Exploration and Exploitation in Complex Networks:

The Influence of Interpersonal Network Structure[1]

Christina Fang
New York University
44 West Fourth Street
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
New York University
40 West Fourth Street
New York, NY 10012
212-477-4213
mschilli@stern.nyu.edu

March, 2006

[1] Authorship is alphabetical; all authors contributed equally.

A previous version of this paper was presented at the Organization Science Winter Conference in February 2006.

Acknowledgements: We are grateful for comments received from Paul Adler, Linda Argote, Jerker Denrell, James G. March, and Nicolaj Seggelkow.

Exploration and Exploitation in Complex Networks:

The Influence of Interpersonal Network Structure

Abstract

We use a simulation study to extend James March's (1991) classic model of exploration and exploitation by allowing for direct interpersonal learning. As individuals interact with one another in an organization, they evaluate each other's performance and alter their own solution routines by imitating other superior performers. In this setup, we are interested in whether a subgroup structure, which is popular in many organizations in the form of team or department structures, fosters the diversity of ideas and solution routines in an organization, thereby improving learning outcomes. We find that this structure is, indeed, conducive to learning, as long as there is a small fraction of random, cross-group links. Such semi-isolation allows the organization to learn moderately fast with the highest learning outcome. This numerical result is consistent with empirical research, which shows that small-world network properties enhance the performance of a system.

KEY WORDS: Learning; Exploration and Exploitation; Networks; Small World; Evolution

In 1991, James March published a seminal work on organizational learning that demonstrated a trade-off between short-run and long-run concerns in an organization. March developed a model in which an organizational code interacts with individuals who are initially endowed with diverse sets of beliefs. The organizational code learns from individuals with superior knowledge, and individuals, in turn, learn from the code. As a result, the organization's knowledge improves over time. His model illustrated that "fast learners" (those that quickly adopt beliefs from the organizational code) rapidly exploit the firm's existing knowledge, increasing the efficiency of organizational learning. Fast learners, however, are likely to converge prematurely on a homogeneous set of beliefs, thwarting long-run learning and leading the organization to a suboptimal equilibrium. In March's simulations, a slower learning rate, though less efficient, allowed the organization to preserve more diversity of individual beliefs, enabling the firm to explore a wider range of possible combinations of beliefs and increasing the chance of improving the quality of organizational knowledge in the long run. This model demonstrated 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. The trade-off between exploration and exploitation represents a fundamental conflict between short-run and long-run concerns in many adaptive systems (Holland, 1975, Goldberg, 1989).

Since adaptive processes often put organizations at risk of focusing primarily on exploitation to the expense of exploration (March 1991), firms must conscientiously attend to and manage the mechanisms that influence their use of exploitation versus exploration. One such mechanism is the interpersonal network structure that shapes how individuals exchange and create knowledge within the firm.

When individuals in an organization interact, they pool, exchange, and recombine information, resulting in the creation or refinement of the knowledge possessed by each individual (Argote, 1999; Brown & Duguid, 1991; Larson & Christensen, 1993; Schilling & Phelps, 2006b; Wegner, 1987). Furthermore, numerous studies suggest that with whom an individual interacts can have influence well beyond any specific dyadic relationship; through the process of interacting with one another, individuals weave a network of direct and indirect relationships that serve as conduits for information dissemination throughout the firm. The overall pattern of relationships created through interpersonal interaction represents the structure of the interpersonal network. As demonstrated in recent studies, the structure of the interpersonal network within which individuals are embedded influences the likelihood and degree to which information is exchanged among a group of individuals, and the ease at which it may be assimilated and utilized productively (Borgatti & Cross, 2003; Hansen, 1999; Hargadon, 2003; Reagans & McEvily, 2003; Schilling & Phelps, 2006b; Uzzi, 1997; Uzzi & Spiro, 2005).

In this paper, we extend March's 1991 model by allowing for direct interpersonal learning. In particular, we examine learning dynamics in an organization that consists of many, semi-isolated subgroups. This structure, which is commonly observed in the form of team or department structures, allows parallel, isolated learning among subgroups, since individuals are more likely to interact with other individuals within a subgroup than outside it. The key questions we raise are: Does this subgroup structure enhance organizational learning? Could an organization improve the balance of exploration and exploitation in organizational learning by designing its structure properly?

Our numerical analysis demonstrates that the subgroup structure fosters the diversity of ideas and solution routines, thereby enhancing the long-run outcome of learning, as long as there is a small fraction of random, cross-group links—i.e., subgroups are semi-isolated. These links play a role of bridges between different subgroups, dramatically facilitating exchanges of superior knowledge across groups. But, either when there are too many such cross-group links or when they are almost (or completely) absent, the long-run outcomes of learning tend to be suboptimal. All these findings together suggest that the semi-isolated subgroup structure improves the balance between exploration and exploitation. In this context, exploitation arises from learning across subgroups, which facilitates rapid diffusion and assimilation of currently superior knowledge, while reducing heterogeneity of knowledge within the firm. Exploration arises from the parallel, isolated learning among subgroups, which preserves the requisite variety of knowledge necessary for the organization to explore different areas of its fitness landscape.[2]

Our paper is organized as follows. First, we introduce a model of organizational learning with subgroup structures. Second, we present the numerical results. Finally, we discuss the implications of our findings in light of the extant literature.

