

**ON THE EVOLUTION
OF SOCIAL AND ORGANIZATIONAL NETWORKS**

Kathleen M. Carley
Carnegie Mellon 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

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.

Forthcoming: In David Knoke & Steve Andrews (eds.) special issue of Research on the Sociology of Organizations on "Networks in & Around Organizations" JAI Press, Inc.

ABSTRACT

Over time, social and organizational networks evolve. 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 these networks can have dramatic social and organizational consequences. Most models of network evolution overlook the simple fact that networks exist with in 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) and the transactive knowledge network (each individual's perception of the network linking people to their ideas). For example, within organizations, the authority or reporting network (who reports to whom) is interlinked with many other networks including the task structure (which tasks are connected to which), the task access structure (who is assigned to what task). Change in any part of this ecology of networks ultimately affects all other parts and the behavior of the entire system is a function of the specific way in which these networks are inter-linked.

Organizations, and indeed any group or society, is continually in a state of flux. This flux often takes the form of changes in the underlying social and organizational networks. Such change can be, and has been, characterized as network evolution. 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 these networks can have dramatic social and organizational consequences. Continual change does not imply unpredictability. If we are to understand, and possibly even predict, the behavior of these systems, then we will need to understand the socio-cognitive mechanics which bring about the observed change. We will need to understand the mechanics by which these networks change. What is the basis for this understanding of organizational change? I want to suggest that the basis lies in a socio-cognitive quantum mechanics.

In the social world, as in the physical world, there is an underlying quantum mechanics. That is, the behavior of a social system results not just from a simple aggregation of the behavior of isolated individual entities (be they people, groups, organizations, institutions, or societies), but emerges from the capabilities of the entities and the dynamics by which these entities interact. However, individuals in the social world, unlike electrons in the physical world, can learn and so change their fundamental properties over time. Individual learning has consequences for all social systems as it is the primary determinant of change in social and organizational networks.

To develop our understanding of social systems in general, and organizations in particular, we need to develop a socio-cognitive quantum mechanics. In this paper, the basic precepts of such a theory are put forward. These precepts rest on research in a number of disciplines and are related to a variety of theoretical conceptions. Collectively, however, these precepts provide a basis for understanding social action from the ground up. These precepts lead to direct predictions that can be tested empirically. Illustrations of the type of data needed to examine some of the predictions are provided. Additionally, computational models consistent with these precepts can be used to illustrate, explain and theorize about social behavior. As such, one use of these models is to derive a number of hypotheses about social behavior that can then be examined with different types of data. Using such computational models, some of the implications of this theoretical approach are described and illustrated for organizations with particular attention to the evolution of social and organizational networks. These models demonstrate the way in which different types of learning can result in "clashes" which collectively alter the course of network evolution and the resultant form of societies and organizations.

FUNDAMENTAL PRECEPTS

The precepts of such a socio-cognitive quantum mechanical theory of social behavior are straight forward: agency, knowledge as structured, action as interaction, synthetic adaptation, an ecology of networks, emergent reality, constraint based action, and the primacy of learning. Let us examine each in turn. Aspects of this discussion will be illustrated using Figure 1 which is a snapshot of the interaction-knowledge space for five individuals in a small hypothetical company at a point in time.

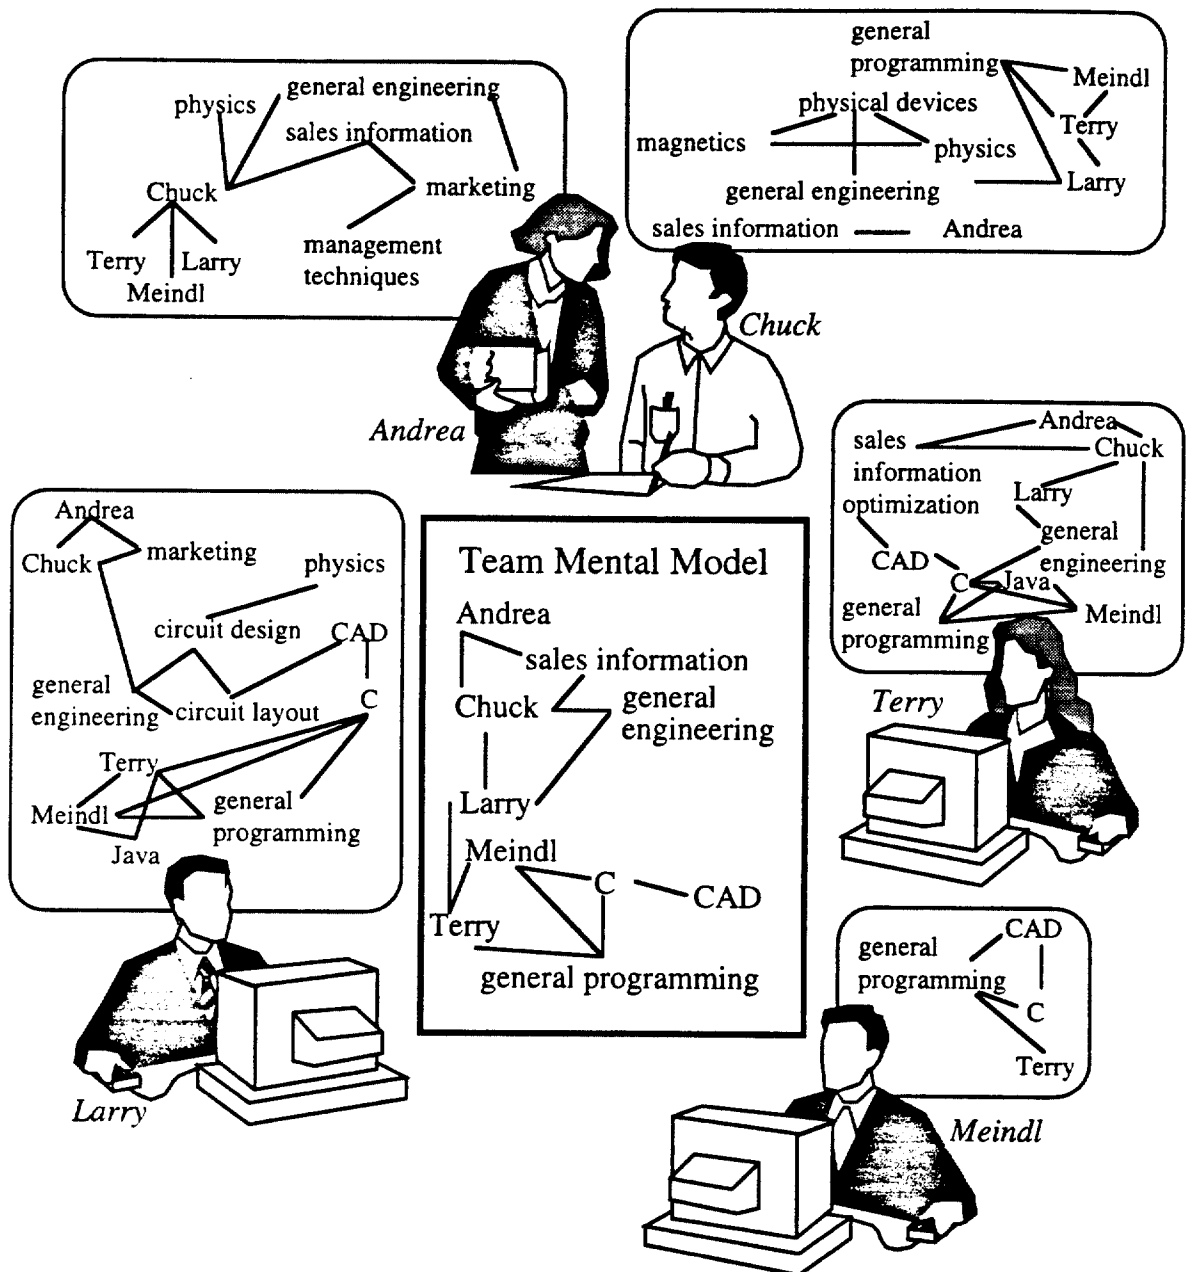


Figure 1. Relating Individual and Group Knowledge

[1] Agency

It is non-problematic to talk about people as having agency. Humans, after all have the ability to take action. The fundamentals of agency, however, are wrapped up in what we mean by agent and computational.

