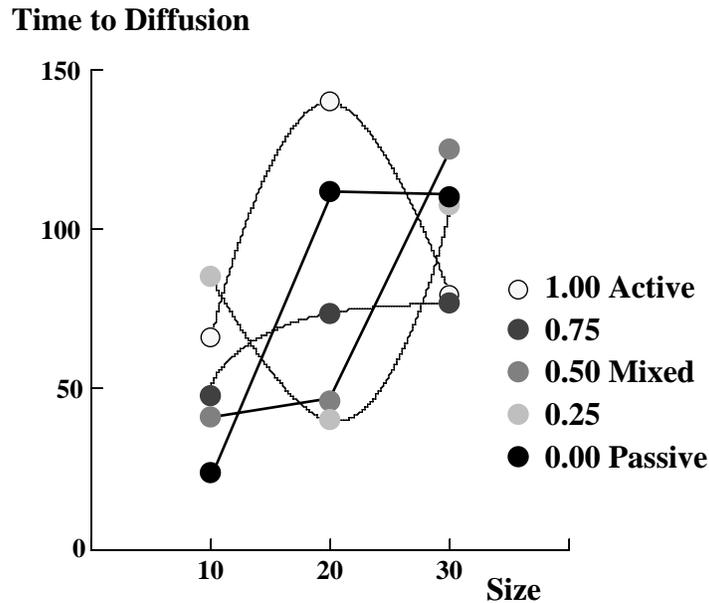


Figure2. Impact of Size and Interaction Style on Time to Diffusion for Differentiated Organizations

Now consider the case of consensus. In this case, as individuals learn new ideas, if what they know initially is very different from what others know, then regardless of how they learn the new idea, such learning is likely to change whom they agree with about what. In larger groups, if the size of the group corresponds to the amount there is to know, then regardless of the way in which individuals learn new information, the larger the group the longer it will take to reach a stable level of consensus. We see, in Figure 3, that this line of reasoning holds up except for two conditions - when a pure passive information reception strategy is pursued and when a mixed 50/50 strategy is pursued. When a purely passive search strategy is used, individuals as groups tend to get the same information. This leads to an alignment of social structure and culture, which then mitigates the effect of size on consensus by enabling large groups of individuals to agree even though they are far from having all the information. In other terms, a community of practice evolves in which those who interact talk the same talk and share the same information. It is precisely in this situation, where culture and structure align, where group-think is likely to prevail.

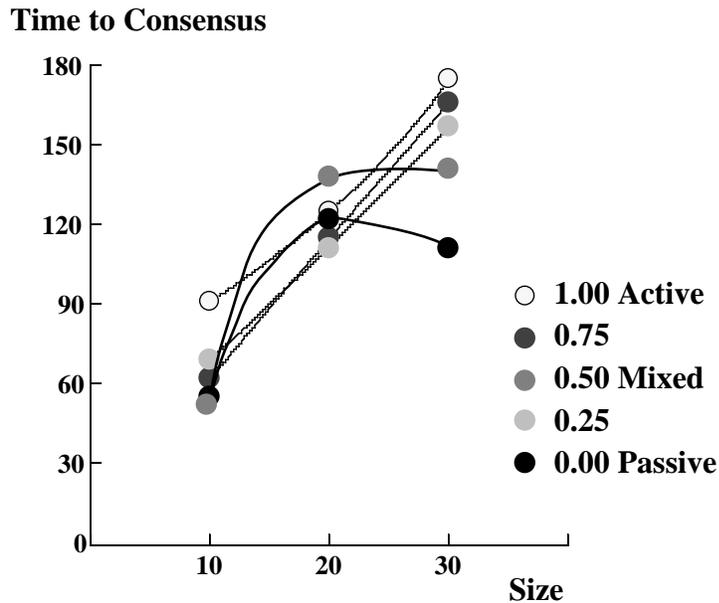


Figure3. Impact of Size and Interaction Style on Time to Consensus for Differentiated Organizations

The stabilization of consensus signals the stabilization of culture. When consensus is varying that means that the beliefs, opinions, attitudes and so on of individuals are changing. This only happens when individuals continue to get new information that contradicts earlier information. The stabilization of triads signals the stabilization of structure. When the number of triads is varying that means that who interacts with whom is changing. This only happens when individuals continue to change what they know, and so who they know. As was noted earlier, in general structure stabilizes after culture, triadic stability occurs after consensus stability. However, as can be seen in Figure 4, there are a few cases where triadic stability leads consensus stability. For example, for small and medium sized unsegmented organizations homophily based interaction enables the organization to reach structural stability (stability in triads) prior to culture stability (stability in consensus).

