

**Balancing Exploration and Exploitation through Structural Design:
The Isolation of Subgroups and Organization Learning**

Christina Fang

New York University
44 West Fourth Street, 7-62
New York, NY 10012
212-998-0241
cfang@stern.nyu.edu

Jeho Lee

Korea Advanced Institute of Science and Technology
207-43Cheongryangri-dong, Dongdaemun-gu,
Seoul 190-102, Korea
+82-2-958-3678
jlee@kgsm.kaist.ac.kr

Melissa A. Schilling (Corresponding Author)

New York University
40 West Fourth Street
New York, NY 10012
212-998-0249
mschilli@stern.nyu.edu

Forthcoming, Organization Science

Authorship is alphabetical and all authors contributed equally. We are grateful to Paul Adler, Linda Argote, Jerker Denrell, Michael Katz, James March, Steve Mezas, Nicolaj Siggelkow, and Gordon Walker for helpful comments and suggestions. We also wish to thank seminar participants at the Organization Science Winter Conference in 2006, New York University and Wharton Mini-Technology Conference in 2007. This research is supported by the National Science Foundation under Grant No. SES-0234075.

Balancing Exploration and Exploitation through Structural Design: The Isolation of Subgroups and Organization Learning

Abstract

The classic tradeoff between exploration and exploitation in organizational learning has attracted vigorous attention by researchers over the last two decades. Despite this attention, however, the question of how firms can better maintain the balance of exploration and exploitation remains unresolved. Drawing on a wide range of research on population and organization structure, we argue that an organization divided into semi-isolated subgroups may help strike this balance. We simulate such an organization, systematically varying the interaction pattern between individuals to explore how the degree of subgroup isolation and inter-group connectivity influences organizational learning. We also test this model with a range of contingency variables highlighted in the management research. We find that moderate levels of cross-group linking led to the highest equilibrium performance by enabling superior ideas to diffuse across groups without reducing organizational diversity too quickly. This finding is remarkably resilient to a wide range of variance in factors such as problem complexity, environmental dynamism, and personnel turnover.

1. INTRODUCTION

In 1991, James March published a now-seminal paper on organization learning that sparked vigorous interest in the tradeoff between exploration and exploitation (e.g., Argote and Greve, 2007; Crosson and Bedrow, 2003; Fang and Levinthal, 2008; Fleming, 2001; Gupta, Smith, and Shalley, 2006; Jansen, Van Den Bosch, and Volberda, 2006; Kane and Alavi, 2007; Lavie and Rosenkopf, 2006; Lewin, Long and Carroll, 1999; Miller, Zhao and Callantone, 2006; Siggelkow and Rivkin, 2006). Most of this line of research has concluded that firms tend to overemphasize exploitation (the use of known solutions) at the expense of exploration (the search for new solutions) because exploitation provides more immediate and certain returns (Denrell and March 2001). This focus on short-term returns can cause firms to become trapped in suboptimal solutions, and render them poorly equipped to adapt to a changing environment. Despite the perilous ramifications of overemphasizing exploitation, and the significant work that has been done, we still do not understand much about how this potentially maladaptive process can be avoided (Gupta, Smith and Shalley 2006).

Recent research has suggested that managers may be able to use organization structure as a lever for improving the balance between exploration and exploitation (e.g., Argote, McEvily and Reagans, 2003; Benner and Tushman 2003; Bower and Christensen 1995; Jacobides, 2007; O'Reilly and Tushman 2004; Lenox 2002; Nickerson and Zenger 2002; Siggelkow and Levinthal 2003). Literature spanning fields as diverse as evolutionary biology, small-world networks, "skunkworks" in innovation, and NK search have all yielded some clues that dividing the organization into subgroups might help the organization explore a diverse solution space, while still ensuring that the best solutions are promulgated throughout the organization, subsequently achieving better long-term learning performance outcomes. This argument is perhaps best articulated by Sewall Wright, who argued that if a population of a species existed in a single large community, the species would not be adaptable. Wright observed that instead, evolving populations are typically divided into small subgroups that are sufficiently isolated from one another to limit mating across subgroups. This isolation permits subgroups to maintain genetic diversity and to explore more diverse solutions in the space of possibilities. At the same time, new genes can

propagate throughout the whole system by occasional migrations of organisms from one subgroup to another. In the organization literature, scholars have independently made symmetric claims about organization or industry structure (e.g., March, 2004; O'Reilly and Tushman, 2004; Schilling and Phelps, 2007; Siggelkow and Levinthal, 2003; Fleming and Sorenson, 2001), emphasizing the benefits of some degree of structural isolation. For example, O'Reilly and Tushman argue that for an organization to both encourage the development of new product lines (i.e., explore) and manage its existing product lines efficiently (i.e., exploit), the organization may need an “ambidextrous” structure whereby separate divisions of the firm utilize different rules, norms, and incentives. An R&D division, for instance, may need to be both geographically and culturally distinct from the larger organization to prevent the paradigms of the larger organization from quashing the R&D division’s heterodox ideas.

On the other hand, management scholars usually do not believe that any one organization structure is universally beneficial. A long and rich history on organization structure, loosely organized under the rubric of “contingency theory,” has argued vigorously that the optimal organization structure is dependent on a host of other factors such as the firm’s technology, strategy, or the dynamism of the environment (e.g., Burns and Stalker, 1961; Chandler, 1962; Eisenhardt and Brown, 1999; Galbraith, 1973; Lawrence and Lorsch, 1967; Thompson, 1967). We thus might expect that the semi-isolated subgroup structure only provides benefits under a particular set of conditions, such as when the organization faces particularly complex problems or when the environment changes rapidly, as either of these conditions might require the firm to search a more diverse solution space.

We thus examine two interrelated research questions here: Does a semi-isolated subgroup structure improve the balance of exploration and exploitation, leading to superior long-term learning performance outcomes? And, if the semi-isolated subgroup structure is beneficial, are its benefits dependent on such factors as the complexity of the problems faced by the organization, the degree of environmental dynamism, or the rate of personnel turnover?¹

¹ We use the term “subgroup” throughout the paper to refer to subdivisions of an organization such as product teams or divisions. The degree to which these subgroups are distinct and/or isolated from the rest of the organization

To answer these questions, we use a simulation to compare the learning dynamics of organizations that have varying degrees of subgroup structure, ranging from nearly isolated subgroups of individuals that are only minimally connected to each other, to randomly connected individuals with almost no distinct subgroups. We begin with an organization with nearly isolated subgroups and then systematically vary the interaction pattern by replacing some fraction of the existing links with cross-group links. We find that a semi-isolated subgroup structure, with moderate levels of cross-group linking, is associated with the highest equilibrium performance while either very low or very high levels of cross-group links are associated with lower equilibrium levels of performance. The intuition is that the semi-isolation of groups fosters the diversity of ideas, thereby enhancing long-run learning outcomes. We next examine whether this finding holds under varying levels of problem complexity, environmental turbulence, and personnel turnover. Remarkably, the semi-isolated subgroup structure yields the best long-term learning outcome in most conditions –a finding that both surprised and intrigued us. Our results are thus more consistent with the Wrightian perspective, while not completely ruling out the contingency perspective.

The paper is organized as follows. First we review the literature on exploration versus exploitation in organizational learning, and the work that suggests that organization structure may serve as a potential mechanism to balance this tradeoff. In so doing, we show the line of reasoning that leads to our two primary research questions. Second, we build a simulation model of organizational learning that parameterizes the degree of subgroup structure and connectivity, and present our initial findings. Third, we subject this model to a series of tests that attempt to explore whether our findings are sensitive to a range of contingency variables that have been highlighted in the management research. Finally, we discuss the implications of our findings for management and future research.

2. LITERATURE REVIEW AND RESEARCH QUESTIONS

In the management and complex adaptive systems literatures, *exploration* refers to the search for new, useful adaptations, and *exploitation* refers to the use and propagation of known adaptations (Holland

(“subgroup structure”) is an outcome of organizational design. This is a different conceptualization of subgroup structure from that used in work on team subgroups (e.g., Gibson and Vermeulen 2003), in which subgroups refer to the self-organization of individuals into smaller cohorts within a team based on demographic similarities.

1975; Goldberg 1989; March 1991; Mitchell 1996). Although exploitation yields more certain and immediate returns, it makes the discovery of truly novel solutions unlikely and can lead to obsolescence in the long run (Holland 1992). On the other hand, while exploration can enable the discovery of profoundly novel solutions, it also typically causes a degradation of performance in the short run because searches for novel solutions usually fail. Holland (1992: p. 69) generalized such a dilemma as follows: “Deciding to what degree the present should be mortgaged for the future is a classic problem for all systems that adapt and learn.”

2.1. March’s Seminal Work and Learning Myopia

In 1991, March developed a now seminal model of organizational learning that clearly demonstrated the tradeoffs between exploration and exploitation. In his model, individuals within the organization are initially endowed with diverse sets of beliefs. These individuals interact with an “organizational code” that is meant to represent the organization’s process of socializing members to the organization’s language, beliefs, and practices. The organizational code learns from the best-performing individuals, and individuals, in turn, learn from the code. As a result, the organization’s knowledge improves over time. Individuals that quickly adopt higher performing ideas from the code are termed “fast learners.” Such fast learning efficiently exploits incremental improvements in the organization’s knowledge, and thereby increases the efficiency of organizational learning. However, fast learning also tends to lead to premature convergence on a homogeneous set of ideas or routines, thwarting long-run learning and leading the organization to a suboptimal equilibrium. By contrast, “slow learners,” those individuals that do not quickly adopt the higher performing ideas from the organizational code, allowed the organization to preserve more diversity of individual beliefs, and thus played an important role in enabling the firm to explore a wider range of possible combinations of beliefs. As a result, the presence of slow learners increased the chance of improving the quality of organizational knowledge in the long run. In short, March’s model demonstrated in an organizational context that although exploitation yields more certain and immediate returns, exploration creates and preserves the requisite variety of knowledge necessary for the organization to sustain its learning in the long term (Levitt and March 1988).

The more immediate returns from exploitation tend to cause organizations (and many other adaptive systems) to exhibit a myopic bias whereby exploitation is overemphasized at the expense of exploration (Denrell and March, 2001; Levinthal and March, 1981, 1993). Individuals and organizations tend to pursue solutions similar to already-known solutions because bounded rationality limits their ability to search all possible domains of knowledge (Simon, 1978) and biases them toward more salient areas of their own prior experience (Cyert and March, 1963, Dosi, 1988; Fleming, 2001; Helfat, 1994; March, 1991; Martin and Mitchell, 1998; Nelson and Winter, 1982; Patel and Pavitt, 1997; Stuart and Podolny, 1996). The utilization of existing knowledge, routines and experience increases the likelihood that solutions will be found quickly and at reasonable costs (Dosi, 1988; Fleming, 2001; Nelson and Winter, 1982). As argued by Levinthal and March (1993: 97): “Knowledge and the development of capabilities improve immediate performance, but they often simultaneously reduce incentives for and competence with new technologies or paradigms.” As a consequence, adaptive processes can become self-destructive by leading the organization to become trapped in a suboptimal equilibrium. Despite significant work on this topic, we still do not understand much about how these self-destructive processes can be avoided. This has led to calls for more work on how managers can help firms to better maintain the balance between exploration and exploitation (Gupta, Smith and Shalley 2006; March 2004).