MODEL

We model organizational learning as a process by which individuals within an organization interact to exchange and jointly create knowledge. When individuals first join an organization, they have idiosyncratic sets of beliefs about reality 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. If individuals are able to ascertain that others have belief sets that are better performing in some way, they may update

[2] The term "requisite variety" generally refers to Ashby's (1956) Law of Requisite Variety that states, " the available control variety must be equal to or greater than the disturbance variety for control to be possible." This axiom is widely used to explain the need for variety to ensure evolvability in any entity subject to selection pressure (e.g., species diversity).

their own beliefs to incorporate aspects of the higher performing belief sets. The process of sharing, comparing, interpreting, and updating beliefs helps groups of individuals develop a shared understanding (Argyris & Schon,1996; Crossan, Lane & White, 1999; Daft & Weick, 1984; Huber, 1991, Klimoski & Mohammed, 1994; Sandelands & 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 lead 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 which 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. The focal individual may then decide to update each dimension of his or her own belief set to this majority view with some probability p that reflects the proclivity or ability of individuals to learn from one another. This majority decision rule (and equiprobability otherwise) is 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 & Goodrich, 1971; Davis et al, 1975; Kerr et al. 1976).

It is particularly likely to be used by groups when it is difficult for individuals to assess which alternative is correct.

Main Entities

In our model, there are three main entities, *reality*, *individual*, and *organization*.

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.

Individual. There are *n* individuals in an organization. Each of them holds *m* beliefs about reality at each time step. Each belief for an individual has a value of 1, 0, or -1.

Organization. As mentioned earlier, our model is different from the March model (1991) in that an organization is seen as a complex system where its individual members directly interact with one another. A key to the understanding of the dynamics of such a system lies in its topology of interaction patterns (Strogatz 2001). To represent interpersonal interactions, we modified Watts' (1999) "connected-caveman" network, which is irregular—some nodes have more links than others. As shown on Figure 1a, our modified connected-caveman network (or "nearly-isolated" network) is a regular network, where every node has the same number of links. To construct different interaction patterns, we rewire each link of this subgroup structure by making random connections between subgroups with probability β. When $\beta = 0$, the nearly-isolated network structure is preserved. When β is non-zero but sufficiently small, the subgroup structure is preserved, but there is a few random connections among subgroups. On the other hand, when β grows large and approaches unity, a network becomes random, as shown in Figure 1c.

---------------------Insert Figure 1 About Here --------------------------

Learning

Individual learning occurs when each individual modifies her beliefs through interactions with other individuals. Organizational learning is viewed as a collective outcome of interactions of individual members. An individual is likely to change her beliefs when she is influenced by "superior performers." They are defined as those individuals who are connected to a focal individual and whose performance is superior to that of the focal individual. These superior performers may not have the same beliefs on each of m dimensions. Then, the focal individual adapts to a dominant belief (or a majority view) within the superior performers on any particular dimension. Like March (1991), we also assume that this learning process is probabilistic. That is, each individual adapts to each dominant belief on m dimensions with probability p.

Payoff

The performance of an individual is evaluated at every time step by a given payoff function. Let $\Phi(x)$ denote a generalized 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. In the 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 δ_j denote jth element of the bit string x. Then, the linear payoff function is

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

where $\delta_j = 1$ if jth belief for an individual corresponds with reality on that dimension, and $\delta_j = 0$ otherwise. In this framework, $L(x)$ 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.

Let us consider another extreme case. We consider 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 3^m possibilities, and all the others are characterized by flat surface. 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. The essence of the problem is that its neighboring points, say, 1011111111 or 1101111111, show a zero payoff. 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 every position, there are 3^m -1 ways to move.

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 can represent our generalized payoff function as

$$\Phi(x) = \prod_{j=1}^{s} s\delta_j + \prod_{j=s+1}^{2s} s\delta_j + \ldots + \prod_{j=m-s+1}^{m} s\delta_j \qquad (2)$$

Here, s serves as a tunable parameter that can control the difficulty of the search problems. Note that $1 \leq s \leq m$. When $s = 1$, $\Phi(x) = L(x)$. 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 for $s = 5$. When an individual obtains all of correct beliefs, her payoff is 5. But, when an individual has four correct beliefs and one incorrect one, her payoff is zero. In general, if any single element among s beliefs is wrong, the payoff for the whole subset becomes zero. The bigger the value of s is, the more interdependent the search problem is. In other words, s represents the degree of coupling. When s is small, the search problem is loosely coupled. When s is large, the search problem becomes tightly coupled. When $s = m$, the search problem becomes Hinton and Nowlan's (1987) payoff function. That is,

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

The main benefit of this characterization of the 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.[3]

Simulation Procedure

Preparation Step. All parameter values for the simulation (see Appendix A) are set up at this stage. Reality is determined by randomly assigning a value of 1 or -1 for each of m dimensions. In the runs of the simulation for the current study, m is set to 100. There are 280 individuals in each organization. Each individual is assigned a value for each belief on the m dimensions, randomly drawn from 1, 0, or -1. Also, an organizational network is constructed by connecting each individual to six other individuals within a subgroup, and then rewiring these connections according to the algorithm specified before.

