The diffusion of competing technological innovations in a network: Exploration versus exploitation revisited

Andrea Masini[1)] Michael T. Pich[2)]

1) Ph.D. Candidate, INSEAD (Andrea.Masini@insead.edu)
2) Assistant Professor, INSEAD (Michael.Pich@insead.edu.sg)

Abstract

The main purpose of this study is to examine the effect of network structure on learning processes that occur through the exploration and the exploitation of competing innovations within an organization. We couch this work in the area of organizations and intra-firm diffusion of technology, although it is generally applicable to any group of actors who can be positioned within a social network.

Technology adoption and diffusion occurs through the individual choices of a number of boundedly rational actors that choose among multiple, competing alternatives with unclear characteristics. Due to bounded rationality, these actors must rely on their own experience and on that of adjacent actors in the network to identify and choose the technologies that they feel will maximize their own performance. The repeated implementation of these decisions gives rise to a technology diffusion pattern that ultimately determines the performance of the organization.

We examine the influence of three network structures (central, dense and linear) under three technological environments on three measures of technology diffusion: (i) *performance*, equal to the sum of individual actor performances; (ii) *order*, measured as the extent to which a common technology is exploited across the organization; and (iii) *robustness*, measured as the probability that superior technologies are retained somewhere in the organization. Through numerical simulation we demonstrate that the efficacy of any network structure is contingent on the particular technology environment within which it operates and on the particular measure of diffusion of interest. The results also suggest that in environments that are both *complex* and *dynamic*, the network structures that facilitate the exchange of information may not necessarily be the most effective for guaranteeing the survival of the most promising technologies.

1. Introduction

The increasing turbulence and the high levels of rivalry among competitors that characterize today's markets, together with the impressive rate of change typical of many industry segments, oblige modern firms to continuously search for new and more effective ways to improve their performance. As most of the primary sources of this environmental instability are situated in the technological domain, the ability to quickly identify new technologies[1], to adopt the ones that guarantee superior performance, and to exploit them profitably is now regarded as a necessary condition for the very survival of an enterprise (Doering and Parayre 2000).

However, this is not an easy task. Both the academic literature and the popular press report numerous examples of firms that – after being unable to manage a process of technological innovation – have been inexorably driven out of business (Christensen 1997). Identifying, adopting and exploiting new technologies is "de facto" a critical process of organizational learning. Actors in this domain are confronted with the classic exploration-exploitation dilemma: shall they abandon an old and well-mastered technology for an unknown, though potentially superior one? "Both exploration and exploitation are essential for organizations, but they compete for scarce resources" (March, 1991: 71).

These decisions, which are already challenging when they are restricted to a limited number of well-known alternatives, become significantly more difficult when the technological landscape is complex and uncertain. In complex technological environments, actors have little information about the complete set of technologies potentially eligible for adoption (Milliken 1987; Wholey and Brittain 1989), and are often obliged to base their choices on the experience accumulated by others. Furthermore, in environments of limited rationality, aspiration levels, or targets, become driving forces behind the

---

[1] We use the term "technology" in its broadest sense. That is, a technology may be a tool *per se*, a process, but also a set of procedures or routines to manage an organization or one of its units (Barley 1986).

choice to explore new technologies or to exploit known technologies (March 1991; Cyert and March, 1963).

Regardless of whether one considers the individual agent or the firm as the focal adopter (i.e. regardless of whether one refers to intra or inter-organizational diffusion processes), the topology of the social network within which this actor operates has a profound influence on information exchange and aspiration levels (Coleman, Katz and Menzel, 1966; Burt, 1987), and hence, on an actor's ability to promptly recognize and adopt the most profitable technologies (Rosenkopf 2000). Indeed, it is especially in these instances that network effects become "central to the process of technological development, as they provide the evaluative information that many organizations need to reach decisions" (Midgley, Morrison et al. 1992: 533).

With these premises, it is not surprising that understanding and leveraging network properties for exploring and exploiting new technologies has become a fundamental priority for any firm that intends to navigate successfully in environments of high technological turbulence. Yet, in spite of this, and of the large amount of literature dedicated separately to the diffusion of innovations in networks, few studies have connected these themes to examine innovation diffusion processes that concern multiple technologies with uncertain characteristics.

The purpose of this paper is to contribute to filling in this lacuna by examining the impact of different network structures on innovation adoption processes that occur through sequential choices among *multiple, ambiguous, competing* alternatives. We couch this work in the area of organizations and intra-firm diffusion of technology, although it is generally applicable to any group of actors who can be positioned within a social network, be they people who exchange information with their colleagues or entire organizations that interact with their business partners. As such, the implications are relevant both to organizations that seek to optimize their internal structure and, also, to firms that are interested in understanding how the complex set of relationships with their business partners may support or interfere with their ability to adopt and exploit new technologies.