2.2. Structure as a Solution

Recent research has focused on the possibility of using organization structure as a mechanism for managing the balance between exploration and exploitation (e.g., Benner and Tushman 2003; Bower and Christensen 1995; Ethiraj and Levinthal, 2004; Lenox 2002; O’Reilly and Tushman 2004; Miller, Zhao and Calantone, 2006; Nickerson and Zenger 2002; Siggelkow and Levinthal 2003, 2005; Siggelkow and Rivkin 2005; Taylor and Greve, 2006). Of particular interest has been the role of decentralization and isolation: By decentralizing the learning process to subunits of the organization and providing barriers to the rapid diffusion of ideas and norms across those subunits, managers may be able to encourage the exploration of a more diverse range of solutions. In support of this argument, a body of research on “skunkworks” has shown that there can be significant gains from isolating new product development

teams from the mainstream organization (Benner and Tushman 2003; Bower and Christensen 1995; O'Reilly and Tushman 2004). Separating the team from the rest of the organization permits it to explore new alternatives, unfettered by the demands and norms of the rest of the organization. For example, when Steve Jobs began to develop the Macintosh design in 1980, he did not believe the growing corporate environment at Apple (which was already churning out Apple II personal computers at a fast clip) was conducive to nurturing the kind of revolutionary change that he envisioned. He thus created a separate division for the Macintosh that would have its own unique culture. He tried to instill a free-spirited entrepreneurial atmosphere where individualistic and often eccentric software developers would flourish. The small group of team members was handpicked, and sheltered from normal corporate commitments and distractions. He encouraged the Macintosh team members to consider themselves renegades, and even hung a pirate's skull-and-crossbones flag over their building (Rogers, 2003). Notably, however, the skunkworks literature has tended to focus only on the separation of an organizational subunit that has a specific mandate to innovate, and has particularly emphasized situations where disruptive technological change is sought. This literature has thus primarily demonstrated how firms can improve their exploration processes rather than how firms can better maintain a balance between exploration and exploitation. The latter objective is the focus of a second line of work which simulates an organization's search for solutions on an NK landscape (e.g., Ethiraj and Levinthal, 2004; Siggelkow and Levinthal, 2003, 2005; Siggelkow and Rivkin, 2005). This line of work has examined, among other things, whether and how the decentralization of decision making helps firms to avoid becoming trapped on a suboptimal local peak. This work has generally found that decentralization can help the organization explore a wider range of the solution space, thereby lessening its probability of prematurely converging on a suboptimal solution.

2.3. Semi-Isolated Subgroups

The research on skunkworks and NK search described above both share a strong resemblance to a much earlier line of reasoning in evolutionary theory. Evolutionary biologists have long maintained that the isolation of subpopulations can play a valuable adaptive role. As argued by Sewall Wright in 1932, if a species is characterized by a single large population, it would not be adaptable in the long run. Wright

(1964) noted that most species are divided into numerous small, breeding local populations that are sufficiently isolated to permit differentiation. He argued that this sort of small, semi-isolated subgroup structure was the most effective architecture for evolution by allowing the species to take advantage of ecological opportunities.

Though the work of evolutionary biologists described above pertains to the spread of genes through a population, these ideas have ready application to the spread of ideas through an organization or other community (e.g. Schilling and Phelps, 2007; Fleming and Sorenson 2001). This is aptly illustrated in an example by March (2004). March noted that the international community of organization researchers remains geographically fragmented, with linguistic, cultural, and regional boundaries separating relatively autonomous scholarly communities. He argued that such fragmentation serves an adaptive role in facilitating the resistance of deviant ideas to the homogenizing tendencies of dominant scholarly groups. Fragmentation slows the consolidation of a clear paradigm, and encourages both experimentation and persistence of new ideas. March contends that the parochialism of organizational scholarship helped to protect behavioral ideas about the firm, for example, from premature sacrifice to the hegemony of economic theory (Augier 2003; Williamson 2003).

Recent work on the influence of small-world networks on innovation has also yielded some highly symmetric conclusions: When a community of actors (individuals, firms, etc.) are structured into well defined clusters that are only sparsely connected to each other, this structure can help to create and preserve requisite variety of knowledge in the broader community (Schilling and Phelps, 2007; Uzzi and Spiro, 2005). Densely connected clusters facilitate the local transmission of information by providing multiple pathways between actors, ensuring that information introduced into a cluster will quickly reach other actors in the cluster. Dense clustering can make actors more willing and able to exchange information by engendering trust, reciprocity norms, and shared identity (Ahuja, 2000; Coleman 1988; Dyer and Singh 1998; Granovetter 1992). In addition to stimulating greater transparency, dense clustering may facilitate intense interaction (Uzzi 1997), improving the transfer of tacit, embedded knowledge (Hansen 1999; Zander and Kogut 1995). Thus, clustering enables richer and greater amounts of

information and knowledge to be exchanged and integrated more readily. On the other hand, the sparse connectivity between clusters helps to preserve requisite variety. Though the internal cohesion of a cluster can cause much of the information within the cluster to become homogeneous and redundant (Burt, 1992; Granovetter, 1973), clusters will tend to be heterogeneous across a network in the knowledge they possess and produce due to the different initial conditions and causes for each cluster to form. The links that span those clusters allow the best ideas to migrate throughout the larger community, enabling the heterogeneous ideas to be fruitfully recombined. Small-world properties (high clustering with a few cluster-spanning bridges in an otherwise sparse and decentralized network of actors) thus appear to offer the “best of both worlds” in information diffusion and learning: semi-isolated clusters enable the creation and preservation of heterogeneous ideas, while a few external links help the best ideas to diffuse.

Integrating the lines of reasoning from evolutionary theory and small world networks above with the organizational learning arguments of March suggests that perhaps dividing an organization into subgroups with some barriers to diffusion between them can help the organization maintain the balance between exploration and exploitation. Some amount of isolation of distinct subunits might encourage them to explore diverse solutions, while the connections between those subunits enable superior solutions to eventually be exploited throughout the organization. Thus our first research question is the following: Does the division of an organization into distinct subgroups that are only modestly interconnected lead to superior learning performance outcomes?

2.4. “One Size Fits All” versus Contingency Theory

The argument that any dimension of organizational structure is universally beneficial stands in stark contrast to a long line of research termed “contingency theory.” Dating back to the 1960’s, contingency theorists posited that the optimal organization structure is dependent on such factors as environmental complexity (Burns and Stalker, 1961; Gailbraith, 1973; Lawrence and Lorsch, 1967), organization size (Chandler, 1962; Hickson, Pugh, and Pheysey, 1969), strategy (Chandler, 1962; Child, 1972) or technology (Perrow, 1967; Thompson, 1967; Woodward, 1965). More recent research has supported this perspective by demonstrating that factors such as the nature of organizational knowledge (Birkinshaw,

Nobel and Ridderstrale, 2002; Hansen, 1999) or environmental dynamism (Eisenhardt and Brown, 1999; Lenox, 2002; Schilling and Steensma, 2001) influence the optimal organization structure, though the moniker of “contingency theory” is now often replaced with labels that are more specific to either the contingent factor or the organizational structure (e.g., “modular organizations,” “dynamic tension,” “weak ties”). Most of the research on contingency theory has empirically examined either 1) what structures firms tend to exhibit given other internal or external factors, or 2) performance outcomes in general (e.g., survival, profitability), given a match between a particular structure and other factors. As such, it is difficult to assess the implications of these studies for exploration and exploitation –after all, in the short term, overemphasis on exploitation can be beneficial. Recent studies, however, have begun to examine more directly the relationship between organization structure, environmental factors, and the balance of exploration and exploitation using simulation or computational methods. For instance, Siggelkow and Levinthal (2003) find that a firm that temporarily decentralizes its search activities into two subunits and then reintegrates is less likely to become trapped on a suboptimal peak of an NK landscape than a firm that remains centralized for all search activities. They also find, however, that the exploration enabled by the decentralization can be costly (sometimes the decentralized firm is led astray), and is more likely to be valuable for firms that face significant environmental change than firms that are in stable environments. Similarly, Siggelkow and Rivkin (2005) find that the value of decentralization and the optimal amount of integration via a department head is contingent upon both the amount of environmental turbulence and complexity faced by the firm.

Integrating the research described above with the original findings by March suggests that the relationship between organizational structure and organizational learning may be dependent on factors such as the degree of complexity and dynamism in the firm’s internal or external environment, or the degree of personnel turnover. Thus we are led to our second research question: Do factors such as problem complexity, environmental turbulence, or personnel turnover determine whether dividing the firm into semi-isolated subunits has a beneficial effect on learning outcomes?

Note that the arguments described above are inconsistent. The line of reasoning from evolutionary theory and small-world networks described above imply that a semi-isolated subgroup structure will be universally beneficial for improving the balance between exploration and exploitation. The work on contingency theory and research on search in a NK landscape, described in this section, however, suggest that the benefit of a particular organization structure is dependent on a range of other internal or external factors. To systematically explore these issues, we create a simulation of organizational learning that extends March's original work by incorporating interpersonal learning. This model permits us to treat the isolation of subgroups as a tunable parameter, and thereby systematically explore how the degree of isolation influences organizational learning outcomes. We also examine whether and how these outcomes vary depending on subgroup size, problem complexity, environmental dynamism, and personnel turnover.

3. MODEL

We regard an organization as a complex adaptive system, where individuals interact with other individuals. In particular, we view individuals as carriers of ideas and knowledge, and organizational learning as a property that emerges from interactions among individuals in the organization. Individuals interact with others, who may influence them to adopt new ideas and to discard old ideas. This assumption of interpersonal learning makes our paper distinct from March's (1991) work, and builds on recent work by Miller, Zhao and Callantone (2006).² We model organizational learning as a process by which individuals within an organization interact to exchange and jointly create knowledge. Our model has three main entities: *an external reality*, *individuals*, and *an organization*.

1. **External Reality.** Like March (1991), we describe reality as having m dimensions, each of which has a value of 1 or -1. The probability that any one dimension will have a value of 1 (or -1) is 0.5; values are randomly assigned.

² In March's (1991) model, individuals in an organization do not learn from each other. Instead, there is "mutual learning" between individuals and the organizational code. The learning by and from the organizational code is meant to represent a socialization process whereby an organization socializes recruits to its language, beliefs, and practices, while the code simultaneously adapts to the beliefs of individual that have superior performance. Our adaptation does not dismiss this socialization process, but rather utilizes recent advances in simulation tools to permit socialization and learning to occur in a manner that we believe is more realistic—through interpersonal interaction.

2. **Individual.** There are n individuals in an organization. Each of them holds m beliefs about the corresponding elements of reality at each time step. Each belief for an individual has a value of 1, 0, or -1. A value of 0 means that an individual is not sure of whether 1 or -1 represents reality.

3. **Organization.** As mentioned earlier, our model is different from the March model (1991) in that an organization is seen as a complex system wherein individuals directly interact with one another. A key factor determining the dynamics of such a system is its topology of interaction patterns between members (Strogatz, 2001). In the past, it was not easy to study such complex systems primarily due to the lack of appropriate tools for analyzing their topologies (Lee et al. 2006). However, recent advances in graph theory—especially a quantum advance in our understanding of networks—have offered tools for studying complex social (as well as non-social) networks.

To examine the role of semi-isolated subgroups, we modify Watts’ ‘connected caveman’ model (1999), which represents interaction patterns among members of an organization (see Figure 1, panel a).³ In this stylized model, the network consists of highly clustered subgroups that are minimally connected together. Such a network is highly “clustered,” meaning that the fraction of neighboring nodes that are also connected to one another is very high. Within this structure, subgroups (or “clusters”) are “nearly isolated” because subgroups have a *minimum* amount of connectivity to the rest of the network. For instance, in the connected caveman graph depicted in Figure 1a, clusters have only two links connecting them in serial fashion to one another.