Period 1. The simulation procedure will evaluate each individual's performance by comparing her beliefs with values on m dimensions of reality. This evaluation is based on the payoff function described before. Then, superior performers for each individual will be determined. 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

[3] Note that by the time equilibrium is obtained, almost all individuals will possess identical beliefs so the difference between using the average performance or the maximum performance of any single individual or group of individuals is trivial.

probability p.

Period t. All of individuals at time $t-1$ who have superior performers will modify their beliefs at period t. The learning process will be repeated until no further change in any individual's beliefs arises (equilibrium is obtained).

SIMULATION RESULTS

To investigate whether organizational structure affects learning outcomes, we developed a computational model, which would enable a learning process that is almost identical to the one used in March (1991), except that learning occurs through interpersonal interaction rather through interaction between employees and an organization code. In this setting, we are interested in whether a subgroup structure, which is popular in many organizations in the form of team or department structures, fosters the diversity of ideas and beliefs in an organization, thereby improving learning outcomes. First, consider an extreme type of organizational structure, "nearly-isolated subgroup structure," which is illustrated on the left in Figure 1. We can construct diverse kinds 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 number 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. When $\beta = 1$, subgroup identity disappears due to too many cross-group links. The questions are: which type is most conducive to improving learning outcome in the long run?; can we improve the balance between the long-run outcome and learning speed by designing an organization properly?

Typical Simulation Runs

As a starting point, we show the typical behavior of our learning model. In March's model, the speed of organizational learning is controlled by learning rates for organizational members. Since the key variable of our interest is organizational structure, we keep learning rate constant at $p = 0.3$ for the moment—we relax this assumption later in our sensitivity analysis. The other parameters used in the simulations are specified in Appendix. Figure 2 shows that all of the dynamics reach some steady states in the long run. The results indicate that the speed of organizational learning is affected by structural parameter β, which can be roughly interpreted as a fraction of cross-group links. The bigger the fraction of cross-group links, the faster the speed of organizational learning. By more facilitating the exchange of diverse ideas and beliefs across subgroups, organizational members can learn faster. The learning speed tend to be faster when β approaches 1 or when subgroups tend to lose their identity with too many cross-group links. On the other hand, the learning outcome tends to be lower with any increase in β over an interval of $\beta \in [0.1, 1]$. When subgroups are "semi-isolated" with a small fraction of random cross-group links (e.g., $\beta = 0.1$), the long-run outcome of learning seems to be highest. On the other hand, when $\beta = 0$, or when subgroups are nearly isolated, both the learning speed and outcome are lowest. Due to the randomness in our system, however, we cannot make a conclusion only with typical simulations.

Simulation Experiments

We conducted simulation experiments to systematically investigate the effects of network structure on the long-run outcome of learning. To reduce statistical errors, we repeated each simulation 100 times. All the data in this figure are averaged over 100 simulation runs at the steady states.

As shown in Figure 3, the results of the simulation experiments confirm the previous typical observation that the long-run average outcome of learning tends to decline over the interval of $\beta \in [0.1, 1]$ as β approaches 1. In this range, the variance of observations is roughly equal across different values of β, as shown in the cumulative frequency distributions in Figure 4.[4] One may wonder why the long-run outcome tends to decrease with an increase in cross-group links. To address this issue, we developed a measure of organizational diversity, or what we call "dissimilarity index." To construct such a measure, we make pairwise 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 as follows:

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

where ω_{ij} takes on a value of 1 if two chosen individuals for ith pair have different beliefs on jth dimension, and 0 otherwise.

Figures 5a and 5b show that an organization tends to lose the diversity of ideas and beliefs very fast when networks are characterized by a large fraction of cross-group links (e.g., $\beta = 1$) or the loss of subgroup identity. Here, dissimilarity among organizational members quickly disappears. In our basic model, once organizational diversity is lost, there is no way for an organization to promote it. So, the management of such diversity is crucial for improving the long-run performance. Diversity among individuals is maintained longer in a semi-isolated network structure (e.g., $\beta = 0.1$), where a subgroup structure exists with a small fraction of cross-group links. Indeed, such semi-isolation improves the long-run outcome of learning. Figure 3

[4] An exception is the variance of observations for $\beta = 0$, which is much larger. The cumulative frequency distribution for this case covers a much broader range than other distributions.

indicates that this finding is not parameter-specific, but that there is a broad interval of β over which the long-run outcome is relatively high.[5] These results suggest that subgroup structure, which fosters organizational diversity, is conducive to learning, as long as a small fraction of random, cross-group links exist.

But our study shows that the long-run outcome of learning is low when $\beta = 0$, or when subgroups are nearly isolated. In this case, diverse ideas and beliefs among different groups are well maintained. Figures 5a and 5b demonstrate that dissimilarity among individuals decays very slowly in this network. Then, why do we systematically observe the lower performance in this network? In our model, learning from other superior individuals is key to improving the outcome. The problem with this 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 well is no good unless it influences organizational members across subgroups to improve their knowledge.