AGENT — An agent is simply a computational entity that exists in an interaction-knowledge space.

COMPUTATIONAL — Entities are computational if they have the ability to do any of the following: acquire, process, store, interpret, or provide information and/or the connections among pieces of information. These computational abilities plus the agent's knowledge define an agent's "intelligence."

Humans, the quintessential agent, have the ability to reposition themselves within and to change both the knowledge and interaction dimensions. But there are many other entities that also can be characterized as agents. For example, we can think of books as agents which contain or store information and can provide it to others, albeit passively. We can think of books as having a fixed position in the knowledge dimension but as being moved in the interaction dimension as different agents interact with the book. Robots, groups, and organizations are also agents with varying capabilities in these dimensions. In Figure 1, the individuals, the book held by the woman at the top, the databases and programs on the computers, are all agents. It is sometimes useful to distinguish agents on the basis of meta-properties such as adaptation, learning or evolution which result from particular combinations of other information processing capabilities (such as acquisition and storage). For example, compared to humans, books are a simpler, "less intelligent" class of agents as they are not adaptive (i.e., they cannot learn or evolve).

In a sense, this notion of agency is the familiar information processing approach to human behavior so familiar to organizational theorists (March and Simon, 1958; Simon, 1944; Galbraith, 1973). However, the argument really goes beyond this. To wit, we can only understand the differences in the way in which agents will act if we understand the differences in the way in which they access and process information. This requires more than comments such as humans are boundedly rational. It requires specifying the precise cognitive capabilities and knowledge of the various agents. The more detailed the specification the greater the range and predictive power of the theory. Additionally, this notion of agency goes beyond the basic information processing framework because it rests on the recognition that the need for action results from cognitive limitations (Carley and Newell, 1994). Cognitive capabilities (exactly how the agent handles information) and knowledge (what information the agent has and how it is structured) determines both intelligence and what actions the agent can take and needs to take (Carley and Newell, 1994). Capability plus knowledge define the set of potential actions. The completely omniscient agent has no need to ask questions or engage in information gathering activities of any kind. Only as limitations set in does the agent need to engage in actions. However, these limitations in and of themselves are insufficient to define the set of actions that

are possible, let alone the set that are probable. Rather, possibilities and probabilities are determined by both the context in which the agent is situated and by the agent's physiology and knowledge.

[2] Knowledge as Structured

Knowledge includes both concepts or information and the connections among those concepts or bits of information. Consequently, there is an essential structural aspect to intelligence. An individual's knowledge can be portrayed graphically, as in Figure 1, as the network of concepts known by each person at a particular point in time (network near each person's head). As noted by many researchers, these concepts include meta-concepts (a densely connected set of concepts) and these connections include connections to meta-concepts. Meta-concepts and the connections to them serve to organize the entire content of the individual's knowledge into various schemes, mental models, or hierarchies (Johnson-Laird, 1983; Carley, 1997). Indeed the entire content of the individual's knowledge can be thought of as a meta-concept; e.g., "What Aria knows." In Figure 1, there are implicit connections in each individual's mental model between themselves and all of the concepts in their mental model (including those concepts representing other individuals).

MENTAL MODEL — A mental model is the set of concepts, meta-concepts and the connections among them known by an agent.

As noted, cognitive capabilities include the ability to manipulate information (nodes) and the connections among pieces of information (relations). We think of these capabilities as learning. Agents who are capable of learning may exhibit different behaviors at different times even in the same situation, simply because their knowledge has changed (such as by adding or forgetting concepts, meta-concepts or relations among them).

[3] Action as Interaction

Action results not just from response or opportunity but from the inter-relation among knowledge (including information on context) and capability (including need for possible actions). In most cases, the range of associated knowledge is much greater than the range of action. Although actions are constrained there is a many-to-one mapping between knowledge and action. Thus you cannot uniquely backtrack from the action to the knowledge that led to it. Actions, if they occur within the agent, are not directly observed by others and do not directly affect others. We can think of such actions as cognitive in nature and thus related to learning. Actions that occur external to the agent involve other agents, either directly or because they are observed by others. Such actions are effectively interactions. The agents who interact need not be human. For example, in Figure 1, the two individuals at the top are interacting by talking with each other; other individuals are interacting with various computer based databases and programs (each of which is an artificial agent). Action is not

necessarily purposive. Action can result from the simple passive reception of information or its automatic interpretation. Some agents may have a cognitive architecture that causes them to be completely goal directed. Humans, despite theories of economic man, may not fall in this category.

[4] Synthetic Adaptation

Any entity 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. Such entities are composite agents.

COMPOSITE AGENT — *Any agent formed through synthesis of other agents.*

Synthetic adaptation produces composite agents who interact with, and can perform the same tasks as non-composite agents. For example, in an inter-organizational response unit (Dynes and Quarantelli, 1968; Topper and Carley, 1997) some of the response "organizations" are single companies, some are consortiums of companies, some are network organizations, and some are groups of individuals acting collectively as an institutional unit. Further research is needed on how this synthesis occurs. However, at its heart, synthetic adaptation naturally derives from the structural nature of knowledge.

Importantly, the connections between concepts are not relegated to occur only within the mind. Rather, there are knowledge connections between individuals, such as "shared ideas", "I know that you know", and so on. For example, in Figure 1, Larry, Terry and Meindl all share general programming knowledge, knowledge about C and the relation of C to general programming. Larry "knows" that Chuck and Andrea know about marketing; however, he is mistaken about Chuck. Andrea, "knows" that Chuck interacts with Terry, Larry and Meindl. Having knowledge does not imply being correct. Knowledge exists within and between individuals. Consequently, knowledge must exist within and between any group that contains individuals. As individuals learn they alter the distribution of information and so the knowledge held within and among individuals, which in turn gives the group, and indeed any composite agent, cognitive capabilities and the ability to learn.

To explain organizational behavior we need to understand that the organization, in and of itself, is an intelligent, adaptive and computational entity (Carley and Gasser, forthcoming). As such, it can take action distinct from, and not predicted from, an aggregation of individual actions. This is not to anthropomorphise the organization. Far from it. The organization's intelligence, adaptiveness, and computational capability results from the detailed, ongoing, interactions among and behavior of the member agents. The principles of combination that generate group behavior are more complex than

simple aggregation. Work pointing in this direction includes that on distributed cognition (Hutchins, 1991, 1995) and transactive memory (Wegner, 1987, 1995; Moreland et al., 1996; Moreland, in press). To illustrate this point, consider the nature of social knowledge. For a group, the social knowledge or team mental model can be conceived as the set of concepts and relations that are shared by the majority of the team members. This is represented graphically in the central box in Figure 1 where all concepts and relations shared by 3 or more people are listed. For a group, the potential knowledge might be that which is knowable by the group and therefor is known by at least one team member. Assuredly, there are other conceptions of social knowledge. The potential group knowledge is simply the union of all of the concepts and relations in Figure 1. Notice, even something as seemingly simple as a group's knowledge is more than a simple sum of the parts. For example, in Figure 1, no individual knows all of the information in the team mental model.

[5] An Ecology of Networks

Networks exist within an ecology of networks. These cross-cutting affiliations influence individual and group behavior (McPherson, 1983) and serve to constrain and facilitate change (Granovetter, 1985). Most models of network evolution, however, consider only a single network (Sanil, Banks and Carley, 1995). Yet, the fact that networks are embedded in each other is inescapable. The social network denoting who talks to whom is intertwined with each individual's cognitive network (the way in which each individual links ideas) 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 interlinked with many other networks including the task structure (which tasks are connected to which), the task access structure (who is assigned to what task) (Krackhardt and Carley, 1998). Change in any part of this ecology of networks ultimately affects all other parts and the behavior of the entire system is a function of the specific way in which these networks are inter-linked.