To construct different interaction patterns, we rewire each link of this connected caveman network by making random connections between nodes with rewiring probability β .⁴ When $\beta = 0$, the connected

³ We begin with a connected model as our baseline for two important reasons. First, our intention is to model organizational learning, which necessitates that our subgroups be somehow connected under an organizing structure. If they were disconnected, the mapping to an overarching organization would be lost. Second, the dynamic properties of isolated networks differ drastically from connected networks. If groups are isolated from one another, information cannot diffuse between them, rendering the usual subjects of interest in diffusion studies (e.g., infection rate, path length, etc.) uninteresting. Thus from a graph theoretical viewpoint, isolated networks are qualitatively quite distinct from connected networks. Our minor modification to Watts’ caveman model ensures that each individual is initially connected to the same number of direct neighbors.

⁴ In network studies it is conventional to use random rewiring when there is no ex-ante reason to use a pre-specified connection pattern. The use of random rewiring does not assume that interaction patterns within an organization are random; it only assumes that they are indeterminate.

caveman structure is preserved. When β is non-zero but sufficiently small, the high degree of clustering is preserved, but there are a few random connections among subgroups (as in Figure 1b). This may be interpreted as “semi-isolation” where subgroups have a *moderate* amount of connectivity to the network yet remain well-defined and separate enough from the rest of the network to have a clear subgroup identity. This semi-isolated subgroup structure reveals two important topological properties of small-world networks: small degrees of separation and high clustering. With only a few random links, or shortcuts, between the subgroups, the average degree of separation between any two nodes drops rapidly while the degree of clustering remains high, helping to preserve subgroup identity. When β grows large and approaches unity, a network becomes completely random, as shown in Figure 1, panel c. Shortcuts are ubiquitous, but the subgroup structure is lost.⁵ Note that the two key ingredients of our model, cliquish subgroups and cross-group links (or bridges), are also the basic building blocks of Granovetter’s (1973) conceptualization of social networks, which were empirically confirmed by Onnela et al. (2007).

<Insert Figure 1 about here>

3.1. Learning

When individuals first join an organization, they have idiosyncratic sets of beliefs that are heterogeneous across individuals. However, as these individuals regularly interact with their contacts within the firm, they will compare the performance of their own belief sets to those of their contacts. In our model, learning is a bi-directional process. That is, a link between two individuals indicates that both are sources of learning for each other. Similar to March’s original work, individuals can observe the payoff of their overall belief set (and that of their direct contacts), but they cannot directly observe how each element of the belief set contributes to this payoff. If individuals ascertain that others have belief sets that are better performing, they may update their own beliefs to incorporate aspects of the higher performing belief sets. The process of sharing, comparing, interpreting, and updating beliefs helps groups

⁵ β can be interpreted as the fraction of additional cross-group links (i.e., beyond those initially used to maintain the network’s minimal connectivity) available in an organizational network. The larger the value of β , the more cross-group links there are in the network, and the shorter the degrees of separation among organizational members. In addition, as a result of random rewiring for $\beta > 0$, the number of neighboring nodes for each node varies. This variation in the number of links can be approximated by a Poisson distribution.

of individuals develop a shared understanding (Argyris and Schon 1996; Brown and Duguid 1991; Crossan, Lane and White 1999; Daft and Weick 1984; Huber 1991, Klimoski and Mohammed 1994; Larson and Christensen 1993; Sandelands and Stablein 1987). In this way, the beliefs about reality possessed by individuals within the organization tend to become more homogeneous over time.

Though individuals may be able to determine if other individuals have belief sets that are better performing than their own, they might not be able to determine which aspects of a particular belief set led to better performance. That is, while an individual may believe that another's individual's belief set is more complete or correct than their own, they may not know *how* it is more correct or complete. This is compounded by the fact that an individual's better-performing peers may not have identical belief sets. These factors cause an individual to face ambiguity about whether and how to update their own beliefs. To address this, we adopt a majority decision rule similar to that used in the March (1991) study. Individuals look at all of the other individuals with which they interact, and identify those that are better performing than themselves. The focal individual then identifies what the dominant belief (the majority view) is on each of m dimensions of the belief sets of these higher performing peers. She may then decide to update each dimension of her own belief set to this majority view with some probability $p_{learning}$ that reflects the proclivity or ability of individuals to learn from one another. The majority decision rule helps to simulate a socialization process within the organization whereby an individual's beliefs are influenced by the collective beliefs of their peers. This decision rule is also consistent with a large body of research on social decision schemes and has been supported by numerous studies of social decision making (e.g., Castore, Peterson and Goodrich 1971; Davis et al 1975). It is particularly likely to be used by groups when it is difficult for individuals to assess which alternative is correct (Kameda and Davis, 1990).

3.2. Payoff Function and Problem Complexity

The performance of an individual is evaluated at every time step by a given payoff function. We construct a generalized m/s payoff function where m refers to the dimensionality of the search problem and s is a single parameter with which the interdependence of the search problem can be tuned. By increasing the value of s , the search problem becomes more interdependent. That is, the performance will

not improve unless several beliefs jointly match corresponding parts of reality. For example, consider a case where $s = 5$ for a search problem with five dimensions (i.e., $m = 5$). If the individual identifies the correct belief for each of the five dimensions of the belief set, her payoff is 5. But, when an individual has four correct beliefs and one incorrect one, her payoff is zero. Thus if $s = m$, the search problem is maximally interdependent; if any single element among the m beliefs is wrong, the payoff for the whole set becomes zero. Alternatively, when s is small, the dimensions of the search problem are more independent. If $s = 1$, when an individual has four correct beliefs and one incorrect one, her payoff is four. The bigger the value of s , the more interdependent the search problem is. Note that our m/s payoff function has a generalized form that captures both the linear payoff function (i.e. $s = 1$) used by March (1991) and the needle-in-the-hay-stack problem (i.e. $s = m$) used by Hinton and Nowlan (1987) (considered one of the toughest search problems) as special cases. In our baseline model, we keep $s = 5$ for a search problem with one hundred dimensions (i.e., moderately low interdependence). The details of this payoff function are described in Appendix A. An organization's performance is measured as the average performance across all individuals in the organization.

Our m/s payoff function builds directly from March's (1991) original model, though it also shares some similarities with the NK model. In particular, our parameters m and s are analogous to N and K in the NK model, where N represents the number of decisions and K represents their interdependence for their contributions to the overall payoff (e.g., Kauffman 1993; Levinthal 1997; Rivkin, 2000). The main distinction of our payoff function lies in its computational efficiency, especially when one deals with high dimensional search problems (i.e., a high value of N) with a substantial degree of interaction (a high value of K). In such a setting, the NK model becomes prohibitively expensive, requiring huge memory and extensive computational time for running a simulation—this problem is known as a “NP hard” problem in computer science. In our set up, the cost of increasing dimension m and interdependence s in our payoff function is much lower, which enables the simulation to run faster.

4. SIMULATION RESULTS

We use the model described above to run a series of simulations that examine the impact of interpersonal network structure on organizational-level performance. We first explore the impact of organization structure on learning and performance in a baseline model and then examine the influence of problem complexity, environmental turbulence, and personnel turnover in an extended model.

4.1. Baseline Model

To study how interpersonal network structure affects organizational learning in our baseline model, we first consider an extreme type of interpersonal network structure, the “nearly-isolated subgroup structure,” which is illustrated by Figure 1a. We then vary the form of subgroup structure by randomly rewiring some of the existing links in this structure with probability β . The bigger the value of β , the greater the percentage of random, cross-group links. For example, when $\beta = 0.1$, we have a “semi-isolated subgroup structure” with about 10% of random, cross-group links on average. When $\beta = 1$, subgroup identity disappears due to the ubiquity of cross-group links. By examining the learning performance of each of these structures we address the following questions: Does subgroup structure influence long-run learning outcomes? Can we improve the balance between exploration and exploitation by varying the degree of subgroup structure within an organization?

Effects of Isolation

In March’s model, learning rates for organizational members control the speed of organizational learning. Since our key variable of interest is organization structure, we keep the learning rate constant at $p_{learning} = 0.3$ for the moment – we relax this assumption later in our sensitivity analysis. All parameters used in the simulations and the simulation steps are specified in Appendix B. Figure 2 plots the averages, 10th percentile, and 90th percentile of the equilibrium performance of 200 simulation runs, for various levels of the structural parameter β , which can be roughly interpreted as the fraction of additional cross-group links (i.e., beyond those initially used to maintain the network’s minimal connectivity). As shown, the long-run learning performance is lower when $\beta = 0$, or when subgroups are nearly isolated. It appears

to be highest when subgroups are “semi-isolated” with a small fraction of random cross-group links. The peak performance is associated with $\beta \in [0.05, 0.1]$. As β approaches 1, the long-run average outcome of learning tends to decrease. Figure 2 also shows that the variance of observations (i.e., the spread between the 10th percentile and 90th percentile lines) for $\beta = 0$ covers a much broader range than those of the other distributions. This implies that more extreme behaviors are likely to be observed when $\beta = 0$. For example, the probability of observing performance levels of 90 or 30 is non-zero even though the mean is 63.6. This can be interpreted as suggesting that organizational risk will tend to be higher when solutions to problems are pursued only by highly autonomous teams with few mechanisms for cross-team learning. For $\beta \geq 0.1$, on the other hand, the variance of observations is much smaller and roughly equal across different values of β . This means that extreme behaviors are less likely.

<Insert Figure 2 about here>

As seen in Figure 2, when cross-group linking increases above $\beta \geq 0.1$, long-run performance tends to decrease. We conjecture that this decrease is driven by the loss of diversity in beliefs in an organization. As the number of cross-group links increases, good ideas spread faster throughout the organization. This may prematurely eliminate deviant ideas, reducing the combinatorial possibilities and thereby lowering the organization’s likelihood of achieving its best possible outcome. To confirm this conjecture, we develop a measure of organizational diversity, or what we call a “dissimilarity index.” To construct such a measure, we make pair-wise comparisons of all n individuals. There are $\frac{1}{2}n(n-1)$ pairs. For each pair of individuals, there are m beliefs to be compared. Then, we measure dissimilarity in beliefs of all individuals in an organization as follows:

$$\text{Dissimilarity} = \frac{2}{mn(n-1)} \sum_{i=1}^{\frac{1}{2}n(n-1)} \sum_{j=1}^m \omega_{ij} \quad (1)$$

where ω_{ij} takes on the value of 1 if two chosen individuals for i th pair have different beliefs on j th dimension, and 0 otherwise. Figure 3 shows the degree of dissimilarity of beliefs among organizational

members over time for different levels of β . As shown, an organization tends to lose its diversity of beliefs very fast when its interpersonal network is characterized by a large fraction of cross-group links (e.g., $\beta = 1$). In essence, as subgroup identity is lost, dissimilarity among organizational members quickly disappears. In our baseline model without turnover, once organizational diversity is lost, there is no way for an organization to promote it. The management of such diversity is thus crucial for improving the organization's long-run performance. By maintaining diversity among individuals longer, a semi-isolated subgroup structure (e.g., $\beta = 0.1$) achieves higher long-run performance.⁶

<Insert Figure 3 about here>

Figure 3 also helps us understand why long-run performance tends to be lower for a nearly isolated network ($\beta = 0$). As shown, dissimilarity among individuals decays very slowly in this network. Diverse ideas and beliefs among different subgroups are thus well maintained. Why, then, do we systematically observe the lower performance in this structure? In our model, learning from other superior individuals is key to improving the organization's performance. The problem with the nearly isolated structure is that diverse ideas cannot be exchanged easily across subgroups due to the limited availability of cross-group channels. In other words, maintaining organizational diversity is not beneficial unless the organization's members can exchange diverse ideas across subgroups to improve their knowledge. With very few cross-group links that act as bridges between different subgroups, good ideas cannot be effectively shared and exchanged. As such, even if heterogeneity is preserved, organizations cannot benefit from this diversity if there are too few cross-group links. Would the long-run outcome of organizational learning be worse if subgroups were completely isolated from one another? In numerical analyses not reported here, we observe that the average learning outcome under this setting is close to zero, indicating that complete isolation makes it very hard for an organization to improve its knowledge.