Would the long-run outcome of organizational learning be worse if subgroups are completely isolated from one another? We tested this idea by running another simulation. Our numerical analysis shows that the average learning outcome under this setting is close to zero, indicating that the complete isolation makes it very hard for an organization to improve its knowledge. We also tested another extreme form of organizational structure, a complete network, where every individual is directly connected to every other individual in an organization—in this case, subgroup identity completely vanishes. Our analysis reveals that the average outcome is about 36.7, which is much lower than those in the above sparse networks. This suggests that organizational learning is also very hard in the complete absence of subgroup structure.

[5] We thank James March for bringing this issue to our attention.

DISCUSSION

In this study, we used a series of simulations to examine the influence of the interpersonal network structure on organizational-level learning. We found that a 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. Such semi-isolation allows the organization to learn moderately fast with the highest learning outcome. On the other hand, either the near-isolation or the absence of subgroups does not help organizational members improve learning outcome in the long run.

Our findings speak to the literature on exploration and exploitation. In March's (1991) work, the main source of the tension between exploration and exploitation was learning rate. Exploitation arises as fast learners increase the efficiency of organizational learning, leading to a suboptimal learning outcome in the long run. On the other hand, exploration occurs when an organization becomes inefficient in learning with slow learners, who collectively achieve the better, long-run outcome by preserving the diversity of ideas and beliefs. Our model did show this sort of tension when we varied learning rates. But a more important contribution of our paper is to identify a new source of tension between exploration and exploitation. In our model, exploitation arises from learning across subgroups, which facilitates the faster diffusion of currently superior knowledge. On the other hand, exploration arises from the parallel, isolated learning among subgroups, which better maintains the diversity of solution routines. That is, our study numerically demonstrated that organizational structure is an important source of tension between exploration and exploitation.

Furthermore, our findings address a key question in the literature: How should an organization balance the effort between exploration and exploitation? Our study numerically demonstrated that semi-isolation with a small fraction of random, cross-group links improves the balance between learning speed and long-run outcome. This study thus extends prior work on exploration and exploitation by demonstrating the importance of organizational structure, and offers important implications for organization design.

Implications for Organizational Design and Innovation

Our results have direct implications for managers seeking to design organizations that strike a balance between exploration and exploitation. In our study, the best equilibrium performance levels were obtained when groups of individuals in the organization were semi-isolated. This enabled each group to pursue an isomorphic learning path until progress slowed, and then explore whether the within-group solution could be improved upon through recombination with the solutions created by other groups. These results are consistent with prior work that 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 & Christensen, 1995). Isolating NPD teams and giving them considerable autonomy can help them to pursue new technological possibilities, unfettered from existing organizational paradigms, routines, and incentives. This is especially important when teams are working on disruptive innovations whose features appear inconsistent with current customer requirements. Without such isolation, NPD teams may face undue pressure to conform to existing organizational practices, reinforcing the organization's current capabilities rather than building new areas of competence. 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 other divisions of the firm) is important to enable the leveraging of ideas across teams, fostering the identification of valuable synergies. This is consistent with prior research that suggests that iterating between centralized and decentralized structures may result in greater long-term learning than either centralization or decentralization alone (e.g., Bartlett & Ghoshal, 1990; Siggelkow & Levinthal, 2003).

The results here are also generally consistent with recent work that suggests that small-world networks may be valuable for creativity and innovation. As argued in Schilling and Phelps (2006a, 2006b), the redundant connections created by clustering help to increase the information transmission capacity of a network, improving the speed of access and degree to which information can be meaningfully understood and utilized. Nonredundant connections in a network, on the other hand, help to increase the scope of new information that can be accessed. Small world network properties help to resolve the tradeoff between transmission capacity and scope of accessible information. Consistent with this, Schilling and Phelps (2006a) found that small-world properties in alliance networks were significantly related to the patenting output of firms in such networks.

Fundamental Linkage to Evolutionary Theory

Our study also indicated that when the interpersonal network is characterized by a subgroup structure with a modest amount of random cross-group links, learning within and between subgroups played out in a process that bore remarkable symmetry to Sewall Wright's shifting balance theory. Wright (1964) attempted to explain how a species could make adaptive walks from a lower peak to a higher peak on a fitness landscape—this has been conceived as a fundamental problem in evolutionary biology. Until he proposed this theory, this had been a puzzling problem. Biologists had known that two variation mechanisms, recombination (through mating processes) and mutation, are not very helpful when the species is trapped into a local peak. Since all the organisms in this state become similar to one another, recombination will only generate offspring that look like their own parents. Then, the whole species cannot move out of the local peak. Mutation may be the only hope. But it is often deleterious, especially when mutants are incompatible with existing genes.