What does it mean for the organization to reach structural stability before stability in diffusion, consensus or performance? When diffusion reaches stability before triads that means that changes in what you know are affecting changes in who you know. Thus, experiential learning and information seeking are driving structural learning. In this case, individuals are shaping and reshaping their groups. In contrast, if the triads are stable in the face of the diffusion of innovative information that means that who you know is driving what you know and so structural learning is

driving experiential learning. In this case, the organization and existing groups are shaping the individuals within them. What the results in this paper are suggesting is that even when there is a cycle in which interaction leads to learning and learning to changes in interaction, over time different organizations will change such that one side of this process may come to dominate. These ideas are summarized in Table 5.

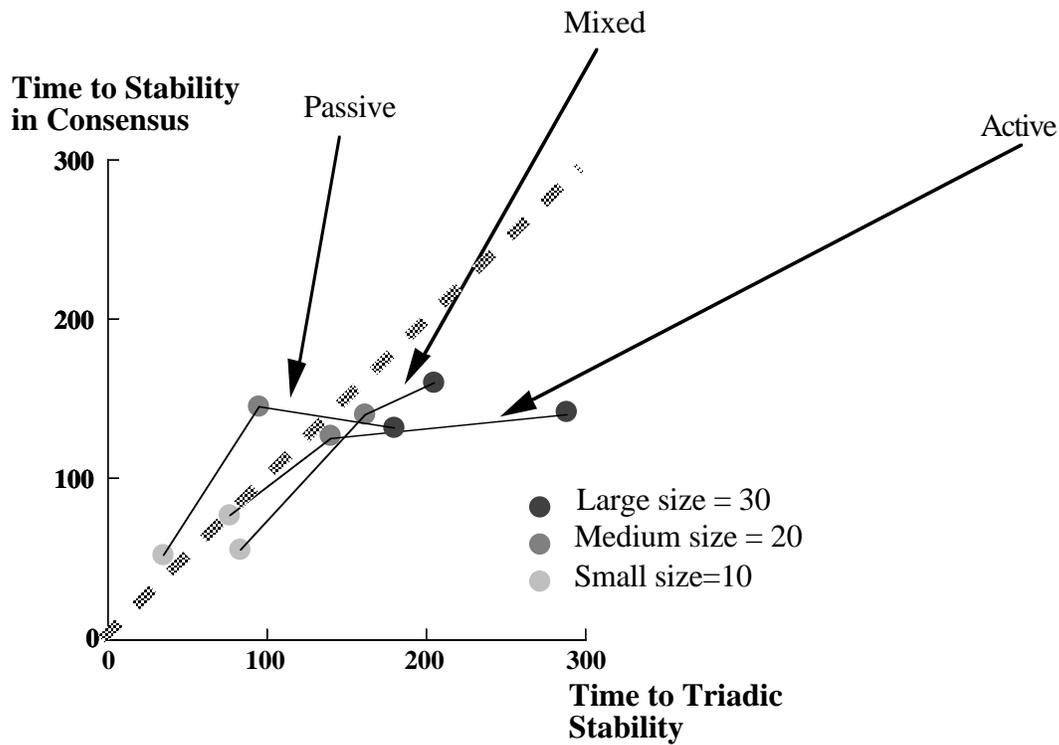


Figure 4. Relative Stabilization of Consensus and Triads in an Unsegmented Organization.