In sum, we find that subgroup structure, which fosters the diversity of ideas or beliefs within an organization, is conducive to learning as long as there is a small fraction of random, cross-group links.

⁶ Figure 2 indicates that this finding is not parameter-specific, but rather that there is a broad interval of β over which the long-run outcome is relatively high. We thank James March for bringing this issue to our attention.

Such semi-isolation allows the organization to learn moderately fast with the highest long-run performance outcome thus improves the balance between learning speed and long-run performance.

Effects of Subgroup Size

Wright's arguments about population structure also emphasized subgroup size: only in *small* semi-isolated subgroups could seemingly non-adaptive traits become preserved. To examine this aspect of Wright's argument we specify a model parameter z that corresponds to the size of each sub-group. We then vary the size of the subgroups from $z = 7$, to $z = 14, 28, 70$, and 140 ⁷, while fixing n .⁸ Figure 4 plots the effects of different subgroup sizes on learning. First, we find that an increase in the size of subgroups lowers learning performance across all forms of organizational structure (i.e. different β s). As the size of subgroups grows bigger, divergent beliefs of individual members are more quickly homogenized, lessening the potential for long-run learning. In results not reported here, we also find that a completely connected network (where every individual is directly connected to every other individual) has a performance of 36.7, which is much lower than those in the sparse networks.

Second, we find that there is little change to our basic pattern of performance across various levels of β s. This indicates that the semi-isolated subgroup structure remains the highest performing structure irrespective of subgroup size. Overall, then, these results are highly consistent with Wright's arguments

⁷ The group sizes we use for our starting points are based on realistic group sizes used in organizations. Devine et al. (1999) report that the average size of new product development teams in the United States is 11 members. The Saratoga Institute's 2001 benchmarking study (*Saratoga Institute Workforce Diagnostic System Benchmarking Reports 2001*) find that the average span of control of supervisors in US companies ranged from four employees per manager (for information services) to 16 employees for healthcare organizations. Needless to say, if the subgroup becomes very large (e.g., approaches the size of the organization), the results approach those that we obtained for a completely connected organization.

⁸ Given that we keep the number of organizational members constant, note that z could also be interpreted as the average connectivity of the organization. Because of the nature of connected caveman graph, an increase in the size of subgroups also increases the average number of links per person. For example, when $z = 7$, the average number of connections per person is 6. When $z = 14, 28, 70$, and 140 , the average number increases to 13, 27, 69, and 139 respectively. If we increase the size of the subgroups but constrain the number of links individuals have within them to something less than the size of the subgroup, we end up replicating results we have already produced. If, for example, we permit each individual to have six links and the links are assigned in a cliquish manner (as would be reasonable given what we know about social processes such as homophily, etc.), this ends up being structurally identical to an organization with 40 subgroups with seven individuals each, with no connectivity or minimal connectivity between them depending on how we set up the initial cliques. If we assign the six links within a large subgroup randomly, this ends up being structurally nearly identical to an organization with 40 subgroups of seven individuals each with relatively high beta.

about effective architecture for evolution: dividing the community into subgroups that are both small and semi-isolated increases the long-run adaptive performance of a community.

<Insert Figure 4 about here>

How the Learning Process Unfolds

To better understand the equilibrium performance results presented above, it is useful to observe the micro-dynamics of the learning process as it unfolds in our simulations. In our simulated organizations, individuals are initially assigned sets of beliefs that are randomly generated, and thus heterogeneous. They then begin to learn from others with which they interact as defined by the underlying interpersonal network. In a clustered organization with well-defined subgroups (i.e., low levels of β), the combination of the clustered structure with a learning mechanism based on majority view tends to lead to the emergence of pockets of relatively homogeneous beliefs (i.e., sets of proximate individuals with highly similar belief strings). Since most of an individual's contacts are also contacts of each other (due to the highly clustered structure), when the individual identifies those of its contacts that are higher performing than itself, that set of higher performing peers will be very similar to the set of higher performing peers that its contacts identify. As each individual looks for the majority view among these high performing peers and adopts their beliefs with probability $p_{learning}$, individuals within a cluster will tend to converge on a common set of beliefs. However, because each cluster is founded with a different group of individuals, each with their own heterogeneous stock of initial beliefs, the common belief sets that emerge within clusters may be different *across* clusters.

Because the clusters are connected to each other (at varying degrees, depending on β), some individuals within the cluster have contacts outside of the cluster. As such, they may identify high performers who have a very different set of beliefs from the sets in their cluster. However, because of the majority view rule, the individual will not learn from high performing "outsiders" until she has a performance that equals or exceeds most of the individuals in her cluster. Only when the number of higher performing "outsiders" meets or exceeds the number of higher performing individuals in her

cluster can the higher performing “outsiders” become the majority or tied for majority. Once the higher performing “outsiders” become the majority or tied for majority, the individual is able to learn from these higher performing individuals. Thus better performing belief sets can diffuse across clusters, but only when they are not competing with a greater number of within-cluster superior peers. Furthermore, if the belief set possessed by the outsider is very different from the focal individual’s belief set, the recombinatorial process can potentially yield a belief set that is quite different (and potentially better performing) than the belief set possessed by *either* individual prior to learning from one another. The preceding highlights the role that the majority view socialization process plays in the learning dynamics; without such a socialization process, subgroup structure may have much less effect on performance. We elaborate on this in Appendix C.

4.2. Extended Model: Testing a Contingency Perspective

As discussed previously, the traditional contingency logic argues that the appropriateness of any organizational design is a direct function of environmental conditions (Burns and Stalker, 1961; Lawrence and Lorsch, 1967). In the previous section, we explore the effect of subgroup structure given a fixed amount of problem complexity and zero environmental turbulence or personnel turnover. We now relax those assumptions to examine whether the dynamics we observed previously change with these contingencies. First, we look at the impact of varying problem complexity by altering both the number of elements (m) that must be considered in the organization’s search, and the degree to which those elements are interdependent (s). Second, we examine the impact of environmental change and personnel turnover.

Effects of Problem Complexity. In the previous analysis, we kept search problem complexity constant by setting parameters m and s to 5 and 100 respectively. However, previous research suggests that problem complexity may be an important contingency variable in organization design. Complex problems with a high degree of interdependency create rugged landscapes with multiple local peaks, and as a result, organizations may become easily stuck (Levinthal, 1997; Rivkin, 2001). In such environments, the organization may benefit from a structure that enables more exploratory search. By contrast, in an environment characterized by simple search problems, it might be the case that most organizational

designs perform satisfactorily, or that designs that encourage exploration of diverse solution spaces are unnecessarily costly, resulting in lower performance than designs that encourage efficient exploitation.

To test this possible contingency, we next vary parameters m ($m = 40, 70, 100, 130, 160$) and s ($s = 1, 2, 3, 5, 6, 7, 8, 9, 10$) jointly. As specified before, m stands for the number of elements that an organization should search for, whereas s represents their interdependence. In results not reported here, we find that the superiority of the semi-isolated subgroup structure holds so long as search problems are moderately difficult (e.g., $2 < s < 10$) and beyond this range there is little variation in performance outcomes. For example, when $s = 1$ or 2 , performance heterogeneity across different structures disappears. All types of organizational structures enjoy the same, maximum level of performance (at around 100, i.e., getting all 100 beliefs correct). In other words, when a search problem is too easy, any type of structure can do well in learning. When $s \geq 10$, or when a search problem becomes too difficult with substantial interdependence among search elements, performance heterogeneity across different structures again disappears. For instance, when $s = 10$, an organization needs to get 10 beliefs correct simultaneously to get any positive performance. Even if it is correct on 9 out of 10 dimensions, its performance is still 0. Not surprisingly, all organizations completely fail to learn and their performances hover around 0. The near zero performance for $s \geq 10$ indicates that an organization or a system cannot evolve when search complexity is beyond a certain threshold.⁹ Research on complex systems suggests that many real-world problems fall into a moderate range of complexity (Kauffman 1995, Waldrop 1992). Given this logic, we fix the value of this parameter to five for the rest of our analysis. Also, we find that m does not greatly affect the pattern we observed in Figure 2. The superiority of a semi-isolated organizational structure is robust to changes in m , as long as the interdependence of search problems is moderate (i.e., $s = 5$). The superiority of the semi-isolated subgroup holds except for very small values of

⁹ This result is reminiscent of the non-evolvability condition, which basically means that systems cannot evolve when search complexity is beyond a certain threshold (Kauffman 1993; 1995). Research on complex systems has indicated that many real-world problems are neither too easy nor too hard, and that the degree of interdependence is rather moderate (Kauffman 1995, Waldrop 1992). Our finding is also highly consistent with prior studies in the management area (Rivkin, 2001; Schoemaker, 1990). For instance, using the NK framework, Rivkin (2001) showed that only in the case of moderate complexity, does the wedge between innovators and imitators exist.

m . In addition, an increase in m seems to make the performance heterogeneity more pronounced because search becomes harder.

Effects of Environmental Change and Turnover. In the previous section, we imposed the assumptions of no environmental turbulence or personnel turnover. Here, we introduce two modifications, while keeping the other parameters the same as before (i.e., $m = 100$, $n = 280$, $z = 7$, $p_{learning} = 0.3$ and $s = 5$). First, we allow personnel turnover to take place: each individual will leave the organization at each time period with probability $p_{turnover}$. When an individual leaves, her position is filled by a new member who starts with randomly assigned beliefs (a value for each of m beliefs is drawn randomly from 1, 0, or -1).¹⁰ Second, we let the environment change periodically. In particular, we perturb each of m dimensions of reality (from 1 to -1, or from -1 to 1) with probability p_{envir} at every T period. For example, we set $T = 200$ and $p_{envir} = 0.1$. That is, in every 200 period, about 10% of m elements in reality change their values. The value of T is tuned to allow sufficient time for the organizations to learn and adapt to new environment. The implicit assumption is that the timescale for environment change is greater than that for organizational learning.¹¹

Figure 7 shows a typical run of the simulation with periodic environmental changes with and without turnover, while holding β constant at 0.1. As shown, after an environmental shift, the organization suffers an immediate drop in performance. However, with personnel turnover, the organization is able to recover from the environmental shift and improve its performance. Without turnover, over time the organization loses its ability to adapt to environmental shifts. Consistent with March (1991), we find that turnover serves a functional role by preserving diversity in the organization.

<Insert Figure 5 about here>

To see whether a semi-isolated structure continues to perform better in the open system, we run 100 simulation runs for each structure ($\beta = 0$, $\beta = 0.1$, and $\beta = 1$) and vary a) the turnover rate and b) the extent of environmental change. The details of these simulation experiments are shown in Table 1.

¹⁰ We find that turnover is detrimental to long-run performance in a stable environment, consistent with prior work (e.g., Argote, et al, 1995; Denrell, Fang and Levinthal, 2004; Ton and Huckman, 2008).

¹¹ We thank Michael Katz for bringing this to our attention.

We find that for each level of turnover rate or environmental change, the learning outcomes are always highest when β equals 0.1, suggesting that the benefits of the semi-isolated subgroup structure are robust to these contingencies. Furthermore, in results not reported here, we find that modest amounts of personnel turnover (when the turnover rate is 0.01%) are consistently beneficial across a range of values for subgroup structure or environmental dynamism. This modest amount of turnover helps to inject the organization with novel ideas and replenish lost diversity necessary to adapt to a changing environment. On the other hand, at turnover rates of 0.1 or above, learning performance drops substantially. This finding is equivalent to what evolutionary biologists call “error catastrophe” (e.g., Kauffman 1995) where too high a mutation rate makes the system too disruptive to evolve.