Wright argued that there should be an additional variation mechanism, what he called “demes,” or semi-isolated local populations. Wright described (1964): “Most species contain numerous small, random breeding local populations (demes) that are sufficiently isolated (if only by distance) to permit differentiation…” Our numerical results are quite consistent with Wright’s intuition. But the critics argued that his conceptualization of the adaptive landscape is not precise enough in characterizing a space. Martin, Lienig and Cohoon (2000: 102-103) argued:

> First and foremost, it is not clear what the topology of the underlying space should be. Wright (1932) considers initially the individual gene sequences and connects genetic codes that are ‘one remove’ from each other, implying that the space is actually an undirected connected graph. He then turns immediately to a continuous space with each gene locus specifying a dimension and with units along each dimension being the possible allelomorphs at the given locus. Specifying the underlying space to be a multidimensional Euclidean space determines the topology. However, if one is to attempt to make inferences from the character of the adaptive landscape (Radcliffe 1991), the ordering of the units along the various dimensions is crucial. With arbitrary orderings the metric notions of *nearby* and *distant* have no clear-cut meaning…

Given no clear-cut meaning of the underlying space, it is hard to make sense out of what semi-isolation means. By adopting the variants of modified connected-caveman graph (i.e., subgroup structure), we clarified what the underlying space would mean. In organizational context, the space means how individuals interact with one another. It also means whether subgroups are nearly isolated or not or whether a subgroup structure exists or not. Furthermore, our study provided a more precise operational definition of semi-isolation: a subgroup structure with a moderate number of random, cross-group links. Put another way, semi-isolation satisfies what complexity theorists would call “small-world” network properties—high clustering with a short characteristic path length. In particular, our numerical analysis indicates that this definition is not specific to an arbitrarily chosen parameter value, but that there is a broad interval of such values where evolutionary outcomes are relatively high. Therefore, we believe we contributed to this literature by proposing how Wright's conceptualization could be numerically reconstructed in a parsimonious way.

CONCLUSION

The research here suggests that organizational learning is a function of both the propensity of individuals to learn from one another, and the network topology that determines who learns from whom. Both dimensions can significantly influence the degree to which an organization exploits well-known competencies or explores unfamiliar terrain. Furthermore, these dimensions interact, enabling levels of one to influence the effect of the other. Future studies of organizational learning should thus attempt to account for both the learning rate and the underlying network topology; doing so can help us greatly advance our understanding of the balance between exploration and exploitation.

REFERENCES

Ahuja, G. & Lampert, C.M. 2001. Entrepreneurship in the large corporation: A longitudinal study of how established firms create breakthrough inventions. ***Strategic Management Journal***, 22: 521-543.

Ancona, D.G. & Caldwell, D.F. 1992. Bridging the boundary: External activity and performance in organizational teams. ***Administrative Science Quarterly***, 37:634-665.

Argote, L. 1999. ***Organizational learning: Creating, retaining and transferring knowledge***. Boston: Kluwer Academic Publishers.

Argyris, C., Schon, D. 1996. ***Organizational learning II: Theory, method, and practice***. Reading, MA: Addison Wesley.

Ashby, W. R. 1956. *Introduction to Cybernetics*. London: Chapman and Hall.

Barabasi, A.L. 2002. Linked***: The new science of networks***. New York, NY: Perseus Publishing.

Bartlett, F.C. 1932. ***Remembering: A Study in Experimental and Social Psychology***. London: Cambridge University Press.

Bartlett, C.A. and Ghoshal, S. 1990. Managing innovation in the transnational corporation. In C.A. Bartlett, Y. Doz, and G. Hedlund (Eds.), ***Managing the Global Firm***. London, NY: Routledge, pp. 215-255.

Bettenhausen, K.L. 1991. Five years of groups research: What we have learned and what needs to be addressed. ***Journal of Management***, 17:345-381.

Boorman, S. 1975. A combinatorial optimization model for transmission of job information through contact networks. ***Bell Journal of Economics***, 6: 216-249.

Borgatti, S.P. & Cross, R. 2003. A relational view of information seeking and learning in social networks. Management Science, 49:432-445.

Bouty, I. 2000. Interpersonal and interaction influences on informal resource exchanges between R&D researchers across organizational boundaries. ***Academy of Management Journal***, 43:50-66.

Bower, J.L. & Christensen, C.M. 1995. Disruptive technologies: Catching the wave. ***Harvard Business Review***, 73(1):43-53.

Brown, J., Duguid, P. 1991. Organizational learning and communities of practice: Towards a unified view of working, learning and innovation. ***Organization Sci.***, 2: 40-57.

Burt, R.S. 1992. ***Structural Holes***. Cambridge, MA: Harvard University Press.

Castore, C.H., Peterson, K., & Goodrich, T.A. 1971. Risky shift: Social value or social choice? An alternative model. ***Journal of Personality and Social Psychology***, 20:487-494.

Crossan, M., Lane, H., White, R. 1999. An organizational learning framework: From intuition to institution. ***Academy of Management Rev***. 24: 522-537.

Daft, R.. Weick, K. 1984. Toward a model of organizations as interpretation systems. ***Academy of Management Rev***., 66: 70-74.