When consensus stabilizes before triads, the organization is faced with a situation where individuals are still learning new information, but the opinions, attitudes, beliefs and so on are not changing even though changes in interaction are occurring. In this case, the group is prone to group think, to cultural stagnation, to rituals, and to the unquestioning application of norms and procedures. The group may also be prone to false consensus. That is everyone may be agreeing on the outcome or decision but for different reasons and so expecting different long run implications. In this study, consensus rarely stabilized before structure. Although it was more common in small and medium sized groups. When triads stabilize first, then the group should fall into stable patterns of behavior in which the pattern of group interaction drives, reflects, and defines

decisions and consensus. How this is done depends, at least in part, on the number of triads. For example, if there are only a few triads that are very stable, then group is responding in a fractionated fashion and consensus is likely to be only consensus among the controlling cliques. When the structure stabilized with a large number of triads, then essentially the group is homogenous as the webs of affiliation create an integrated group. In this case, consensus is likely to be a true consensus of the group, and not forced by a coalition. These arguments are summarized in Table 6.

Table 5: Diffusion and Triadic Stability		
	Diffusion	
Triads	Unstable	Stable
Unstable	Group in formation	Experiential learning drives structural learning. New ideas may be seen as disruptive.
Stable	Structure determines what individual learns. New ideas are used in support of existing structure.	Everyone knows the information. Communication of the information is ritual exchange of information and does not affect the group. Information may become a tacit unspoken part of the underlying culture.

Table 6: Consensus and Triadic Stability		
	Consensus	
Triads	Unstable	Stable
Unstable	Group in formation.	Structure comes to imitate culture. Group think.
Stable	Culture comes to imitate structure. Coalitions or integration may occur depending on number of triads.	Block voting.

Table 7: Performance and Triadic Stability		
	Performance	
Triads	Unstable	Stable
Unstable	Group in formation	Exploitation of competencies exploring new structures. Structural Learning Flexibility Organization may engage in pruning or selected downsizing.
Stable	Directed exploitation Experiential Learning	Mature Organizational Form

Performance stabilizes before structure if no learning is occurring or if people are still learning but such learning is not related to the task, is redundant vis the task. In this case, the only learning

that occurs is through structural learning. Thus, while performance may not be improving efficiency might well be. The organization can be thought of as engaged in a process of exploiting known competencies but exploring new structures. Since structural changes have little effect on performance the organization has the flexibility to explore alternate organizational forms. Further, the organization might see the structure as containing redundancies and might choose to downsize, figuring that this can be done with impunity. Performance improvements would be possible only through strategic changes. In contrast, if the structure stabilizes before performance then the structure will direct what people learn when and so lead to very directed changes in performance. The organization can be thought of as engaged in a type of directed exploitation. The organization in this case relies on the structure for a certain minimal level of performance, and individuals engage in experiential learning on the task. Individual learning enables further performance improvements. How triadic stability affects the way in which performance is reached again may depend on the number of triads. These implications are summarized in Table 7.

The Nature of Triadic Stability

Over time, diffusion, consensus and performance increase more or less, monotonically. In contrast, the number of triads oscillates wildly until it stabilizes (see Figure 5 for accuracy and Figure 6 for triads). Over time, triads behave like a damped second order non-linear system; i.e., there are large oscillations, the spacing between and degree of which changes with time. The number of triadic oscillations increases with the size of the organization and the time to stability. However, proportionally, there is little variation. In figure 6 the number of triads, as a proportion of the number of possible triads is shown. Initially there are only a few triads, however, they are very stable. There are a few cliques and it takes a lot of information exchange to break these apart. Over time, more triads emerge. Eventually, triads are ubiquitous; but, the longevity of any one triad is short. That is, individuals form and reform groups spontaneously and rapidly in response to changing information (and presumably changing tasks). Ultimately, the organization is home to a very flexible structure.

Another key to structural change is the height of the final peak. The larger the group, the longer it takes to reach structural stability and the greater the level of triadic integration before the final stage and the less severe the oscillations. Small groups are thus extremely structurally volatile, even in the absence of turnover, until they reach quiescence. In contrast, large groups exhibit less structural stability over time, even though it takes them longer to reach their final ultimate stability. In this sense, larger groups are more staid.