4.3. Sensitivity Analysis

First, we vary the size of an organization ($n = 140, 210, 350,$ and 420) and find that there is no substantive change in our results. Second, we vary the individual learning rate to see if it affected our main findings. When learning rates are sufficiently small (e.g., $p_{learning}$ is smaller than 0.7), modest amounts of cross-group linking always yield the highest long-run performance. When learning rates are high (e.g., $p_{learning} = 0.7, 0.8,$ and 0.9), however, it becomes harder to trace this pattern as performance heterogeneity among different network structures becomes less pronounced. Furthermore, consistent with March’s (1991), we find that the average outcomes of learning decrease with these higher values of learning rates.¹²

5. DISCUSSION

In this study, we use a series of simulations to examine the influence of interpersonal network structure on organizational-level learning. We focus our analysis on the extent of isolation of small

¹² We also implement a version of our program that closely resembles a formal hierarchical organization. We keep the original set up the same (i.e. same number of nodes, same number of dimensions etc), except that we organize sub-groups hierarchically. Instead of having seven subgroups in a flat organization, we organize the subgroups into four levels (with one subgroup randomly chosen to represent the top management team and two subgroups for each of the lower levels). Once this structure is set up, we again rewire existing links probabilistically according to the value of β . In results not reported here, we find that the general pattern is the same – the highest performance is reached when the probability of rewiring is around 0.1, which corresponds to a semi-isolated subgroup structure. Notably, we do not introduce differential authority or power into our model, which would be an interesting area for future research.

subgroups by varying the degree of cross-group links. We find that modest amounts of cross-group linking (i.e. low but non-zero levels) are associated with higher equilibrium performance levels. The semi-isolation of subgroups helps to preserve heterogeneity of ideas and beliefs in the organizations, permitting organizations to explore a wider range of ideas and belief sets, while the small number of random cross-group links permits the best performing ideas and belief sets to diffuse throughout the organizations. Our study also finds, rather surprisingly, that the semi-isolated subgroup structure achieves superior long-run learning outcomes across a wide range of parameter conditions, including contingencies such as varying problem complexity, environmental change, and personnel turnover. The role of subgroup structure, however, appears to be influenced by our socialized learning process (i.e., majority decision rule). Furthermore, performance variation across different subgroup structures is only observable for problems of moderate complexity – if problems are too simple every organization easily achieves the maximum score and if problems are too difficult, every organization performs poorly.

5.1. Structure as a Mechanism for Balancing Exploration and Exploitation

As noted previously, prior research has found that organizations tend to emphasize exploitation at the expense of exploration and that organization structure may help to attenuate this problem. Our study provides additional insight into this stream of research by demonstrating how subgroup structure influences learning processes within the organization. In our simulated organizational context, *parallel, isolated learning within each subgroup* enables exploration, and the preservation of variety of knowledge in an organization. *Learning across subgroups* enables exploitation by facilitating the rapid diffusion and assimilation of currently superior knowledge. Our study thus suggests that a productive balance between exploration and exploitation can be achieved by breaking an organization down into small semi-autonomous subunits with a small fraction of cross-group links.

Our results resonate with a growing body of research on network structure (Bell and Zaheer, 2007; Fleming, King and Juda, 2007; Hansen 1999, 2002; Reagans and McEvily 2003; Schilling and Phelps 2007; Uzzi 1997; Uzzi and Spiro 2005). Research on social networks argues that dense clustering provides trust and reciprocity norms, while also facilitating intense interaction (Coleman, 1988, Uzzi,

1997). On the other hand, weak ties (i.e., those that extend beyond an actor's group of densely connected affiliations) create advantages for actors by providing access to resources that are different from those found in the actor's more immediate social network (Granovetter 1973; Hansen 1999). Similarly, studies of inter-firm networks suggest that the redundant connections created by clustering help to increase the information transmission capacity of a network, while connections across clusters help to increase the scope of new information that can be accessed. Thus networks that have both clustering and some amount of random linking between them are valuable for creativity and innovation (Schilling and Phelps 2007). Both the social network and inter-firm network literatures thus emphasize the local transmission advantages of clustering, and the information scope advantages of cross-cluster connections. Our study, however, highlights another key aspect of this dynamic: it is the *preservation of diverse ideas within small groups* that increases the information scope of the overall network. In essence, the subgroups shelter heterodox ideas, enabling them to survive and be refined, rather than quickly extinguished through competition in the larger population.¹³

Our central finding above is also consistent with the literature on technology management. This literature has suggested that it is beneficial to isolate new product development (NPD) teams, at least temporarily, from each other or from the rest of the organization (e.g., Bower and Christensen 1995). Isolating NPD teams and giving them considerable autonomy allows them to pursue new technological possibilities, unfettered from existing organizational paradigms. This is especially important when teams are working on disruptive innovations whose features appear inconsistent with current customer requirements. On the other hand, our results also suggest that teams should not be completely isolated. A modest degree of connection between NPD teams (or between NPD teams and other divisions of the firm) is important to enable the leveraging of ideas across the organization, fostering the identification of

¹³ It is important to note, however, that even though a semi-isolated structure may be optimal, individuals within this structure may see opportunities to create bridging ties that will enhance their personal performances. Incentives in real organizations may not be well aligned.

valuable synergies. Small-world connectivity of this sort can be achieved by creating inter-team liaison roles or utilizing practices such as personnel rotation or inter-divisional task forces.¹⁴

5.2. Resolution to Inconsistent Views

The evolutionary work by Wright (and others) has argued that certain structural designs have selective advantage under a wide range of environmental settings. Research on small-world networks has suggested a similar point, arguing that particular network structures have inherent advantages that yield knowledge creation benefits irrespective of many of the finer details of the system or its context. The ubiquity of small-world structures in diverse physical and social systems has provided empirical evidence that strengthens this position (Watts, 1999). Management research, however, has long emphasized a contingency perspective that asserts that contextual or environmental conditions almost always moderate relationships between organizational design and organizational performance. This body of research would tend to resist the notion that any structural dimension would be *universally* advantageous.

Our evidence is more or less consistent with the evolutionary/small-world perspective, while not completely ruling out the claims of the contingency perspective. The semi-isolated subgroup structure achieves superior performance despite a wide range of contingency variations we explore such as problem complexity, environmental change, or personnel turnover. The benefits of a semi-isolated structure appear to be remarkably robust when the search problem is of moderate complexity ($2 < s < 10$). However, this robustness may be bounded to the “majority decision” rule that we adopted and modified from March’s original study. It appears that this rule plays a pivotal role in the isolation and preservation of diverse ideas in different subgroups. Because the majority of an individual’s contacts are within their subgroup, the majority decision rule causes individuals to initially be primarily influenced by members of their subgroup. A contact outside the subgroup only becomes the “majority” when there are no superior peers within the subgroup. Needless to say, in an environment where an individual’s learning is less influenced by the collective beliefs of their peers, subgroup structure is likely to matter far less. While one study

¹⁴ Our results also echo several recent simulation papers which study learning and search using different models (e.g., Davis, Eisenhardt and Bingham, 2008; Winter, Cattani and Dorsch, 2007).

clearly cannot provide sufficient evidence to assert that one particular structural dimension or archetype is universally beneficial, we believe that this might suggest an important area for future research. We are hopeful that future studies will continue to explore whether the learning benefits of the semi-isolated subgroup structure are bounded to particular contextual settings.

5.3. Methodological Innovations

Our work offers a set of methodological innovations for future research. First, we develop a model of organizational learning that departs substantially from the original March model (1991). As we noted before, March (1991) sidesteps the structural complexity inherent in interactions among organizational members (i.e., individual members are not assumed to learn from each other) in order to highlight the dynamical complexity of organizational learning. We tackle the structural complexity of interpersonal learning by building on and modifying Watt's caveman model (Watts 1999). By modeling interpersonal learning via connected caveman networks with a tunable rewiring parameter, we expose an untapped link between evolutionary theory and the burgeoning literature on small-world networks.

Our second methodological contribution lies in the development of a generalized m/s payoff function which incorporates some well-known payoff functions (e.g., March's linear payoff function) as special cases. Like the NK model, our payoff function produces the complexity of a search problem that arises not only from the number of elements to search but also from their interdependence. The complexity parameter in our generalized payoff function (s) effectively tunes the level of interdependency (i.e., K in the NK model). In this sense, our payoff structure is similar to the NK model. Compared to it, however, our payoff function dramatically reduces the computational burden of running a computer simulation. Though NK models have become popular for management simulations, they are not typically used in computer science because the computational burden of an NK model increases exponentially with dimensions. Our baseline model studies learning in a high dimensional problem ($m = 100$); perhaps when computing power renders it feasible, one could replicate this study with an NK model, though we conjecture that the main results would not change much.

To conclude, our paper provides an optimistic answer to the challenge of balancing the opposing forces of exploration and exploitation, a key concern in organizational learning, adaptation and performance. We show that there exists a particular structural design that may help thwart the self-destructive tendencies of adaptive processes in organizations. While our finding is more consistent with evolutionary perspectives and the recent literature on small world networks, we do not believe we have the 'final word'. Future research should continue to explore the impact of contextual and contingent variables on structure and we hope that the work we present here represents an important stepping stone to uncovering more nuanced underpinnings between structure and performance.

REFERENCES

- Ahuja G. 2000. Collaboration networks, structural holes, and innovation: A longitudinal study. *Admin. Sci. Quart.* **45** 425-455.
- Argote, L., H. Greve. 2007. A behavioral theory of the firm—40 years and counting: introduction and impact. *Org. Sci.* **18** 337-349.
- Argote, L., C. Insko., N. Yovetich., A. Romero. 1995. Group learning curves: The effects of turnover and task complexity on group performance. *J. Appl. Social Psych.* **25** 512–529.
- Argote, L., B. McEvily., R. Reagans. 2003. Managing knowledge in organizations: An integrative framework and review of emerging themes. *Management Sci.* 49(4), 571-582.
- Argyris, C., D. Schon. 1996. *Organizational Learning II: Theory, Method, and Practice*. Addison Wesley, Reading, MA.
- Augier, M. 2003. *The evolution of behavioral economics*. Unpublished Manuscript, Stanford University.
- Bell, G., A. Zaheer. 2007. Geography, networks, and knowledge flow *Org. Sci.* **18** 955-972.
- Benner, M. J., M. Tushman. 2003. Exploitation, exploitation, and process management: The productivity dilemma revisited. *Acad. Management Rev.* **28** 238-256.
- Birkinshaw, J., R. Noble, J. Ridderstrale. 2002. Knowledge as contingency variable: Do the characteristics of knowledge predict organization structure? *Org. Sci.* **13**(3) 274–299.
- Bower, J. L., C. M. Christensen. 1995. Disruptive technologies: catching the wave. *Harvard Bus. Rev.* **73**(1) 43-53.
- Brown, J., P. Duguid. 1991. Organizational learning and communities of practice: Towards a unified view of working, learning and innovation. *Org. Sci.* **2** 40-57.
- Burns, T., G. M. Stalker. 1961. *The Management of Innovation*. Tavistock, London, UK.
- Burt, R. 1992. *Structural Holes: The Social Structure of Competition*, Harvard University Press, Cambridge, MA.
- Castore, C. H., K. Peterson, T. A. Goodrich. 1971. Risky shift: Social value or social choice? An alternative model. *J. Personality Soc. Psych.* **20** 487-494.
- Chandler, A. 1962. *Strategy and Structure: Chapters in the History of Industrial Enterprise*. MIT Press, Cambridge, MA.
- Child, J. 1972. Organization structure, environment and performance: The role of strategic choice. *Sociology.* 6 1–22.
- Coleman, J. 1988. Social capital in the creation of human capital. *Amer. J. Soc.* 94 95-120
- Crosson, M. M., I. Bedrow. 2003. Organizational learning and strategic renewal. *Strat. Mgmt. J.* **24** 1087-1105.