Davis, J.H., Kerr, N.L., Atkin, R.S., Holt, R. & Meek, D. 1975. The decision processes of 6- and 12-person mock juries assigned unanimous and two-thirds majority rules. ***Journal of Personality and Social Psychology***, 32:1-14.

De Sola Pool, I., & Kochen, M. 1978. Contacts and Influence. ***Social Networks,*** 1: 5-51.

Dosi, G. 1988. Sources, procedures, and microeconomic effects of innovation. ***Journal of Economic Literature***, 26: 1120-1171.

Erdos, P. & Renyi, A. 1959. On random graphs. ***Pulicationes Mathematicae***, 6:290-297.

Fleming, L. 2001. Recombinant uncertainty in technological search. ***Management Science,*** 47:117-132.

Goldberg, D.E. 1989. ***Genetic Algorithms in Search, Optimization and Machine Learning*** <http://www.amazon.com/exec/obidos/ASIN/0201157675/illigallab> Addison-Wesley Pub. Co.

Granovetter, M. 1973. The strength of weak ties. ***American Journal of Sociology***, 78:1360-80

Grant, R.M. 1996. Prospering in dynamically-competitive environments: Organizational capability as knowledge integration. ***Organization Science***, 7: 375-386.

Guare, J. 1990. ***Six Degrees of Separation: A Play***. New York: Vintage Books.

Hansen, M.T. 1999. The search-transfer problem: The role of weak ties in sharing knowledge across organization subunits. ***Administrative Science Quarterly***, 44: 82-111.

Hansen, M.T. 2002. Knowledge networks: Explaining effective knowledge sharing in multiunit companies. ***Organization Science***, 13: 232-250.

Hargadon, A. 2003. ***How breakthroughs happen: The surprising truth about how companies innovate***. Boston: Harvard Business School Press.

Hartl, D.L. & Clark, A. G. 1997. ***Principles of population genetics***, third edition. Sunderland, MA: Sinauer Associates.

Hinton, G.E. & Nowlan, S.J. (1987). How learning can guide evolution. In Complex Systems, 1, 495-502.Holland, J.H. 1975. ***Adaptation in natural and artificial systems***. Ann Arbor, MI: University of Michigan Press.

Huber, G. 1991. Organizational learning: The contributing processes and the literatures. ***Organization Sci.***, 2: 88-115.

Kerr, N.L., Atkin, R.S., Stasser, G., Meek, D., Holt, R.W., & Davis, J.H. 1976. Guilt beyond a reasonable doubt: Effects of concept identification and assigned decision rule on the judgments of mock jurors. ***Journal of Personality and Social Psychology***, 34:282-294.

Klimoski, R., Mohammed, S. 1994. Team mental model: Construct or metaphor? ***J. of Management***, 20: 403-437.

Larson, J.R., Christensen, C. 1993. Groups as problem-solving units: Toward a new meaning of social cognition. ***The British Psychological Society***, 32:5-30.

Lazarsfeld P.F. & Merton, R.K. 1954. Friendship as a social process: a substantive and methodological analysis. In M. Berger (Ed.), ***Freedom and control in modern society***, 8–66. New York: Van Nostrand.

Levinthal, D. & March, J.G. 1981. A model of adaptive organizational search. ***Journal of Economic Behavior and Organization***, 2: 307-333.

March, J.G. 1991. Exploration and exploitation in organizational learning. ***Organization Science***, 2:71-87.

March, J. & Simon, H. 1958. ***Organizations***. New York: Wiley.

Martin, W.N., Lienig, J. & Cohoon, J.P. Island Models: Evolutionary algorithms based on punctuated equilibria. In Back, T., Fogel, D.B. & Michalewicz, Z. E. (eds.) ***Evolutionary Computation 2***: ***Advanced Algorithms and Operators***, Philadelphia, PA: Institute of Physics Publishing, 101-124.

McPherson, M., Smith-Lovin, L., & Cook, J. 2001. [title] ***Annual Review of Sociology***, 27(1):415-439.

Mezias, S.J. & Glynn, M.A. 1993. The three faces of corporate renewal: Institution, revolution, and evolution. ***Strategic Management Journal***, 14: 77-101.

Milgram, S. 1967. The small world problem. ***Psychology Today***, 2,60-67.

Nelson, R.R. & Winter, S. 1982. ***An Evolutionary Theory of Economic Change***. Cambridge, MA: Harvard University Press.

Newman, M. E. J. 2000. Small worlds: The structure of social networks. Santa Fe Institute working paper 99-12-080

Perkins, D.N. 1995. Insight in minds and genes. In Sternberg, R.J., Davidson, J.E. (eds.) ***The nature of insight***. Cambridge, MA: MIT Press, 495-534.

Reagans, R. & McEvily, B. 2003. Network structure and knowledge transfer: The effects of cohesion and range. ***Administrative Science Quarterly***, 48:240-267.

Reagans, R. & Zuckerman, E. 2001. Networks, diversity, and productivity: The social capital of corporate R&D teams. ***Organization Science***, 12:502-517.

Reuf, M., Aldrich, H.E. & Carter, N. 2003. The structure of organizational founding teams: Homophily, strong ties, and isolation among U.S. entrepreneurs. ***American Sociological Review***, 68(2):195-222.