Traditional structural analyses explore social and organizational issues by focusing on networks composed of one type of agent — only people (Marsden and Lin, 1983; Wellman and Berkowitz, 1988; Wasserman and Galaskiewicz, 1994) or only firms (Mizruchi and Galaskiewicz, 1994). However, not all agents in a network are necessarily of the same type (contrary to most currently collected data sets). One reason for this is synthetic adaptation; another is that agents have different cognitive capabilities (Kaufer and Carley, 1993). Different classes of agents (people, books, webbots, teams, or organizations) despite differences in capabilities, knowledge, or composition may still be connected in the same way (e.g., they may all still communicate with each other). Relation type and not

node type becomes the primary boundary determinant when collecting network data.

Consequently, traditional structural analysis, which bounds the network by the type of agent, rather than the type of connection, may result in erroneous or misleading conclusions about the role of the network in producing social change. Thus, 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) all classes of agents, all sources of information need to be considered. For Kapferer's (1972) tailor shop, the two classes of agents were simply managers and workers, and taking both into account and the differentials in their knowledge makes it possible to better explain the changing potential to mobilize for a strike (Carley, 1990). For Krackhardt's (1992, 1995) firm the only non-human "agents" were union pamphlets distributed at a group meeting and considering these pamphlets should produce a better account of the event. Another implication of this view is that an examination of the emergence of network organizations that focuses only on corporations will result in a misperception of how the collectivity operates.

[6] Emergent Reality

All social, cultural and individual behavior emerges out of the ongoing interactions among intelligent adaptive agents. There are several parts to this argument. The first part is that the fundamental social act is interaction and that without such interaction there can be no social behavior (Carley, 1991; Carley and Newell, 1994). Second, sets of intelligent agents are self-organizing due to capability and knowledge constraints (Padgett, 1997; Epstein and Axtell, 1997; Kaufman, 1995). Third, learning from others and the presence of networks can speed the rate of team learning (Prietula and Carley, 1994), generate misperceptions, and result in cycles of beliefs. Thus, regularities in behavior across, among, and over time for these agents and composite agents emerge as agents interact. Fourth, the social and cultural world is, in part, continually reconstructed through these on-going interactions as individuals create, manipulate, utilize, and exchange information and so symbols (Stryker, 1980). Fifth, since norms, regulations, institutions and so on emerge as agents interact, the specifics of who interacts with whom when determines which norms, regulations, and institutions emerge, when, and their effect. In this sense, social change emerges, to an extent, from the right person being in the right place at the right time (White, 1992). It also means that there is no ontological imperative that gives one agent (be it person, group, company or institution) more apriori import on an individual's or company's future than another.

[7] Constraint Based Action

A primary feature of any socio-cognitive quantum theory must be a recognition of constraints. Unlike a pure constructivist framework, where all is

socially constructed, here it is recognized that there are temporal and physical constraints (Mayhew, 1984; Latane et al., 1995). At each level of agency, factors exogenous to that level serve to constrain what actions are possible. These constraints reduce the set of potential actions to the set of acceptable actions. These constraints can be so severe that they eliminate all but one course of action. At the simplest level this is a matter of saying that there are temporal and physical limitations (such as there are only 24 hours in a day and people only have two hands). At a more detailed level, however, this means that within an organization the actions taken by an individual are typically constrained by factors external to them including time, physical constraints, task, technology, social structure, and culture. In this way these factors have a direct effect on actions. Additionally, each of these factors may have an indirect effect on action as mediated by cognition and so mediated by the knowledge held by the agent. For example, individual actions are constrained not by culture per se', but by culture as perceived by the individual.

Specifying these constraints in detail is necessary to predict and explain action. At the individual level, models, such as GOMS (Card et al., 1983; John and Kieras, 1996), that predict human performance by taking into account the interaction between cognitive processing and the timing of physical actions at the micro-level (such as eye-movement and key-stroke) have demonstrated high predictive accuracy. At the group and organizational level, attention to the constraints imposed by tasks and the communicative properties of technologies is increasing the predictive power of organizational level models (Levitt et al., 1994; Kaplan and Carley, 1998). Specifying constraints is a familiar notion to organizational theorists in the information processing tradition (Galbraith, 1973; Salancik and Pfeffer, 1978). However, specifying constraints requires more than a recognition that constraints exist. It requires specifying the precise set of tasks, networks, institutions, resources, knowledge, agents and technology that affect the flow of information. The issue here is not simply a matter of accuracy. The more theories, and associated models, take into account specific constraints (tasks, structures, etc.) the greater the range of predictions and the more specific the predictions possible (Carley and Preitula, 1994).

[8] Primacy of Learning

Perhaps the main distinction between a socio-cognitive quantum mechanics and the quantum mechanics of physics is that people, unlike electrons, learn. In fact, it is probable that people cannot stop themselves from learning. The ubiquity of learning means that there is always the possibility of change. Individuals appear to learn in a variety of ways. For example individuals learn by observing others, by problem solving, by generating expectations, by interacting with others and so on. In part, the difference in these learning "mechanisms" is a difference in the feedback or information that the individual

uses.¹ Three types of learning are particularly interesting from an organizational perspective: communication based, experience based, and expectation based. In communication based learning individuals learn about tasks, people, organizations, etc. by observing or being told. The information garnered in this way is expected to be new or novel to the learner. Experiential learning has its basis in task repetition and feedback. There are several sources for this experience: the communication of previous results, increased familiarity, increased physical skill, prior problem solving. Finally, expectation based learning occurs when individuals engage in planning, thinking ahead about the future, and then use these expectations as a basis for future reasoning.

From a network perspective, learning results in the construction of nodes and relations. Consequently, learning can and does occur at multiple levels. For organizations, it is useful to distinguish between individual and structural learning (Carley and Svoboda, 1996). Individual learning occurs within the individual. At this level, learning results in the individual changing his or her mental model by adding or dropping either concepts and/or relations among concepts. These changes, however, may precipitate changes in interaction. Thus, cognition mediates interaction and individual learning mediates structural learning. Structural learning occurs within or among composite agents (such as groups, organizations, institutions). At this level, learning results in the adding or dropping of agents (individual or composite) and/or the relations among such agents.

OVERARCHING FRAMEWORK

An overarching framework is shown in Figure 2. We can think of Figure 2 as a conceptual device for illustrating how the two specific models used herein are related to each other and the precepts even while recognizing that many alternative models can be built to instantiate the processes implied by Figure 2. Figure 2 is also illustrative of the way in which the precepts as a collectivity form a processual description of agent and social change. Let us consider how each precept plays out in this framework.