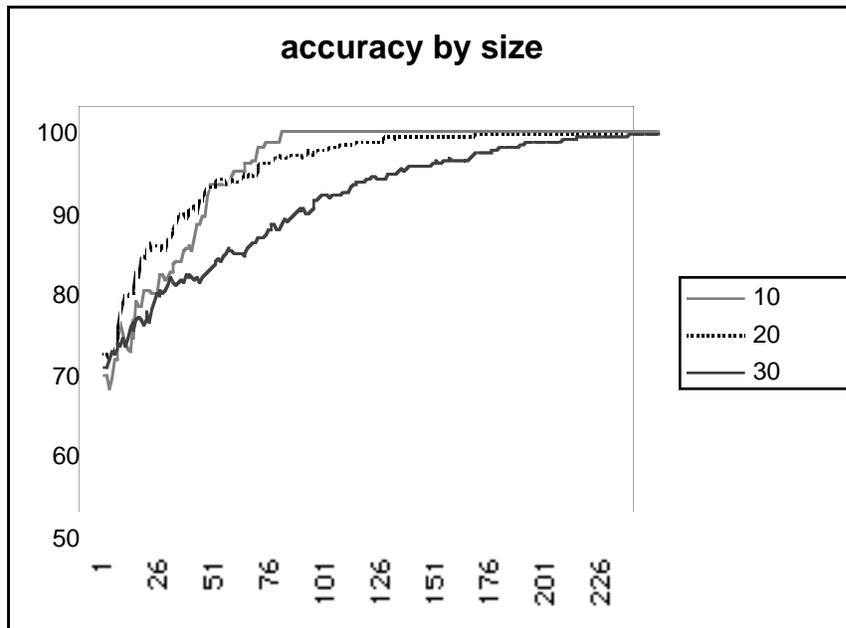


Figure 5. Illustration of rise in accuracy over time, by size of organization for an unsegmented organizational structure.

DISCUSSION

Computational modeling, like any formal modeling technique, enforces a certain level of rigor. Assumptions need to be laid bare and hypotheses are formally derived. Computational modeling also enables the researcher to address real world empirical phenomena. Many computational models generate results that on the surface appear to have vague general empirical similarity to actual phenomena. Such models cannot be used, in any sustained sense, to provide detailed guidance to researchers, managers or policy analysts. Carley and Prietula (1994) have argued that to go beyond such superficial similarity organizational models should develop ACTS models – models of organizations where the agents are cognitively bounded, task oriented and socially situated. CONSTRUCT-O is an ACTS model, and as such has the ability of being used to generate a large number of detailed and empirically verifiable hypothesis, well beyond those discussed herein.

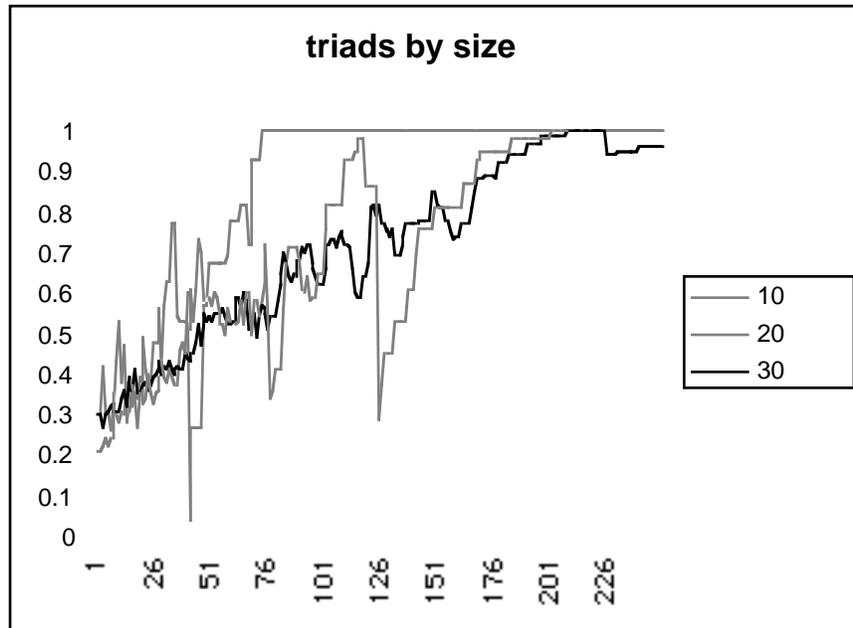


Figure 6. Illustration of change in the proportion of triads over time, by size of organization for an unsegmented organizational structure.