Sandelands, L., Stablein, R. 1987. The concept of organization mind. ***Research in the Sociology of Organizations,*** 5: 135-161.

Schilling, M.A. & Phelps, C. 2006a. Interfirm knowledge networks and knowledge creation: The impact of "small-world" connectivity. New York University working paper.

Schilling, M.A. & Phelps, C. 2006b. Small-World Networks and Knowledge Creation: Implications for Multiple Levels of Analysis. New York University Working Paper.

Siggelkow, N. & Levinthal, D.A. 2003. Temporarily Divide to Conquer: Centralized, Decentralized, and Reintegrated Organizational Approaches to Exploration and Adaptation. ***Organization Science***, 14:650-669.

Solomonoff, R. & Rapoport, A. 1951. Connectivity of random nets. ***Bulletin of Mathematical Biophysics*** 13:107-17

Strogatz, S. 2001. Exploring Complex Networks, ***Nature***, 410: 268-276

Stuart, T. & Podolny, J. 1996. Local search and the evolution of technological capabilities. ***Strategic Management Journal***, 17(summer special issue): 21-38.

Uzzi, B. 1997. Social structure and competition in interfirm networks: The paradox of embeddedness. ***Administrative Science Quarterly***, 42: 35-67.

Uzzi, B. & Spiro, J. 2005. Collaboration and creativity: The small world problem. Forthcoming in *American Journal of Sociology*.

Watts, D.J. 1999a. Networks, dynamics, and the small-world phenomenon. ***American Journal of Sociology***, 105: 493-528.

Watts, D.J. 1999b. Small worlds: ***The dynamics of networks between order and randomness***. Princeton, NJ: Princeton University Press.

Watts, D. & Strogatz, S. 1998. Collective dynamics of 'small-world' networks. ***Nature*** 393:440-42.

Wegner, D.M. 1987. Transactive memory: A contemporary analysis of the group mind. In B. Mullen, G. R. Goethals (Eds) ***Theories of group behavior.*** New York: Springer-Verlag.

Williams, K. & O'Reilly, C.A. 1998. Demography and diversity in organizations: A review of 40 years of research. ***Research in Organizational Behavior***, 20:77-140.

Wright, S. 1932. The roles of mutation, inbreeding, crossbreeding and selection in evolution. ***Proceedings of the VI International Congress of Genetics***: 356-366.

Yamaguchi, K. 1994. The flow of information through social networks: Diagonal-free measures of inefficiency and the structural determinants of inefficiency. ***Social Networks***, 16: 57-86.

Appendix. Parameters for Simulations

Figure	n	b	m	β	p	s
2	1000	7	100	Variation	0.3	5
3	1000	7	100	Variation	0.3	5
4	1000	7	100	Variation	0.3	5
5a	1000	7	100	Variation	0.3	5
5b	1000	7	100	Variation	0.3	5