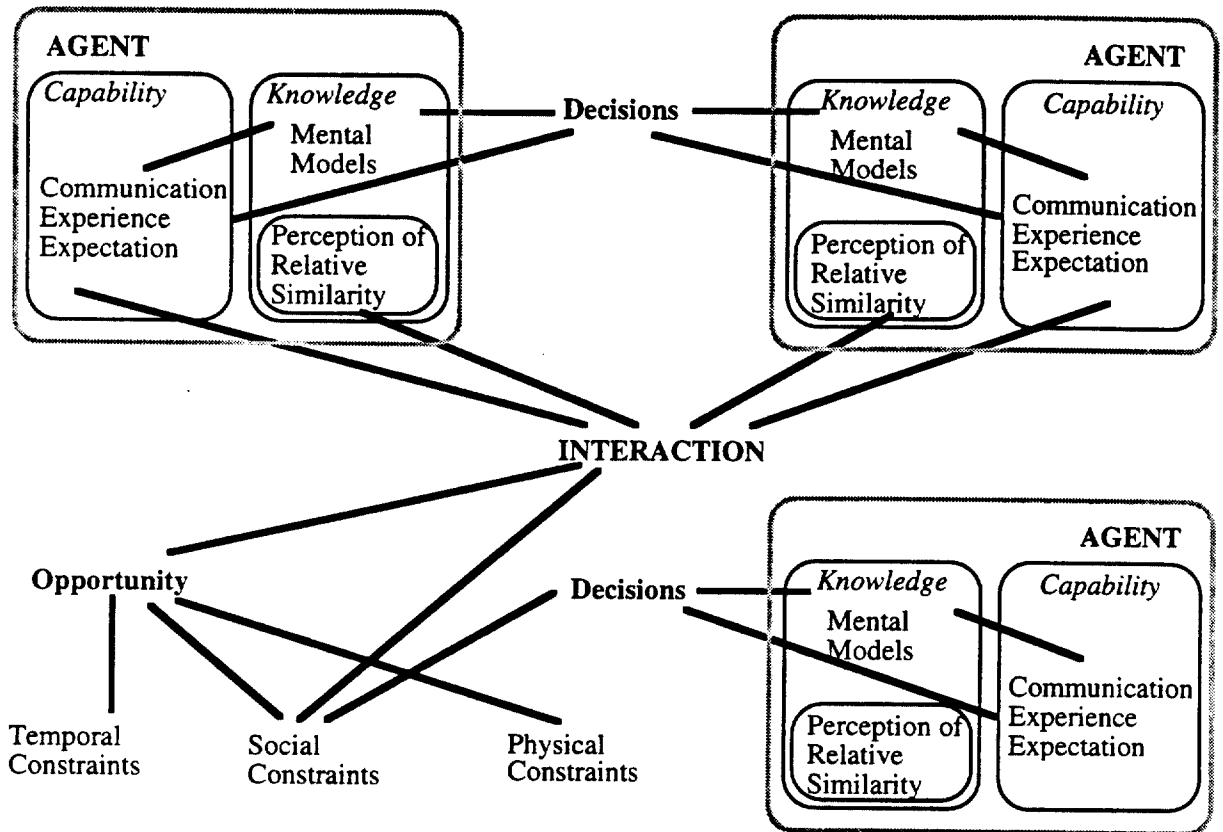


Figure 2. General Framework

Agency and Synthetic Adaptation: All actors are agents and so can be individual human beings, computational agents, or composite agents such as groups, teams, or organizations.

Knowledge as Structured: Agents' mental models includes their perception of the knowledge network (who they think knows what), their perception of the social network (who they think interacts with whom), task knowledge, knowledge of others both in particular and general, knowledge of norms, beliefs, customs and so on of both self and others. Because each agent's mental model includes knowledge of others it includes the cognitive social structure (Krackhardt, 1987, 1990; Krackhardt and Kilduff, 1990) and the cognitive knowledge network (Monge and Contractor, in press).

Action as Interaction: Agents engage in interactions with other agents. Agent's decisions are a function of their extant mental models. Some agents, for various reasons (such as organizational, cultural, or institutional logics) may be in a position where their decisions directly influence other agents' interaction. An example of this is a manager terminating the employment of a subordinate.

An Ecology of Networks: There are knowledge networks within and between agents, networks formed of interactions, and networks formed by joint decisions.

Emergent Reality: Agents learn and take action (including choice of interaction partner) concurrently. Through this concurrent behavior agents, composite agents, and the socio-cultural environment co-evolve. Changes in interaction can influence others knowledge and so attitudes and beliefs (Krackhardt and Porter, 1985).

Constraint Based Action: Cognition mediates action. Interaction is a function of external constraints (opportunity), cognitive and knowledge constraints (perception of their relative similarity to others). External constraints include temporal constraints (e.g., there are only 24 hours in a day), physical constraints (e.g., biological and location), and social constraints (e.g., position, institutional, etc.).

Primacy of Learning: When agents interact they learn (assuming they are cognitively capable of learning). Learning influences the agents knowledge or mental model which includes their perception of relative similarity.

Even at this meta-level several points stand out. First, interaction is the primary defining social act. Second, this conception is inherently dynamic as interactions or decisions lead to learning and change in mental models which in turn leads to change in interactions or decisions. And third, the emphasis is on mechanism or process. This meta-model is a valuable framework for thinking through the impacts of agency and action on social change. Given the complex dynamic nature of this framework, theorizing about the effects in any specific setting is difficult. Consequently, computational models are particularly valuable tools for doing theorizing in this area. Such models, because they capture process can be used to examine a wide range of phenomena.

A variety of models can be, and have been, built that are consistent with this overarching framework. The two models that I will use to illustrate the impact of thinking about groups and organizations in this way are CONSTRUCT (interaction among top two agents in Figure 2) and ORGAHEAD (bottom constraints and agents with no relative similarity judgments in Figure 2).² Each of these computational models has been previously described in detail in the literature. Thus, a full description of these models will not be repeated. Rather, only a few remarks will be made as to how these two models fit into this overall framework.

CONSTRUCT (Carley, 1990, 1991, 1995; Kaufer and Carley, 1993) is a model of social change in response to the diffusion of information among individuals as they interact and communicate with each other. A key model component is that agents communicate a piece of their knowledge when they interact with others. Individuals learn as they interact (through communication) and so their body of knowledge grows. Individuals try to engage in interactions with those with

whom they are most relatively similar in terms of shared knowledge. Individuals' mental models contain both a knowledge network (who knows what) and an impression of the social network (who knows who). Under this perspective, culture is the distribution of information across individuals and social structure is the distribution of interactions across individuals. CONSTRUCT can be 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.

ORGAHEAD (Carley and Svoboda, 1996; Carley, forthcoming) is a model of organizational adaptation. A key component of this model is that organizational performance is affected by actions of the CEO or executive committee at the strategic level and actions of the members of the organization at the operational level. Operational personnel learn as they work on tasks through both experience and communication. Their position or role in the organization influences their decisions and to whom they communicate their decisions about the task. The CEO learns through both experience and expectation. The CEO makes a decision about the task for the organization as a whole. CEOs can alter the organizational structure and so alter the personnel's access to information. Individuals' mental models contain a task model and a knowledge network for subordinates (which subordinates know what). The CEO's mental model includes information on performance, who knows what, previous changes, and expectations about alternative structures. ORGAHEAD can be used to look at a variety of organizational issues including the impact of different change strategies, constraints on organizational re-design, and the impact of rate of changes on performance for organizations faced with either stable or changing environments.

IMPLICATIONS

A variety of implications for social and organizational change follow from this approach. Three will be focused on: the evolution of identity, the trend toward stability, and social differentiation. For each focus area an example from the micro or individual level and from the macro or organizational level will be given.

Evolution of Identity

As agents interact they can continually form and reform their identity and their perceptions of other's identities. An agent's identity is made manifest in their mental models and in their pattern of interactions with others. Changes in identity can be measured as changes in what concepts and connections are known by the agent and changes in the set of interaction partners and the pattern of interaction in which the agent is embedded. Identity has a cultural component in terms of the pattern of knowledge held by the agent and the

agent's knowledge ties to others. Identity also has a structural component in terms of the pattern of interactions. Identity has a subjective component - the agent's own perception of his or her position in the overall interaction-knowledge space - and a social component — other's perception of the agent's position in the interaction-knowledge space. This is illustrated in Figure 3.