Using computational models for theory development and hypothesis generation is one of several ways in which computational models can and are being used. Alternative ways of using computational models include but are not limited to the following. Computational models can be used to tune laboratory experiments and field studies. They can be used to logically reconstruct events and provide greater insight by ruling out or in alternative explanations (see for example Carley, 1990). Computational models can be used to test and compare competing theories (e.g., Larsen and Lomi, 1999; Sastry, 1997; 1999, this volume; Sterman, 1999). They can be used for reasoning about current events by matching them to the patterns observed in the past (Baligh, Burton and Obel, 1990). Finally, these models can be used to provide a “history friendly” social lab for reasoning about social and organizational phenomena (Nelson, Winter, et al., 1999).

Agent Based Models

In this paper, an agent based model is used. Each agent is individually modeled, has its own knowledge, makes its own decisions, has its own unique history of interaction, and so on. Culture and social structure are viewed in terms of distributions across agents. This can be generalized to a more macro level, where the organization as a whole and not the individual component members are modeled. For example, we could capture culture as the distribution of knowledge as a functional form, the mean and standard deviation in much the same way as is done by Harrison and Carroll

(1999, this volume). In terms of long run behavior and meta-results, the two approaches are complimentary and will yield similar results. The distinction is in terms of short term behavior, the ability with an agent based approach to observe over time path dependencies, to examine the impact of individual differences, and to link to findings in cognitive psychology. Whereas, the equation based approach is, at times, simply due to computational power, more extensible to examining a larger number of organizations at one time and to linking to macro-level indicators and the findings of institutional theory. It is likely that in the future, multi-level models that include agent based components and non-agent based components will be used to develop social and organizational theory.

Implications for Network Analysis

Traditional network analyses explore organizational issues by focusing on networks composed of one type of node “N”. “N” usually represents people (Wasserman and Galaskiewicz,1994; Krackhardt, 1992) or firms (Mizruchi and Galaskiewicz,1994). Consequently, traditional analysis, which bounds the network by the type of node may result in erroneous or misleading conclusions about the role of learning in producing social change. If trying to understand whether or not employees in a company will be mobilized to participate in a strike (Kapferer, 1972) or to unionize (Krackhardt,1992; 1995) both agents and knowledge need to be considered. For Kapferer’s tailor shop, models that attend to both the social network and the knowledge network are better able to explain the changing potential to mobilize for a strike than are models that only account for the social network (Carley, 1990).

The organization as a synthetic agent, is an intelligent, adaptive and computational entity (Carley and Gasser, 1999). Consequently, it can engage in behavior that is distinct from an aggregation of the participant agents’ behaviors. The organization’s intelligence, adaptiveness, and computational capability results from the detailed, ongoing, interactions among and behavior of these participant agents as they move through the interaction-knowledge space. This means that a variety of measures can be calculated on such a meta-network (Carley and Butts, 1998). Such measures can be calculated on the meta-network in its entirety or on sub-networks. For example, the overall complexity of the organization as a system can be defined in terms of the number of linkages. The number of triads could be calculated from the social network and the degree of agreement from the knowledge network. Herein, we focus on measures based on the sub-networks. The analysis presented here focuses on four outcomes: groups, diffusion, consensus, and performance. These behaviors are important to examine as they capture socially and economically important ways in which organizations may be affected by change in the meta-network.