- Crossan, M., H. Lane, R. White. 1999. An organizational learning framework: From intuition to institution. *Acad. Management Rev.* **24** 522-537.
- Cyert, R. M., J. G. March. 1963. *A Behavioral Theory of the Firm*. Prentice Hall, Englewood Cliffs, NJ.
- Daft, R., K. Weick. 1984. Toward a model of organizations as interpretation systems. *Acad. Management Rev.* **66** 70-74.
- Davis, J., K. Eisenhardt., C. Bingham. 2008. Complexity theory, market dynamism and the strategy of simple rules. Working Paper. Stanford University.
- Davis, J. H., N. L. Kerr, R. S. Atkin, R. Holt, D. Meek. 1975. The decision processes of 6- and 12-person mock juries assigned unanimous and two thirds majority rules. *J. Personality Soc. Psych.* **32** 1-14.
- Denrell, J., J. G. March. 2001. Adaptation as information restriction: the hot stove effect. *Org. Sci.* **12** 523-538.
- Denrell, J., C. Fang., D. Levinthal. 2004. From T-mazes to labyrinths: Learning from model based feedback. *Management Sci.* **50**(10) 1366–1378.
- Devine, D.J., L.D. Clayton, J.L. Philips, B.B. Dunford, S.B. Melner.1999. Teams in organizations: Prevalence, characteristics, and effectiveness. *Small Group Research.* 30:678-711.
- Dosi, G. 1988. Sources, procedures, and microeconomic effects of innovation. *J. Econ Lit.* **26** 1120-71.
- Dyer, J., H. Singh. 1998. The relational view: cooperative strategy and sources of interorganizational competitive advantage. *Acad. Management J.* 23 660-679.
- Eisenhardt, K. M., S. L. Brown. 1999. Patching: Restitching business portfolios in dynamic markets. *Harvard Bus. Rev.* 77 72–82.
- Ethiraj, S., D. Levinthal. 2004. Bounded rationality and the search for organizational architecture: an evolutionary perspective on the design of organizations and their evolvability. *Admin. Sci. Quart.* **49** 404-437.
- Fang, C., D. Levinthal. 2008. The near term liability of exploitation: exploration and exploitation in multi-stage problems. *Org. Sci.*. Forthcoming.
- Fleming, L., C. King., A. Juda. 2007. Small worlds and regional innovation. *Org. Sci.* **18** 938-954.
- Fleming, L. 2001. Recombinant uncertainty in technological search. *Management Sci.* **47** 117-132.
- Fleming, L., O. Sorenson. 2001. Technology as a complex adaptive system: Evidence from patent data. *Res. Pol.* **30** 1019-1039.
- Galbraith, J. R. 1973. *Designing Complex Organizations*. Addison-Wesley, Reading, MA.
- Gibson, C., F. Vermeulen. 2003. A healthy divide: Subgroups as a stimulus for team learning behavior. *Admin. Sci. Quart.* **48** 202-239.

- Goldberg, D. E. 1989. *Genetic algorithms in search, optimization and machine learning*. Addison-Wesley.
- Granovetter, M. 1973. The strength of weak ties. *Amer. J. Soc.* **78** 1360-80.
- Granovetter, M. 1992. Problems of explanation in economic sociology, in Nohria, N. and Eccles, R.(eds), *Networks and Organizations: Structure, Form and Action* (Boson, MA: Harvard Business School Press)
- Gupta, A. K., K. G. Smith, C. E. Shalley. 2006. The interplay between exploration and exploitation. *Acad. Management J.* **49** 693-706.
- Hansen, M. T. 1999. The search-transfer problem: The role of weak ties in sharing knowledge across organization subunits. *Admin. Sci. Quart.* **44** 82-111.
- Hansen, M. T. 2002. Knowledge networks: Explaining effective knowledge sharing in multiunit companies. *Org. Sci.* **13** 232-250.
- Helfat C. 1994. Evolutionary trajectories in petroleum firm R&D. *Management Sci.* **40** 1720-1747.
- Hickson, D., D. S. Pugh, D. Pheysey. 1969. Operations technology and organization structure: an empirical reappraisal. *Admin. Sci. Quart.* **14** 378-98.
- Holland, J. 1975. *Adaptation in natural and artificial systems*. U. of Michigan Press, Ann Arbor, MI.
- Holland, J. 1992. Genetic algorithms. *Scientific Amer.* **267**(1) 66-72.
- Hinton, G., S. Nowlan. 1987. How learning can guide evolution. *Complex Systems*, **1** 495--502.
- Huber, G. 1991. Organizational learning: The contributing processes and the literatures. *Org. Sci.* **2** 88-115.
- Jacobides, M. 2007. The inherent limits of organizational structure and the unfulfilled role of hierarchy: Lessons from a near-war. *Org. Sci.* **18** 455-477.
- Jansen, J. J. P., F. A. J. Van Den Bosch, H. W. Volberda. 2006. Exploratory innovation, exploitative innovation, and Performance: Effects of organizational antecedents and environmental moderators. *Management Sci.* **52** 1661-1674.
- Kameda, T., J. Davis. 1990. The function of the reference point in individual and group risk decision making. *Organizational Behavior and Human Decision Processes.* **46**: 55-76.
- Kane, G. C., M. Alavi. 2007. Information technology and organizational learning: An investigation of exploration and exploitation processes. *Org. Sci.* **18** 796-812.
- Kauffman, S. 1993. *The origins of order*. Oxford University Press.
- Kauffman, S. 1995. *At Home in the Universe*. Oxford University Press.
- Klimoski, R., S. Mohammed. 1994. Team mental model: construct or metaphor? *J. Mgmt* **20** 403-437.

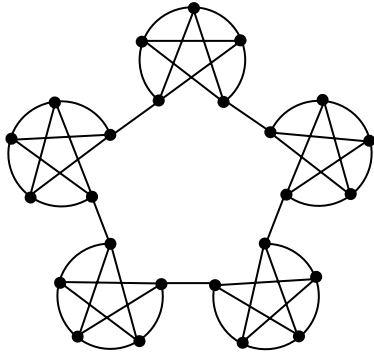
- Larson, J. R., C. Christensen. 1993. Groups as problem-solving units: Toward a new meaning of social cognition. *Brit. Psych. Society* **32** 5-30.
- Lavie, D., L. Rosenkopf. 2006. Balancing exploration and exploitation in alliance formation. *Acad. Management J.* **49** 797-818
- Lawrence, P. R., J. W. Lorsch. 1967. *Organization and Environment*. Harvard Business School Press, Boston, MA.
- Lee, E., J. Lee, J. Lee. 2006. Reconsideration of the winner-take-all hypothesis: Complex networks and local bias. *Management Sci.* **52** 1838-1848.
- Levinthal, D. A. 1997. Adaptation on rugged landscapes. *Management Sci.* **43** 934–950.
- Levinthal, D., J. G. March. 1981. A model of adaptive organizational search. *J. Econ. Beh. Org.* **2** 307-333.
- Levinthal, D., J. G. March. 1993. The myopia of learning. *Strat. Management J.* **14** 95–112.
- Levitt, B., J. G. March. 1988. Organizational learning. *Ann. Rev. Soc.* **14** 319–340.
- Lenox, M. 2002. Organizational design, information transfer, and the acquisition of rent-producing resources. *Computation and Mathematical Organization Theory.* **8**(2) 113-131.
- Lewin, A.Y., C. P. Long, T. N. Carroll. 1999. The coevolution of new organizational forms. *Org. Sci.* **10** 535-550.
- March, J. G. 1991. Exploration and exploitation in organizational learning. *Org. Sci.* **2** 71-87.
- March, J. G. 2004. Parochialism in the evolution of a research community: The case of organization studies. *Management Org. Rev.* **1** 15-32.
- Martin, X., W. Mitchell. 1998. The influence of local search and performance heuristics. *Res. Pol.* 26(7/8) 753-772.
- Miller, K. D., M. Zhao, R. J. Calantone. 2006. Adding interpersonal learning and tacit knowledge to March's exploration-exploitation model. *Acad. Management. J.* **49** 709-722.
- Mitchell, M. 1996. *An Introduction to Genetic Algorithms*. MIT Press, Cambridge, MA.
- Nelson, R. R., S. G. Winter. 1982. *An Evolutionary Theory of Economic Change*. Harvard University Press, Cambridge, MA.
- Nickerson, J. A., T. R. Zenger. 2002. Being efficiently fickle: A dynamic theory of organizational choice. *Org. Sci.* **13** 547–566.
- Onnela, J. P., J. Saramäki, J. Hyvönen, G. Szabó, D. Lazer, K. Kaski, J. Kertész, and A.-L. Barabási. 2007. Structure and tie strengths in mobile communication networks. *Proc. National Acad. Sciences of the USA* **104** 7332-7336.
- O'Reilly, C.A., M. L. Tushman. 2004. The ambidextrous organization. *Harvard Bus. Rev.* **82**(4) 74-81.

- Patel, H. K. Pavitt. 1997. The technological competency of the world's largest firms: complex path-dependent, but not much variety. *Res. Pol.* **26** 141-156.
- Perrow, C. 1967. A framework for the comparative analysis of organizations. *Amer. Soc. Rev.* **32** 194-208.
- Reagans, R., B. McEvily. 2003. Network structure and knowledge transfer: The effects of cohesion and range. *Admin. Sci. Quart.* **48** 240-267.
- Rivkin, J. W. 2000. Imitation of complex strategies. *Management Sci.* **46** 824-844.
- Rivkin, J. 2001. Reproducing knowledge: replication without imitation at moderate complexity. *Org. Sci.* **12** 274-293.
- Rogers, M. 2003. It's the Apple of his eye. *Newsweek*, March 1.
- Sandelands, L., R. Stablein. 1987. The concept of organization mind. *Res. in the Soc. of Org.* **5** 135-161.
- Schilling, M. A., C. Phelps. 2007. Interfirm collaboration networks: The impact of large scale network structure on firm innovation. *Management Sci.* **53** 1113-1126.
- Schilling, M., K. Steensma, K. 2001. The use of modular organizational forms: An industry level analysis. *Acad. Management J.* **44** 1149-1169.
- Schoemaker, P. 1990. Strategy, complexity, and economic rent. *Management Sci.* **36**(10) 1178-1192.
- Simon, H. 1978. Rational decision-making in business organizations. In Nobel Lectures, Economics 1969-1980, Assar Lindbeck (Ed), World Scientific Publishing, 1992
- Siggelkow, N., D. Levinthal. 2003. Temporarily divide to conquer: centralized, decentralized, and reintegrated organizational approaches to exploration and adaptation. *Org. Sci.* **14** 650-669.
- Siggelkow, N. D. Levinthal. 2005. Escaping real (non-benign) competency traps: linking the dynamics of organizational structure to the dynamics of search. *Strategic Org.* **3** 85-115.
- Siggelkow, N., J. Rivkin. 2005. Speed and search: designing organizations for turbulence and complexity. *Org. Sci.* **16** 101-122.
- Siggelkow, N., J. Rivkin. 2006. When exploration backfires: unintended consequences of multi-level organizational search. *Acad. Management J.* **49** 779-795.
- Stuart T, Podolny J. 1996. Local search and the evolution of technological capabilities. *Strategic Management J.* **17** 21-38.
- Strogatz, S. 2001. Exploring complex networks. *Nature* **410** 268-276.
- Taylor, A., H. Greve. 2006. Superman or the fantastic four? Knowledge combination and experience in innovative teams. *Acad. Management J.* **49** 723-740.