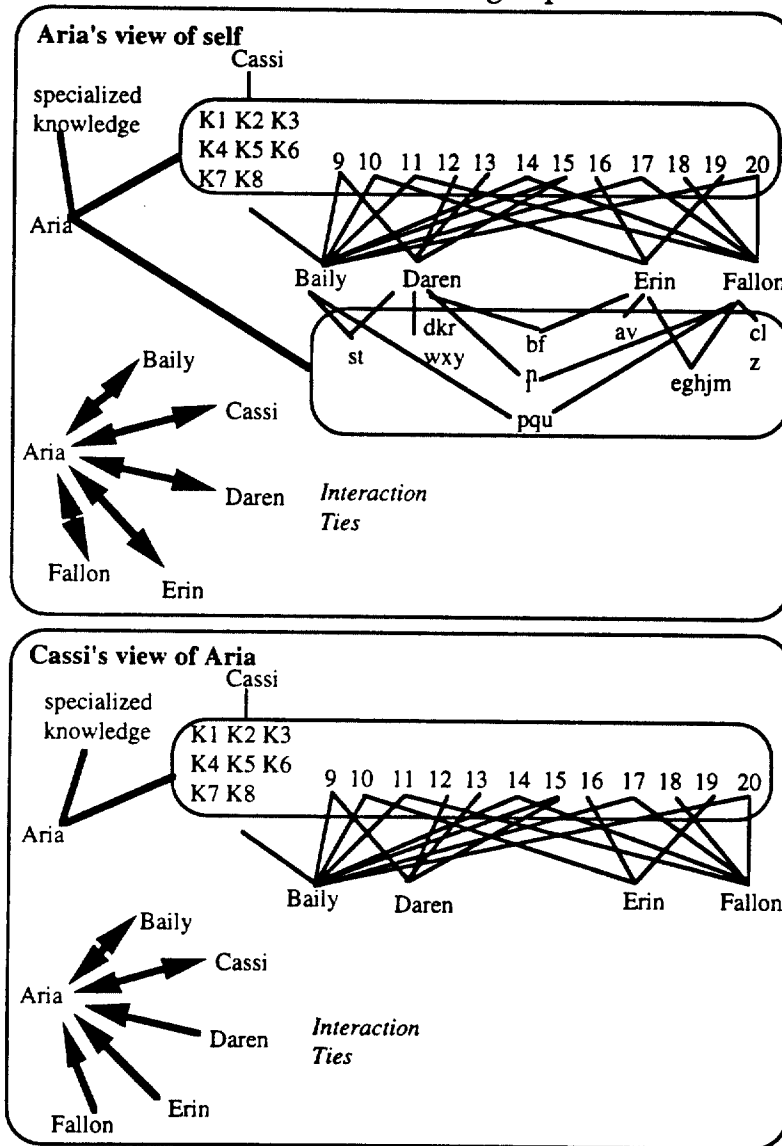


Figure 3. Aria's and Cassi's view of Aria's Identity

A simulation was set up to run using CONSTRUCT. In this simulation, the group was modeled as six people roughly divided into two groups (group 1: Aria, Baily and Cassi; group 2: Daren, Erin and Fallon). Each individual was

characterized as having a certain percentage of the available 110 pieces of information. Additionally, Aria, Baily and Cassi are modeled as having a set of things in common that are readily observable (such as gender or race) and that set them apart from the other three individuals. At the top of Figure 3 is a graphical portrayal of Aria's initial information, whom she thinks she shares what information with, and her perception of whom she is likely to interact with. At the bottom of Figure 3 is a graphical portrayal of what information Cassi knows she shares with Aria, which information she thinks Aria shares with others, and whom she thinks Aria is likely to interact with. Aria sees her self as integrated into the group and choosing to interact equally with all of the others, being equally chosen by all others to interact, and with sharing, different, but similar amounts of information with all others. Cassi, on the other hand, sees Aria as being tightly connected to Baily and Cassi and having little in common with Daren, Erin and Fallon. Although Cassi also thinks that all group members will be equally likely to choose to interact with Aria. Differences in what is known by Cassi and Aria lead to differences in expectations about who will interact with whom.

At the organizational level identity is also defined in terms of the interaction and knowledge ties. Thus, organizations like human can also change their identities by moving about in the interaction-knowledge space. For organizations their position in the knowledge dimension can be thought of as capabilities and position in the interactions dimension can be characterized in terms of alliances, agreements, and various institutional arrangements. We can see changes in organizational identity by examining changes in alliances and capabilities. Westinghouse, for example, transformed its identity from an electronics and appliance company into a small player in the communication and media industry.³ Unlike humans, organizations can convert interactions into knowledge. A merger is a movement converting an organization's interactions with another organization into connections between one organization and a set of capabilities or knowledge. This is a type of learning though consolidation. Organizations, also unlike humans, can de-convert knowledge into interactions. For example, this can occur through breakups and divestitures. Organizational strategy thus involves the converting of one type of tie into another and the timing of these conversions.

For humans most changes in identity occur insidiously as people engage in interactions with others, age, make new acquaintances, change jobs, ..., and learn. For a composite agent, such as a group or an organization, the gradual change in the identities of the component agents results in a gradual change in the identity of the composite. The idea here is similar to the notion of a bottom-up or grass-roots basis to cultural change. Such change to the composite's identity will typically take a long time to manifest itself. In that sense, we can say that the rate

of change in identity decreases as the degree of composition increases. A change in a composite agent's identity, however, can have traumatic impact on the identities of the component agents. For example, when companies dissolve the personnel move to new companies and take on new identities as members of different companies and different teams. A prediction that follows from the precepts is that the magnitude of the change in identity to a composite entity is magnified in the components as the degree of composition decreases.

Ultimate Stability

This continual creation of identity occurs against a backdrop in which there is a general movement amongst agents toward stability. This movement occurs in the long run and in the absence of change events (such as discovery, innovation, catastrophes, or forgetting). On the one hand, this trend toward stability can be thought of as a natural result of learning mechanisms and a fixed body of knowledge to learn. Learning can occur during passive exchanges, such as when individuals garner new ideas from the mass-media. Many actions and interactions are accidental. Consequently, stability can result from non-purposive behavior.

On the other hand, stability can also result from goal directed or purposive behavior. Goal directed behavior can be characterized in terms of optimization; i.e., in an unchanging environment stability is reached when some function is optimized. For example, Behrens (1997) suggests that individuals try to maximize the stability of their networks. Hence, under periods of stress, individuals will drop from their networks those whom they perceive as creating local instabilities. Evidence for this thesis is provided by an in-depth analysis of individuals who are tested for HIV. Behrens shows that, even though they do not think their networks are changing, individuals are altering their networks by reducing interactions with those whom they perceive as unstable.

This trend to stability can be seen in the interaction patterns among people. Using CONSTRUCT the interaction and learning behavior of a small group of six people was simulated. In this simulation each individual agent, or person, could interact, communicate, and learn. Initially, the six people were roughly divided into two groups of three such that group members were more likely to share information with each other than non-group members and all group members were highly likely to interact with Cassi who acts as a bridge between the two groups. Figure 4 is a graphical portrayal of changes in Cassi's (the initially most central person) likelihood of choosing to interact with each of the other five people in the group. As can be seen, Cassi's probability of interacting with specific others oscillates over time and then eventually stabilizes. As Cassi comes to share all the information known by others her likelihood of interacting with others approaches $1/N$ and she becomes equally likely to interact with all group members.