Over time, this meta-network evolves. These networks have a great deal of influence; e.g. they effect the rate of information diffusion among individuals and within organizations, the ability of individuals to acquire and use information, and the speed, quality, and accuracy of organizational

decisions. Consequently, the change or evolution of the meta-network can have dramatic organizational consequences. Most models of network evolution overlook the simple fact that the social network does not exist independent from the knowledge network. However it is a fundamental aspect of the human condition that the social network denoting who talks to whom is intertwined with the knowledge network denoting who can talk about what. These networks co-evolve as the individuals interact, communicate, and so learn new information (Carley, 1991). Herein we have seen that although the social network and the knowledge network co-evolve they can reach stability at different times, with formidable implications for group and organizational behavior.

SUMMARY AND THEORETICAL IMPLICATIONS

The story we have told is that organizations are essentially synthetic agents composed of intelligent adaptive agents constrained and enabled by their positions in networks linking agents and knowledge and by the tasks in which they are engaged. Many types of learning occur within the organization including individual experiential and structural. Interactions between structural and learning and experiential learning affect the stability of the organization at the structural and cultural level. Which process dominates depends on the history of interaction and communication within the organization. Since the basic processes are non-linear complex and unanticipated behaviors emerge. As a result, large organizations can behave in ways that appear qualitatively different than small; adding strategies for searching for information can radically alter the effects of size; and so on.

This story, this theoretical perspective, emerged through computational analysis. In addition, a number of specific results appear, such as moving from a group of size 20 to 30 when individuals act as passive receptors of information, doubles the length of time to triadic stability. These specific results should be viewed with caution. The story that the simulation enables us to tell is not in the specific values, but in the general trends and patterns of behavior. General trends, such as the increasing number of oscillations and decreasing degree of oscillation with the size of the organization, are the robust results.

Extrapolating from this study needs to be done with caution. A major limitation of this study is that the size of the organization and the size of the underlying knowledge base were keyed off each other and so perfectly correlated. Additional studies need to be done on the differential effects of increasing the number of people and increasing the amount there is to know. Earlier work suggests that these two dimensions should have very different effects on diffusion. That suggests that differential effects will emerge for consensus, performance and triadic stability as well. This study demonstrates that although structure is generally the last item to stabilize, it can stabilize while new ideas are still diffusing, while the group has not reached a stable culture of consensus, while

performance has not stabilized. The culture and behavior of groups and organizations is affected by the relative rates with which these factors (ideas, consensus, structure and performance) stabilize. These findings complement and build on those of Harrison and Carroll (1999, this volume). They find that failure to achieve a strong culture increases the likelihood of organizational mortality. For our model, Harrison and Carroll's strong culture is achieved when the organization stabilizes at the knowledge, the consensual, the structural level and the performance level. They note that failure to achieve a strong culture increases the chances of organizational mortality. Our results add that early stabilization of knowledge, consensus, structure or performance creates a set of vulnerabilities in how the organization responds to problems. Thus, organizational mortality may be as a function of poor management practices during these periods of vulnerability.

Understanding how the structure, processes, and knowledge within a group determines the relative stability of these factors is the first step in providing a formal and testable theory of culture creation and maintenance, and it will provide managers interested in altering culture with a concrete set of guidelines for altering those aspects they are interested in. This work is a step in this direction.

ACKNOWLEDGEMENT

The research reported herein was supported in part by the National Science Foundation NSF IRI9633 662 and in part by the Office of Naval Research (ONR), United States Navy Grant No. N00014-97-1-0037.

References

- Argote, L., Beckman, S., and Epple, D. 1990. "The persistence and transfer of learning in industrial settings," Management Science, 36, 140-154
- Argote, Linda, 1996, "Organizational learning curves: Persistence, transfer and turnover" International Journal of Technology Management, 11(7,8): 759-769.
- Argote, Linda, 1999. Organizational Learning: Creating, Retaining and Transferring Knowledge. Boston, MA: Kluwer.
- Baligh, H.H., R.M.Burton, and B.Obel. 1990. "Devising expert systems in organization theory: The organizational consultant," in Organization, Management, and Expert Systems, M. Masuch, (Editor), Walter De Gruyter, Berlin. pp. 35-57.
- Carley, Kathleen M. 1990. "Group Stability: A Socio-Cognitive Approach." Pp. 1-44 in Advances in Group Processes: Theory and Research Vol. VII , edited by E. Lawler, B. Markovsky, C. Ridgeway and H. Walker. Greenwich CN: JAI Press.
- Carley, Kathleen M. 1991. "A Theory of Group Stability." American Sociological Review 56(3): 331-354.
- Carley, Kathleen 1992, "Organizational Learning and Personnel Turnover." Organization Science, 3(1): 20-46.
- Carley, Kathleen 1995, "Communication Technologies and Their Effect on Cultural Homogeneity, Consensus, and the Diffusion of New Ideas." Sociological Perspectives , 38(4): 547-571.
- Carley, Kathleen M. 1997, "Extracting Team Mental Models Through Textual Analysis." Journal of Organizational Behavior, 18: 533-538.