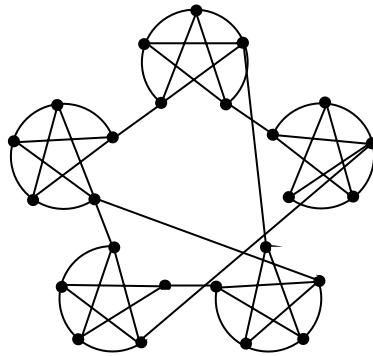
- Thompson, J. D. 1967. *Organizations in Action*. McGraw-Hill, New York.
- Ton, Z., R. Huckman. 2008. Managing the impact of employee turnover on performance: the role of process conformance. *Org. Sci.* 19(1) 56–68.
- Uzzi, B. 1997. Social structure and competition in interfirm networks: the paradox of embeddedness. *Admin. Sci. Quart.* 42 35-67.
- Uzzi, B., J. Spiro. 2005. Collaboration and creativity: the small world problem. *Amer. J. Soc.* 111 447–504.
- Waldrop, M. 1992. *Complexity: The emerging science at the edge of order and chaos*. New York: Simon & Schuster.
- Watts, D. J. 1999. *Small Worlds: The Dynamics of Networks Between Order and Randomness*. Princeton University Press, Princeton, NJ.
- Williamson, O. 2003. Examining economic organization through the lens of contract. *Indus. Corp. Change* 12 917-42.
- Winter, S., G. Cattani, A. Dorsch. 2007. The value of moderate obsession: Insights from a new model of organizational search. *Org. Sci.* 18 403-419.
- Woodward, J. 1965. *Industrial Organization: Theory and Practice*. Oxford University Press, London.
- Wright, S. 1932. The roles of mutation, inbreeding, crossbreeding and selection in evolution. *Proceedings of the VI International Congress of Genetics* 356-366.
- Wright, S. 1964. Stochastic processes in evolution. J. Gurland, ed. *Stochastic Models in Medicine and Biology*, University of Wisconsin Press, Madison.
- Zander, U., B. Kogut. 1995 Knowledge and the speed of the transfer and imitation of organizational capabilities: an empirical test. *Org. Sci.* 6(1) 76-92

FIGURE 1 Variations on Watts' (1999) Connected Caveman Model^a

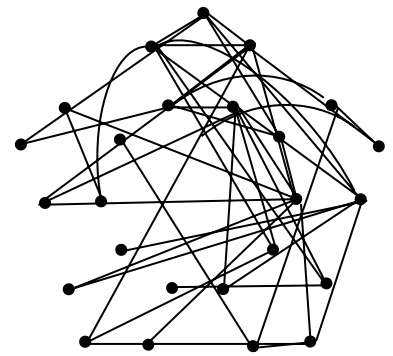
a) Nearly-isolated subgroup structure



b) Semi-isolated subgroup structure with randomly rewired links

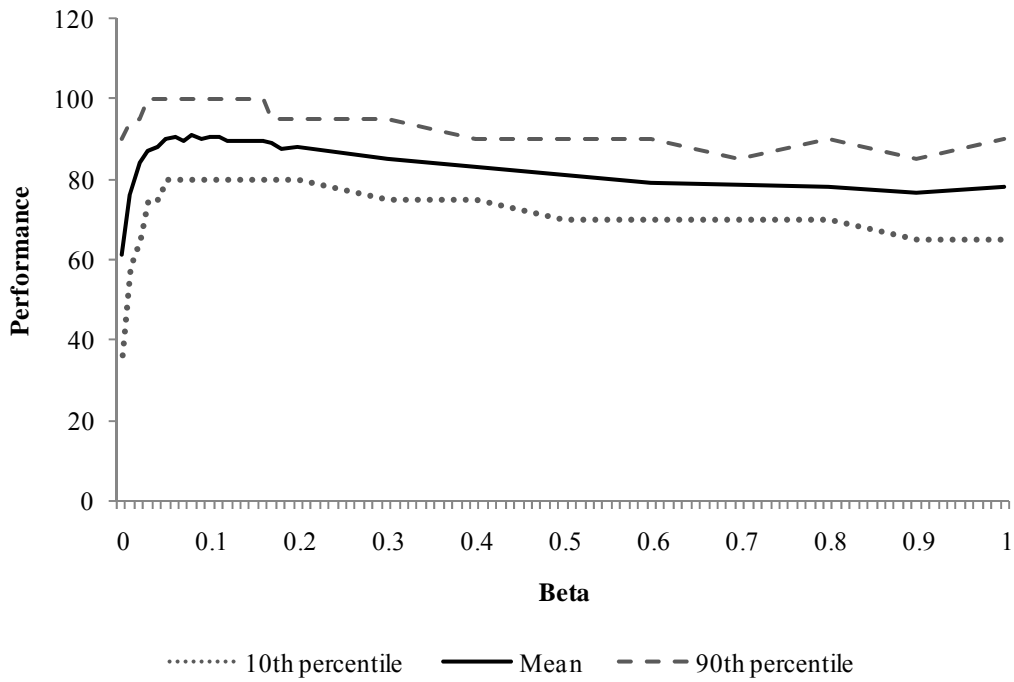


c) Random network: Network structure without subgroup identity



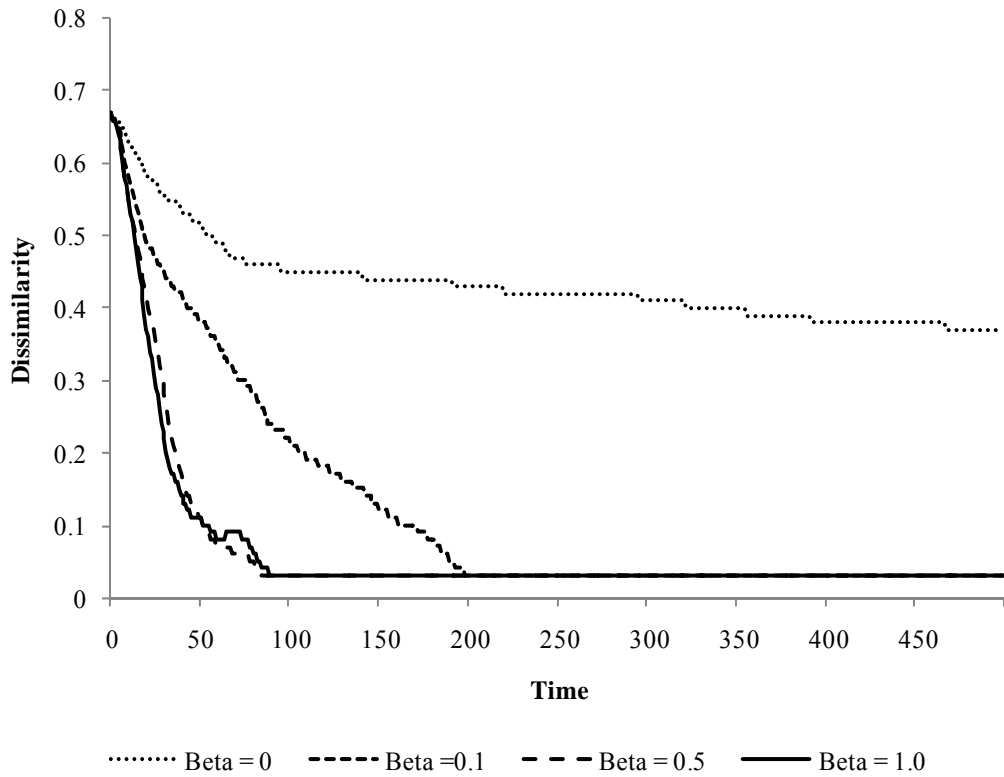
^a The actual models we used had 280 individuals and subgroups of seven individuals each, but for visual clarity, simplified models are shown here.

FIGURE 2 Organizational Structure and Equilibrium Performance Outcome



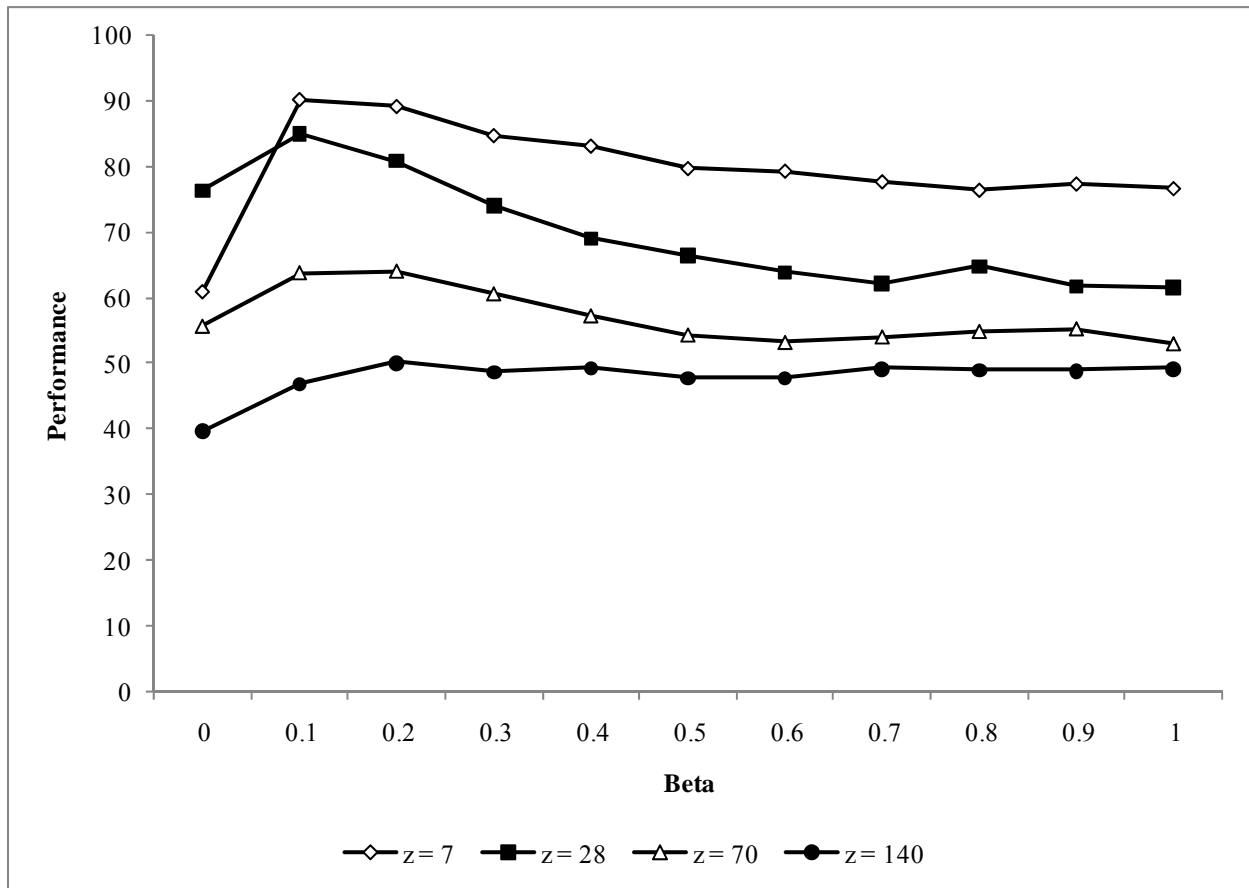
Note: Each data point here is the average over 200 simulation runs observed at equilibrium points. The organization reaches a long run equilibrium performance when its performance does not improve anymore. As superior ideas diffuse, there comes a point where the entire organization converges upon a set of beliefs and there is no more superior performing individuals to learn from.