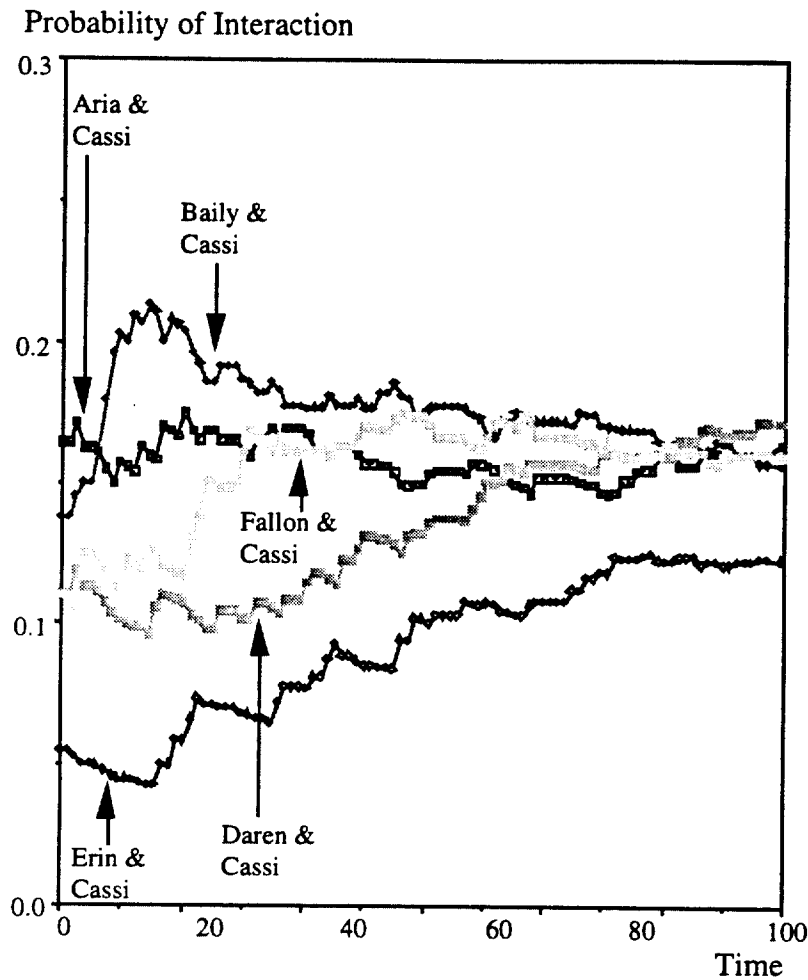


Figure 4. Change in probability of interaction for five dyads over time

This trend to stability can also be seen at the group or organizational level. Using ORGAHEAD 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. In each organization the individual agents, or people, could learn and the organization as a whole engaged in structural learning. The resultant behavior for five top performing organizations is shown in Figure 5. This figure illustrates that despite individual and structural learning, over time the variation in performance decreases for these organizations. Over time 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 stability, in a stable environment can be beneficial for those organizations who locate satisfactory structures.

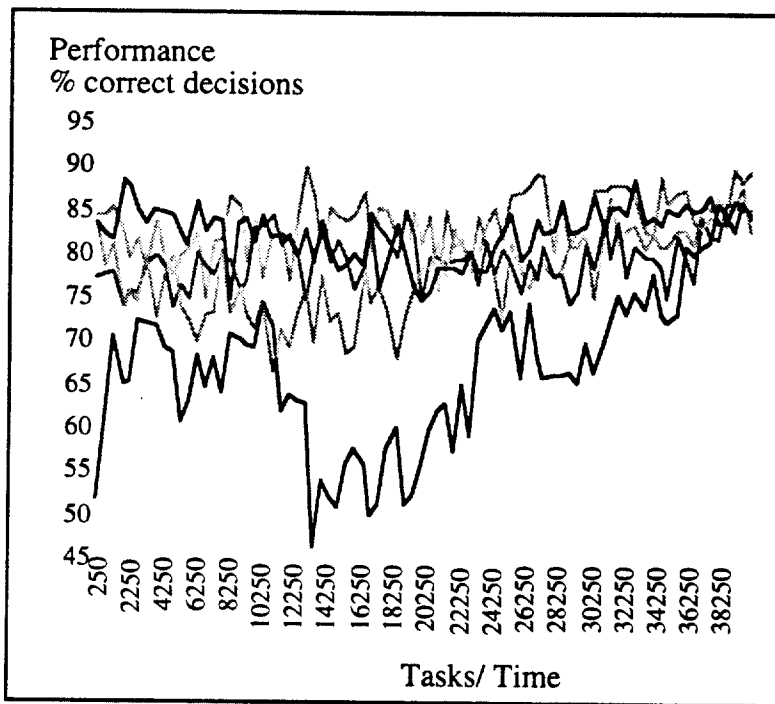


Figure 5. Over time behavior for five hypothetical organizations faced with a stable environment

At both the individual and organizational level we see a general trend to stability. This general trend is less interesting than the path of getting there. At both levels of analysis behavior is irregular in the short term. These oscillations appear random. That is, they do not occur in a specific period and the size of the oscillations varies. In the case of individuals specific interactions and the communication of specific information determines the exact pattern of oscillation observed. In the case of organizations, specific strategic changes, specific tasks, and at the operational level the specifics in who communicates what to whom and who knows what determines the pattern of oscillation. In both cases, there is a path dependency to behavior that results from the instantiation of specific socio-cognitive mechanisms in particular situations. Consequently, despite appearances, the behavior is not random.

As an aside, it is worth noting that most models of network change predict an ultimate stability (Sanil et al., 1995). Thus the trend toward ultimate stability is not the novel implication of the two models discussed herein. Rather, the novelty lies in demonstrating that such stability emerges from simple socio-cognitive mechanisms. It is also worth noting that the trend toward ultimate stability does not imply a trend toward a single outcome. In fact, in most situations there are multiple endpoints and stability can mean that the collection

of agents will ultimately divide into self-reinforcing groups. For example, ORGAHEAD simulations 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 sector. Moreover, the internal structure of the high and low performance organizations will be different (Carley, forthcoming). Even though organizations begin very similarly they will tend to take different paths and get locked into structures where the internal knowledge and interaction networks that develop lock them into a particular pattern of performance. High and low performers are employing the exact same learning mechanisms. Divergence occurs because different organizations learn different things which affects whether and when the various learning mechanisms collide and so what meta-learning strategies evolve. Ultimate differences result not from capabilities (mechanism) but from the content of what is learned, from local choices that move the organization through the interaction-knowledge space.

Social Differentiation

This trend toward ultimate stability simply points the way that societies and organizations tend to move on average. In most cases, disruptions will occur long before stability is reached. The issue then is not how long does it take to reach stability, but how fast can agents recover from disruptions. These disruptions can take many forms including changes in the environment, new institutional regulations, natural disasters, technological accidents, discoveries, ..., or changes in the population of agents.

What happens along the way to ultimate stability? In the short run, this trend toward stability may not be noticeable. However, it is through these short term oscillations that stratification emerges and shakeouts occur. The basic socio-cognitive mechanisms that cause short term oscillations in interaction and performance also cause individuals and organizations to enter into tracks through which they become increasingly differentiated. In part, these oscillations are a natural result of learning. However, and of critical import, many oscillations (and the resultant shakeouts) are a result of collisions between the "types of learning." For example, at the individual level learning through observation and learning through interaction may lead to "shakeouts" in personal friendships. Imagine a group of people, put into a new situation, such as freshmen in a dormitory. Initially, friendships will form on the basis of relative similarity in observable characteristics (gender, race, etc.). In fact, homophily based interactions are common (McPherson and Smith-Lovin, 1987). At issue is whether these superficial or observable ties are sufficient to support future interaction. CONSTRUCT suggests that in groups, although initial interactions may be homophily based, over time, as individuals interact, if their

knowledge is not co-aligned with these observable characteristics then people will learn about deeper similarities based on other non-observable knowledge and those ties will come to dominate the interaction.

For example, the group of individuals shown in Figure 3 were simulated using CONSTRUCT for 100 time periods. Initially, the set of core knowledge shared by group members Aria, Baily and Cassi is sufficient to focus the majority of their interactions. Over time, as individuals interact and begin to share other information, the pattern of interaction becomes complicated (time 50) and then finally, when the new shared knowledge comes to outstrip the initial shared knowledge (time 100) the pattern of alliances changes (see Figure 6). Thus, while people initially interact due to observable or obvious characteristics as they interact and learn more about what other information they share, and as they go through joint experiences, they will change their interaction partners. In Figure 7 this is shown with Aria initially interacting with Baily, developing a relationship with Fallon, and then dropping the interaction with Baily. In other words, as experiential knowledge comes into conflict with initial observational knowledge changes in interaction partners occur. This can appear as a "friendship" shakeout when initial interaction partners are forsaken and longer term stable relations form. If on the other hand, the underlying knowledge is aligned in the same way as the observable information then future interaction will serve to re-enforce observable differences and stronger degrees of stratification will result.