- Carley, Kathleen M. 1998, "Organizational Adaptation." Annals of Operations Research. 75: 25-47.
- Carley, Kathleen M. 1999a. "On Generating Hypotheses Using Computer Simulations." Systems Engineering, 2(2): 69-77.
- Carley, Kathleen M. 1999b. "On the Evolution of Social and Organizational Networks." In Steven B. Andrews and David Knoke (Eds.) Vol. 16 special issue of Research in the Sociology of Organizations. on "Networks In and Around Organizations." JAI Press, Inc. Stamford, CT, pp. 3-30.
- Carley, Kathleen M., forthcoming, "Organizational Adaptation in Volatile Environments" Ch. 3 in C. L. Hulin & D. R. Ilgen (Eds.) *Computational Modeling in Organizational Behavior: The Third Scientific Discipline*. Washington, DC: American Psychological Association.
- 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," Pp. 55-87 in Kathleen M. Carley & Michael Prietula (Eds.), *Computational Organization Theory* .Hillsdale, NJ: Lawrence Earlbaum Associates.
- Carley, Kathleen M. and David Krackhardt, 1996, "Cognitive inconsistencies and non-symmetric friendship." Social Networks , 18: 1-27.
- 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 David M. Svoboda. 1996. "Modeling Organizational Adaptation as a Simulated Annealing Process." Sociological Methods and Research 25(1): 138-168.
- Carley, Kathleen M. and Les Gasser. 1999. "Computational Organization Theory." Ch. 7 in Distributed Artificial Intelligence edited by Gerhard Weiss. Cambridge, MA: MIT Press.
- Chatman, J. 1991. "Matching people and organizations: Selection and socialization in public accounting firms. Administrative Science Quarterly 36, 459-484.
- Cohen, Michael and Lee Sproull, 1996. Organizational Learning, Newbury Park, CA: Sage.
- Cooke, R and Rousseau, D.M. 1988 "Behavioral norms and expectations" Group and Organization Studies 13 (3) 245-273.
- Deal, T. and Kennedy A. 1982 Corporate Cultures: The rites and rituals of corporate life Reading, MA Addison-Wesley.
- Epstein, Joshua and Rob Axtell. 1997. Growing Artificial Societies. Boston, MA: MIT Press.
- Galbraith, Jay 1973. Designing Complex Organizations. Reading, MA: Addison-Wesley.
- Granovetter, Mark. 1985. "Economic Action and Social Structure: The Problem of Embeddedness," American Journal of Sociology 91: 481-510
- Gundry, L. and Rousseau, D. 1994. "Critical incidents in communicating culture to newcomers: The meaning is in the message" Human Relations 47 (9), 1063-1088.

- Hill, V. 1998. *An Examination of Culture Creation and Maintenance in Temporary Placement Agencies*" Unpublished doctoral dissertation, Carnegie Mellon University, Pittsburgh, PA.
- Hill, V. 1999 "Making friends before you need them: The influence of social relationships on the creation and maintenance of organizational culture" Unpublished manuscript, University of Arizona, Tucson, AZ.
- 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. 1995. Cognition in the Wild. Cambridge, MA: MIT Press.
- Kapferer, Bruce. 1972. Strategy and Transaction in an African Factory: African Workers and Indian Management in a Zambian Town. Manchester, England: Manchester University Press.
- Kathleen M. Carley and Carter Butts, 1997, "An Algorithmic Approach to the Comparison of Partially Labeled Graphs" in Proceedings of the 1997 International Symposium on Command and Control Research and Technology. June. Washington, DC.
- Kauffman, Stuart A. 1995. At home in the universe : the search for laws of self-organization and complexity. New York, NY: Oxford University Press.
- Krackhardt, David. 1987. "Cognitive Social Structures," *Social Networks* 9: 109-134.