FIGURE 1

Various Types of Organizational Structure

a) Nearly-isolated subgroup structure

b) Semi-isolated subgroup structure with randomly rewired links

c) Random network: Network structure without subgroup identity

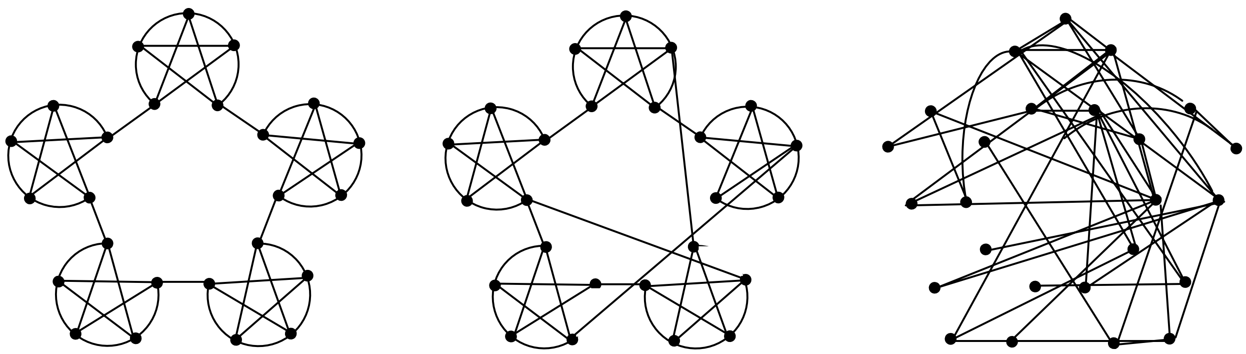

FIGURE 2

Typical Simulations

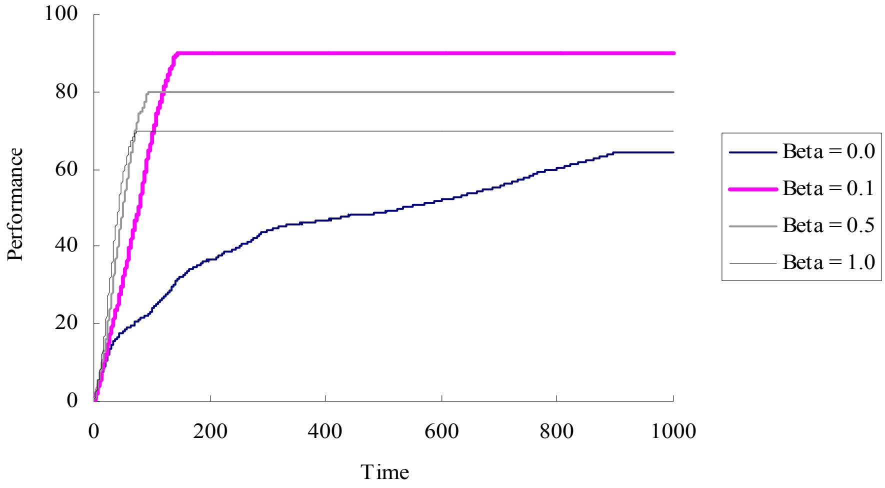

FIGURE 3

Organizational Structure and Learning Outcome

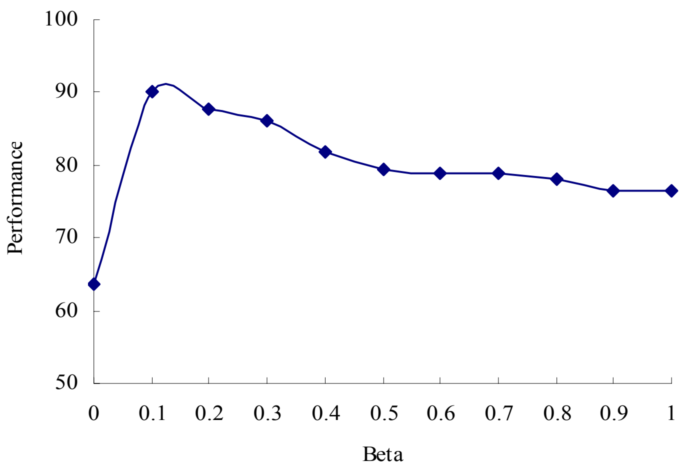

FIGURE 4

Cumulative Frequency Distributions of Learning Outcomes

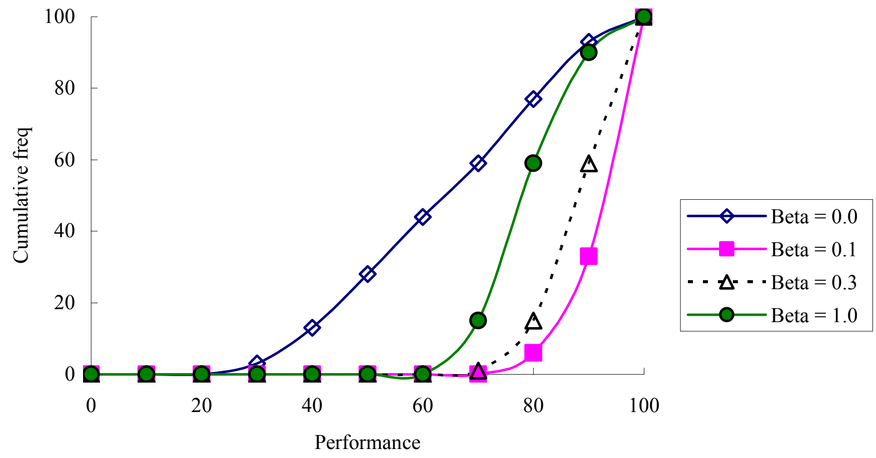

FIGURE 5

(a) Diversity of Beliefs over Time:
Typical Simulation Run

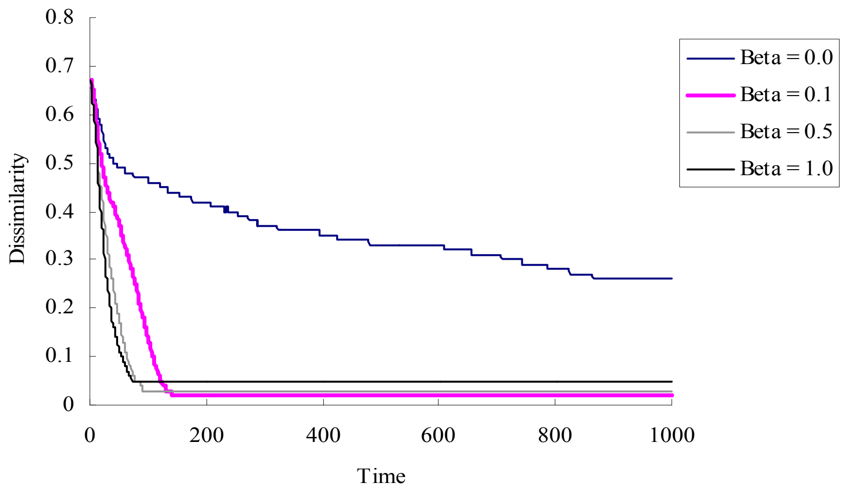

(b) Ten-period Average of Dissimilarity