FIGURE 3 Diversity of Beliefs over Time



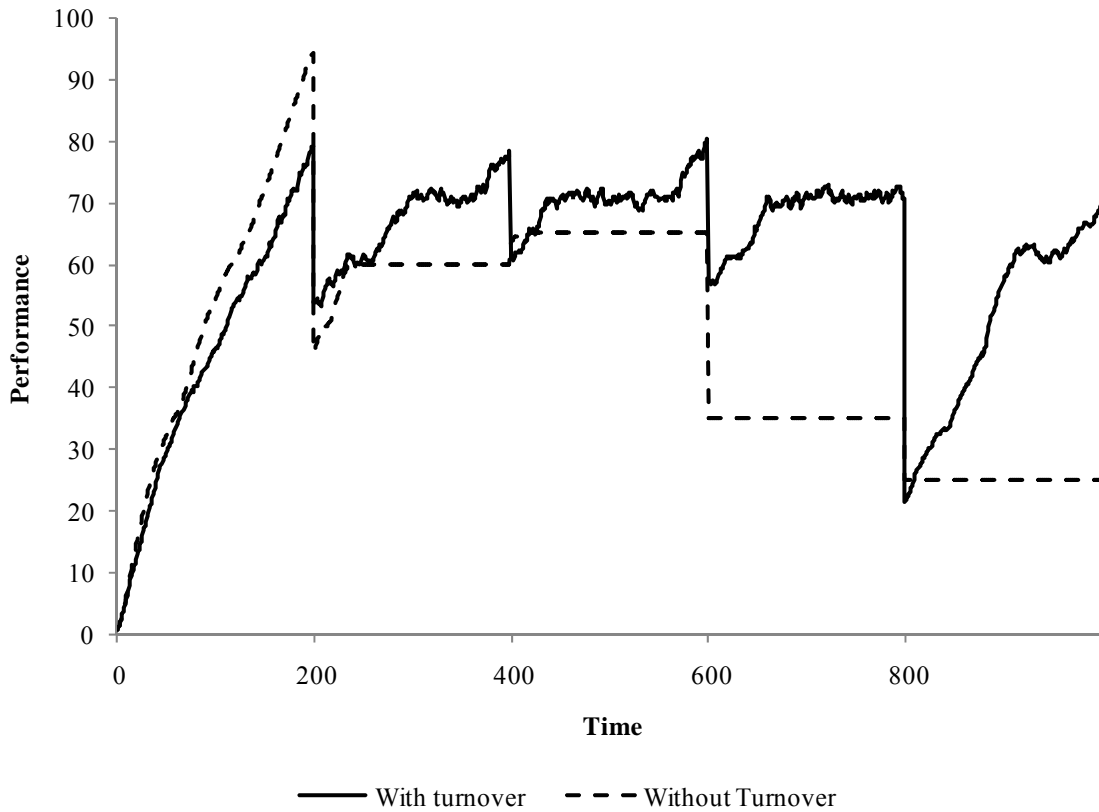
Note: The figure shows a typical simulation run. The dissimilarity index indicates how individuals in an organization are different from each other in viewing reality with m dimensions. This index was constructed by making pair-wise comparisons of all n individuals across m beliefs.

FIGURE 4 Effect of Subgroup Size on Organizational Performance



Note: The figure shows the effects of different subgroup sizes on the long-term learning performance, and z represents the size of a subgroup. Each data point here is the average over 200 repeated simulations.

FIGURE 5 Effect of Turnover in Changing Environment



Note: The graph with turnover shows a typical simulation run when 10% of the elements of reality change their values every 200 periods and when 1% of the organizational members are replaced by new members in each time period. The graph without turnover shows a typical simulation run with the same periodic change of reality but with no membership change.

TABLE 1 Mean Performance of Organizations under Different Rates of Turnover and Environmental Change

	Rates of Turnover			Rates of Environmental Change		
	No ($p_{turnover} = 0$)	Low ($p_{turnover} = 0.01$)	High ($p_{turnover} = 0.1$)	Low ($p_{envirr} = 0.1$)	Medium ($p_{envirr} = 0.4$)	High ($p_{envirr} = 0.7$)
Nearly-isolated	31.17	44.45	23.32	44.45	26.05	23.41
Semi-isolated	36.42	61.49	41.71	61.49	38.13	34.22
Random	32.39	54.47	38.63	54.47	32.11	28.64
F-value	21.18	1127.23	4187.54	1129.23	361.31	257.55
P-value	0.001	0.001	0.001	0.001	0.001	0.001

Note: The table reports the levels of organizational performance across three levels of turnover and three different levels of environmental change. For the results where turnover is varied, the value for p_{envirr} is fixed at 0.1. For the results where environmental change is varied, $p_{turnover}$ is fixed at 0.01. The results show that the semi-isolated subgroup structure is always the highest performer across different turnover and environmental change conditions, and the performance differences among different structures are statistically significant at $p < 0.01$.

Appendix A: Formal Description of the m/s Generalized Payoff Function

Let $\Phi(x)$ denote the payoff function for a bit string x with its dimension m . We assume that $\Phi(x)$ is characterized by a continuum between two polar ends. On one end of the continuum is a linear payoff function $L(x)$ —this is equivalent to the one March (1991) used in his learning model. Let x_j denote j th element of the bit string x . Then, the linear payoff function is

$$L(x) = \sum_{j=1}^m \delta_j \quad (1)$$

where $\delta_j = 1$ if x_j corresponds with reality on that dimension; $\delta_j = 0$ otherwise. In our formulation, $L(x)$ is a special case of $\Phi(x)$. When the payoff function is characterized by (1), it is rather easy for an organization to search for higher payoff points because a payoff of each problem is independent from others.

Consider the other extreme on the continuum: Hinton and Nowlan's (1987) payoff function, which is known as a "needle-in-a-haystack" search problem. Let $H(x)$ denote this payoff function. There is only one peak in a space of 2^m possibilities. Unlike March's linear search problem, this problem is very hard because the search landscape provides no cue for guiding evolution to the peak. For instance, suppose $m = 10$. For simplicity, suppose that a bit string representing the highest payoff is 1111111111. There are zero payoffs to even minor deviations from this configuration, for example, 1011111111 and 1101111111. That is, even when the organization is one step away from the peak, there is no clue for the organization to infer where the peak is. Only random trial and error can lead it to the peak. At this level of interdependency, it is extremely hard for organizations to improve their performance—a manifestation of the "non-evolvability condition" discussed in Kauffman (1993; 1995).

Between the two extreme cases, there is a middle ground, where an m -bit string is partitioned into l independent subsets. Within each subset, there are s bits, whose performance is coupled. Note that $l = m/s$. Formally, we represent our generalized payoff function as

$$\Phi(x) = s \left(\prod_{j=1}^s \delta_j + \prod_{j=s+1}^{2s} \delta_j + \dots + \prod_{j=m-s+1}^m \delta_j \right) \quad (2)$$

Here, s serves as a tunable parameter that can control the difficulty of the search problems, and $1 \leq s \leq m$. When $s = 1$, $\Phi(x) = L(x)$. By increasing the value of s , the search problem becomes more interdependent. The bigger the value of s , the more interdependent the search problem. When $s = m$, the search problem becomes Hinton and Nowlan's (1987) payoff function. That is,

$$\Phi(x) = \prod_{j=1}^m m \delta_j = H(x) \quad (3)$$

For instance, suppose reality is represented by a string of 11111 11111 and we set our parameter s to 5. The payoff to one individual A whose beliefs are 11100 11111 is: $5 \times 1 \times 1 \times 1 \times 0 \times 0 + 5 \times 1 \times 1 \times 1 \times 1 \times 1 = 5$. The payoff to another individual B whose beliefs are 11111 11111 is: $5 \times 1 \times 1 \times 1 \times 1 \times 1 + 5 \times 1 \times 1 \times 1 \times 1 \times 1 = 10$. If we set our parameter s to 10, A's payoff will be 0 while B's will remain at 10. Thus the higher the s parameter, the more complex is the problem.

The main benefit of this characterization of the m/s payoff function is that we can control the difficulty of a search problem with only a single parameter s . An organization's performance is measured as the average performance across all individuals in the organization.

Appendix B: Model Parameters and Steps

Parameters

Parameters	Remarks	Range of parameter values analyzed
n	number of individuals in the organization	140, 210, 280, 350, 420
m	dimensions of beliefs	40, 70, 100, 130, 160
z	size of a subgroup	7, 14, 28, 70, and 140*
β	rewiring probability	Varies from 0 to 1
$p_{learning}$	probability of individual learning from the majority view	0.1, 0.2, 0.3, ..., 0.8, 0.9
s	degree of complexity	1, 3, 5, 7, 10
$p_{turnover}$	probability of turnover	0, 0001, 0.005, 0.01, 0.05, 0.1, 0.5, 1.0
p_{envir}	probability of change in each element of reality	0.1, 0.4, 0.7
T	interval for environment change	200

* Changing the size of a subgroup entails corresponding changes to values of n because the number of individuals in the organization (i.e., n) has to be divisible by the size of a subgroup.

Simulation Steps

Preparation Step. All parameter values for the simulation are set up at this stage. There are 280 individuals in each organization (i.e., $n = 280$). The number of dimensions in the beliefs is set to 100 (i.e., $m = 100$). Reality is determined by randomly assigning a value of 1 or -1 for each of 100 dimensions, while each dimension of an individual's belief set is determined by assigning a value randomly drawn from 1, 0, or -1. Each organization consists of 40 clusters of individuals (i.e., each cluster has seven individuals, though we will later explore the effect of varying the size of the cluster). This organizational network is implemented by first connecting each individual to six other individuals in a pattern analogous to the connected caveman pattern in Figure 1a. We then rewire these connections according to β . For each focal individual, we draw a random number from a uniform distribution [0, 1]. If the random number drawn is smaller than β , we eliminate one of the individual's links and randomly choose another individual outside of the cluster to whom the focal individual will have a new connection. We repeat this for all the existing connections of the focal individual and then repeat for all individuals in the organization.

Period 1. The simulation procedure evaluates each individual's performance by comparing her beliefs with the m dimensions of reality. This evaluation is based on the payoff function described previously. Then, superior performers to which the focal individual is directly connected are identified. When there is no superior performer, the focal individual keeps all of her previous beliefs. When there are two or more superior performers, a majority belief among them will be determined for each of m dimensions. The focal individual adapts to each majority belief with probability $p_{learning}$. When there are ties (i.e., the number of superior performers who believe in -1 is equal to that of those who believe in 1), the focal individual keeps her existing belief. All the individuals learn simultaneously from one another based on beliefs at the beginning of the period.

Period t . Individuals update their beliefs at period t . The learning process is repeated until no further change in any individual's beliefs occurs (i.e. when equilibrium is obtained).

Appendix C: Comparison of Majority Decision Rule versus Tournament Selection

Note: Because we are interested in modeling organizational learning, we utilized a majority decision rule that simulates the socialization process by which an individual is influenced by their peers. This is closely analogous to the setup in the original March (1991) model. Not surprisingly, subgroup structure matters much less when the model does not incorporate socialization. If, for example, one employs a “tournament selection” rule (an individual identifies the best performing individual to which they are connected and emulates each of their beliefs with probability $p_{learning}$), subgroup structure has little effect on performance. Such a learning rule outperforms the majority decision rule when β is extremely low (nearly isolated subgroups) or very high (when subgroup structure is lost), but underperforms the majority decision rule when subgroups are semi-isolated, as shown below.