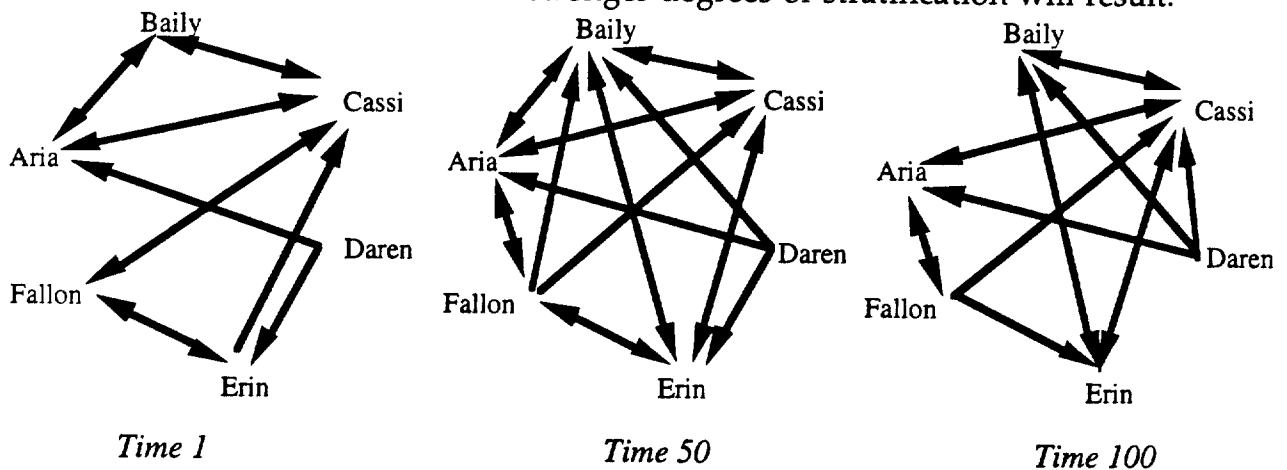


Figure 6. Change in interaction partners over time

At the organizational level, an ORGAHEAD virtual experiment demonstrates that clashes between structural and experiential learning can cause increasing differentiation in organizations that are initially very similar. A set of 100 organizations were simulated for 100 time periods in a stable environment. Initial structure was chosen randomly. In Figure 7, changes in performance and design for the best/worst organizations is shown. Over time, oscillations occur

in behavior, performance, and design. Over time organizations diverge in performance (top of Figure 7) and in structure (bottom of Figure 7).

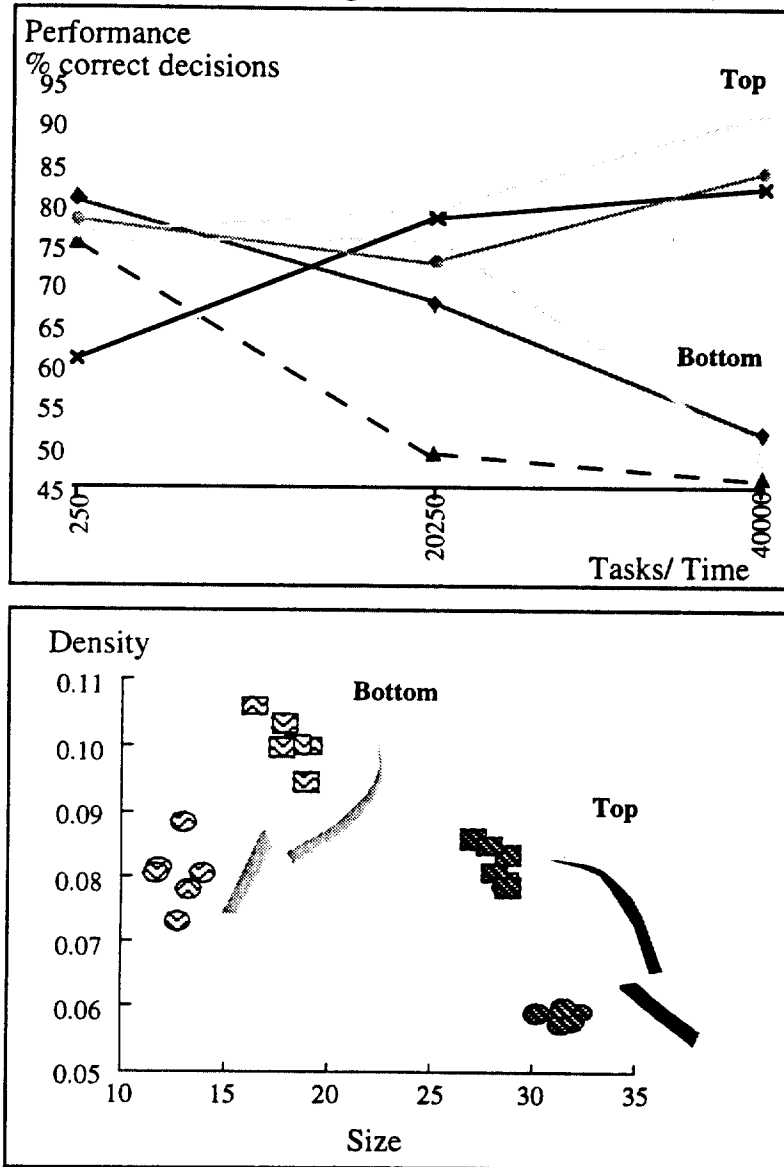


Figure 7. Increasing divergence in organizational performance and structure

Additionally, the way in which these organizations change also diverges. For these organizations the top performers grow slightly in size and become less dense, whereas the bottom performers become less dense and decrease in size. Ultimately, poor performing organizations engage in a pattern of strategic operation where they typically oscillate between bouts of hiring and lay-offs. This pattern of strategic change results in a loss of operational knowledge as

highly trained personnel are laid-off and a loss of structural knowledge as hiring reduces remaining personnel's knowledge of who knows what or who knows who knows what. These changes represent large shifts of the organization and its' personnel in the interaction-knowledge space. Whereas, high performing organizations are more likely to engage in strategic activities which re-engineer the firm through altering connections among personnel, information, or both. These small shifts of the organization and its personnel in the interaction-knowledge space result in higher performance and a different structural profile. These results should not be taken as saying that successful organizations will be, in general, small, low in density, and employ re-engineering rather than lay-offs. These specific findings may be a product of the task being performed by the organizations in the model. The main point is that the exact same learning mechanisms will 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 may be failing not 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.

TOWARD A SOCIO-COGNITIVE PHYSICS

The argument here is quite simple, a unified theory of social and organizational behavior should be possible if it has at its basis a socio-cognitive quantum mechanics. Principles underlying such a mechanics would include: agency, knowledge as structured, action as interaction, synthetic adaptation, an ecology of networks, emergent reality, constraint based action, and a primacy of learning. Complex social and organizational behaviors result from simple learning mechanisms operating within a system of constraints and an ecology of changing networks.

An important aspect of this argument is that the theory is in the details. That is, in order to predict behavior the mechanics and specific constraints will need to be specified at the quantum level (i.e., at the level of the basic agents in the system). There are several implications. First, the higher the level of detail in the models and theories the higher their predictive capability. Second, methodologically this means that the computational modeler's best friend is the ethnographer. Only in the details of a thorough case study are there sufficient details for the computational theorists to begin to glimpse answers to the numerous questions about constraints, networks, distributions of information, and so on that are needed to develop even a relatively simple computational model. A second aspect of this argument is that resultant explanations derive from the duality of knowledge and action. At the social level what this means is that the basic socio-cognitive mechanisms when carried out across and within large numbers of individuals results in the co-evolution of social structure and

culture (Carley, 1995). At the organizational level the duality between action and knowledge leads to a co-evolution of the organization's form and the underlying knowledge networks, distributed cognition and transactive memory.

The advantage of this approach lies in the ability to theorize about a multitude of social and organizational behaviors from the ground up. Only a few of the possible implications were presented herein. At a purely speculative level, let us consider what this theoretical approach is likely to say about other basic social and organizational processes.