- Krackhardt, David and Kathleen Carley. 1998. "A PECANS Model of Structure in Organizations." In Proceedings of the 1997 International Symposium on Command and Control Research and Technology. June. Monterey, CA.
- Krackhardt, David. 1992. "The Strength of Strong Ties: The Importance of Philos in Organizations." Pp. 216-239 in Networks and Organizations: Structure, Form, and Action, edited by Nitin Nohria and Robert Eccles. Boston, MA: Harvard Business School Press.
- Krackhardt, David. 1995. "Entrepreneurial Opportunities in an Entrepreneurial Firm: A Structural Approach," Entrepreneurship Theory and Practice 19: 53-69.
- Kaufers, 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.
- March, James G. and Herbert A. Simon. 1958. Organizations. New York, NY: Wiley.
- Martin, J. 1992 Cultures in organizations: Three perspectives Oxford University Press
- Martin, J. Feldman, M., Hatch, M. and Sitkin, S. 1983 "The uniqueness paradox in organizational stories" Administrative Science Quarterly 28 438-
- McPherson, J. Miller. 1983. "An Ecology of Affiliation," American Sociological Review 48:519-532.
- Mizruchi, Mark and Joseph Galaskiewicz. 1994. "Networks of organizations," Pp. 230-254 in Advances in Social Network Analysis, edited by S. Wasserman, S. and J. Galaskiewicz. Thousand Oaks, CA: Sage.

- Monge, Peter R. and Noshir S. Contractor. forthcoming. "Emergence of communication networks." In Handbook of Organizational Communication (second edition), edited by L. Putnam and F. Jablin. Newbury Park, CA: Sage.
- 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.
- O'Reilly, C, Chatman, J. and Caldwell, D. 1991. "People and organizational culture: A profile comparison approach to assessing person-organization fit" Academy of Management Journal 34, 487-516.
- Ostrom, T. M., 1988. "Computer Simulation: The Third Symbol System," Journal of Experimental Social Psychology, 24: 381-392.
- 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.
- Peters, T. and Waterman, H. 1982. In search of excellence: Lessons from America's best run companies Harper and Row, NY.
- Posner, B. , Kouzes, J., and Schmidt, W. 1985. "Shared values make a difference: An empirical test of corporate culture" Human Resource Management 24 (3) , 292-309
- Rousseau, D.M. 1990 "Normative Beliefs in Fund-Raising Organizations" Group and Organization Studies 15 (4), 448-460.

- Sastry, Anjali M. 1997. "Problems and paradoxes in a model of punctuated organizational change," *Administrative Science Quarterly* 42(2): 237-275.
- Schein, E. 1983 "The role of founder in organizational culture" *Organizational Dynamics*, 12, 5-23.
- Simon, Herbert A. 1944. "Decision-making and administrative organization," *Public Administration Review* 4: 16-31.
- Simon, Herbert A. (1955). "A behavioral model of rational choice". *Quarterly Journal of Economics*, 69: 99-118
- Simon, H.A., 1981 "Studying human intelligence by creating artificial intelligence," *American Scientist*, 69(3): 300-309.
- Trice H. and Beyer J. 1984 "Studying organizational cultures through rites and ceremonials" *Academy of Management Review* 9 (4), 653-669.
- Trice H. and Beyer J. 1987 "How an organization's rites reveal its culture" *Organizational Dynamics* 15, 4-25.
- Wasserman, Stanley and Joseph Galaskiewicz (Eds). 1994. *Advances in Social Network Analysis: Research in the Social and Behavioral Sciences*. Newbury Park, CA: Sage.
- 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.