For example, how do normative structures that guide actions become institutionalized? The argument here rests on the idea that at the heart of institutionalization is the structuration of patterns of information and interaction in a form that is communicable and maintainable beyond the lifetime of its creator or author and in a way that does not depend on any one individual. The theoretical approach outlined herein would suggest that this process hinges on the diffusion and learning processes which when enacted across a society result in norms as shared alignments between knowledge and action. More specifically, as individuals interact they learn a relation between some pattern of knowledge and some sequence of actions. Exactly how an individual learns this relation is not of importance. As individuals learn a relation between knowledge and action they communicate this to others. In general the pattern of knowledge will be broad and complex and the set of related actions will be overly constrained. That is, add or drop one piece of information and the same action is in general still likely to result. Since the pattern of knowledge and actions that is learned is complex individuals may not learn all aspects at the same time. Nevertheless, within a group there will be a core of individuals such that across that core the same relations among patterns of knowledge and actions is more or less shared. This produces a tendency to act in a common way, though perhaps for different reasons. At this point the relation between knowledge and action takes on a life of its own independent of specific group members.

People also communicate by creating artificial agents, such as books, laws, and web-pages. As information becomes communicated through artificial agents it can be communicated faster and can gain strength and added legitimacy. Essentially, these agents take on a life of their own and continue the normative message independent of the author thus increasing the institutionalization of the norm. These artificial agents are particularly legitimizing if, as agents, they cannot learn. Thus, the message that is communicated from such sources is fixed and so continually re-affirmed. This lack of change in what an agent knows increases the likelihood that others will learn that message (in this case the relation between knowledge and action). This can increase the rate at which

normative structures form and become institutionalized and the rate at which institutionalized norms become disparate from the population's view.

When individuals in positions of authority alter the underlying social structure, e.g., by hiring, firing, enacting various constraints on childbearing, education, etc. they are in a position to either re-enforce the extant structure or destroy it and thereby re-enforce or mitigate existing norms. For example, constraints on child-bearing, requirements of "fit" in the hiring process rather than merit, and educational programs which direct the flow of information can serve to reinforce the existing normative basis for particular actions.

Whether the forces of authority and technology serve to facilitate the institutionalization of specific normative structures is dependent on the underlying interaction-knowledge space and the positioning of agents within it. The logic described here focuses on agents, task, interactions (structure), and knowledge (culture) as the primitives. An important part of this argument is that the boundaries around agents are to an extent mutable, particularly for composite agents. Think of a particular set of agents, tasks, interactions and knowledge as a particular "configuration." There may be multiple configurations to achieve any objective, such as populating the stars. To gain a complete understanding of the sociology of these situations it is important to banish assumptions about what constitutes an agent and to completely map out and contrast alternative configurations.

¹ A wide number of models of individual learning exist (for a review see Pew and Mavor, forthcoming). In some cases the proponents claim 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 is beyond the scope of this paper.

² CONSTRUCT is written in standard C, the code is available from the author upon request. ORGAHEAD is written in a combination of C, PERLSCRIPT, and C++, interested readers should contact the author to determine the most feasible mode of access.

³ This figure is based on data collected by Sumit Chowdhury (1998) as part of his Ph.D. dissertation on the changing network structure of the tele-electronics and communication industries.

REFERENCES

- Banks, David and Kathleen Carley. 1994. "Metric Inference for Social Networks." *Journal of Classification* , 11: 121-149.
- Behrens, Dean M. 1997. *Self-Isolation and The Structuring of Support: The Relationship Between Stress and Network Evolution*. Ph.D. dissertation, Social and Decision Sciences Department, Carnegie Mellon University, Pittsburgh, PA.
- Boorman, Scott A. and Paul R. Levitt. 1980. *The Genetics of Altruism* . New York, NY: Academic Press.
- Card, Stuart K., Thomas P. Moran, and Allen Newell. 1983. *The Psychology of Human-Computer Interaction*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Carley, Kathleen M. forthcoming. "Organizational Adaptation." *Annals of Operations Research*.
- Carley, Kathleen M. 1997. "Extracting Team Mental Models Through Textual Analysis." *Journal of Organizational Behavior*, 18: 533-538.
- Carley, Kathleen M. 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. 1991. "A Theory of Group Stability." *American Sociological Review* 56(3): 331-354.
- 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. and David Banks. 1994. "Nonparametric Inference for Network Data." *Journal of Mathematical Sociology* 18(1): 1-26.
- Carley, Kathleen M. 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.
- Carley, Kathleen M. and Les Gasser. forthcoming. "Computational Organization Theory." Ch. 7 in *Distributed Artificial Intelligence* edited by Gerhard Weiss. Cambridge, MA: MIT Press.
- 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 Erlbaum Associates.

- Carley, Kathleen M. and David M. Svoboda. 1996. "Modeling Organizational Adaptation as a Simulated Annealing Process." *Sociological Methods and Research* 25(1): 138-168.
- Chowdhury, Sumit. 1998. *Struggle for Power and Influence: A Network study of Mergers, Acquisitions, Strategic Alliances and Competition in the Communication, Information and Entertainment Industries*. Ph.D. Thesis, H. John Heinz III School of Public Policy and Management, Carnegie Mellon University, Pittsburgh, PA.
- 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
- 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
- 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.
- John, Bonnie E. and David E. Kieras. 1996. "The GOMS family of user interface analysis techniques: Comparison and Contrast," *ACM Transactions on Computer-Human Interaction*, 3 (4): 320-351.
- Johnson-Laird, Philip N. 1983. *Mental Models: Toward a Cognitive Science of Language, Inference, and Consciousness*, Harvard University Press, Cambridge, MA.
- Kapferer, Bruce. 1972. *Strategy and Transaction in an African Factory: African Workers and Indian Management in a Zambian Town*. Manchester, England: Manchester University Press.
- Kaplan, David J. and Kathleen M. Carley. 1998. "An Approach to Modeling Communication Technology in Organizations." Ch. 9 in *Simulating Organizations: Computational Models of Institutions and Groups*, edited by Michael Prietula, Kathleen Carley and Les Gasser. Boston, MA: AAAI Press / The MIT Press.
- 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, Daniel H. 1993. "The Link Between Individual Learning and Organizational Learning," *Sloan Management Review* Fall issue, 37-50.
- Krackhardt, David. 1987. "Cognitive Social Structures," *Social Networks* 9: 109-134.
- Krackhardt, David. 1990. "Assessing the Political Landscape: Structure, Cognition, and Power in Organizations," *Administrative Science Quarterly* 92: 142-154.
- 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.
- 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 and Martin Kilduff. 1990. "Friendship Patterns and Culture: The Control of Organizational Diversity," *American Anthropologist* 92(1): 142-154.
- 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.
- Latane, Bib, J. H. Liu, A. Nowak, M. Bonevento and M. Long. 1995. "Distance matters: Physical space and social impact," *Personality and Social Psychology Bulletin*, 21(8): 795-805.
- 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.
- March, James G. 1996. "Exploration and Exploitation in Organizational Learning." In *Organizational Learning*, edited by M. D. Cohen and L. S. Sproull. Thousand Oaks, CA: Sage.
- March, James G. and Herbert A. Simon. 1958. *Organizations*. New York, NY: Wiley.
- Marsden, Peter and Nan Lin (Eds.). 1983. *Social Structure and Network Analysis*. Thousand Oaks, CA: Sage.
- Mayhew, Bruce. 1984. "Baseline Models of Sociological Phenomena," *Journal of Mathematical Sociology* 9: 258-281.

- McPherson, J. Miller. 1983. "An Ecology of Affiliation," *American Sociological Review* 48:519-532.
- McPherson, J. Miller and Lynn Smith-Lovin. 1987. "Homophily in voluntary organizations: Status distance and the composition of face-to-face group," *American Sociological Review* 52: 370-379.
- 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.
- 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.
- Pew, Richard and Anne Mavor. forthcoming. *Modeling Human and Organizational Behavior: Application to Military Simulations*. Washington DC: National Academy of Science Press.
- 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.
- Simon, Herbert A. 1944. "Decision-making and administrative organization," *Public Administration Review* 4: 16-31.
- Simon, Herbert A. 1956. "Rational choice and the structure of the environment," *Psychological Review* 63: 129-138.
- 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.
- 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.
- Wellman, Barry and Steve Berkowitz (Eds). 1988. *Social Structures: A Network Approach*. New York NY: Cambridge University Press.
- White, Harrison C. 1992. *Identity and Control: A Structural Theory of Social Action*. Princeton, NJ: Princeton University Press.