The remainder of the paper is organized as follows. In section 2 we first review the existing literature on the topic, and then outline the structure of our work and its contribution to this literature. Sections 3 to 5 contain the fundamental building blocks of our investigation. These are, respectively: an agent-based model of technology adoption, a framework to characterize the technological environment in which the process occurs, and a stylized description of the different network structures that are to be compared. In Section 6 we derive and discuss a set of operational measures that are used to further analyze the behavior of the different structures. Finally, in Section 7 we present the major findings and discuss managerial implications.

2. The diffusion of innovations in networks

The study of the diffusion of innovations enjoys a long history in the social sciences. Disciplines as disparate as sociology, marketing, technology management, strategy and communication have all examined the spread or adoption of innovations within a population of actors, be they people, organizations, communities or schools (Mansfield 1961; Bass 1969; Rogers 1976; Teece 1980; Di Maggio and Powell 1983; Tolbert 1983). The nature of this research is very diverse. While some scholars have dedicated attention to mathematical models with the general objective to identify the form that best mimics historical diffusion patterns (Fisher and Pry, 1971; Blackman 1972; Mahajan, Muller et al. 1990), other studies have examined the predictors of the adoption of a *particular* innovation over time, arguing that certain characteristics of the innovation or of the actors affect the pattern and timing of diffusion within the target group (Robertson and Gatignon 1986; Labson and Gooday 1994).

A subset of these researches has extended the study of general processes of diffusion to examining the role of social networks on patterns of adoption. Since the pioneering work of Coleman, Katz and Menzel (1966) and Burt (1987) who identified, respectively, cohesion and structural equivalence as the primary drivers of the diffusion of an innovation through a population of doctors, others have carried on the tradition by examining the spread of a single innovation through networks of doctors (Strang and Tuma 1993), corporate directors (Davis 1991), managers (Galaskiewicz and Wasserman 1989), and even cities (Knoke 1982).

The increasing importance of these phenomena, the call for further research on intra-organizational diffusion processes (Dierickx, Cool et al. 1997), and the acceptance of computer simulation as an appropriate tool for analyzing social systems have recently induced scholars to make extensive use of this instrument to better characterize the influence of social networks on diffusion processes and to compare the behavior of alternative structures, both in general (Midgley, Morrison et al. 1992; Abrahamson 1997) and in relation to the internal structure of organizations (Watkins and DeCanio 1998; DeCanio, Dibble et al. 2000).

For instance, using the traditional Bass model to mimic innovation adoption, Midgley (1992) investigated the impact of different network topologies and alternative models of social contagion on the spread of a single innovation. Watkins and DeCanio, (1998) simulated the diffusion of an innovation through four different structures that allow different levels of information exchange to occur among agents, and suggested that the configurations that guarantee maximal transfer of information are also optimal for diffusing an innovation *only* when the processing capabilities of the agents are sufficiently large.

Abrahamson and Rosenkopf (1997) also focused on the diffusion of a *single* innovation through a population of interconnected actors and suggested that the structural properties of the network may explain some lock-in effects that are not justified by traditional theories of diffusion. Although they did not aim at comparing alternative configurations, Abrahamson and Rosenkopf made an explicit effort to provide a more precise characterization of the adoption process, which occurs as a result of an individual choice of the agents - based on their threshold level of adoption - rather than as the consequence of an exogenous probabilistic event.

This stream of literature is symptomatic of a general urge to provide decision-makers with sound instruments that, by incorporating the features of real diffusion processes, may help them handle complex innovation decisions and identify structures that are most beneficial to this purpose. However, as distinguished as they are, all the models developed so far display certain limitations that leave interesting questions unanswered, hence indicating avenues for further research.

First and foremost, most studies analyze or model the diffusion of a *single* innovation. However, in reality firms are confronted with multiple innovations that "compete" simultaneously for adoption. This is particularly true when the innovations that compete for diffusion are technologies. In these instances, firms typically need to make the fundamental choice between the further *exploitation* of technology already in use and the *exploration* of new technologies that are potentially more attractive but also relatively unknown to the adopter.

Second, as suggested by (Rogers 1976) these studies often "lack concepts and propositions that reflect a process orientation" (p.294), thus calling for increased study of the behavioral processes that underlie adoption decisions and, hence, patterns of diffusion. For instance, with the exception of a few studies (c.f. Abrahamson and Rosenkopf, 1997) , network studies typically mimic adoption processes at a very aggregated level—that is, without accounting for individual choices made by potential adopters—thus neglecting behavioral phenomena that often are among the primary responsible for many observed diffusion patterns.

Third, most papers do not use performance indicators to compare the effectiveness of alternative network structures, thus implicitly assuming that the spread of an innovation is always beneficial for the population of adopters (see DeCanio, et. al. (2000) for a notable exception).

Fourth, the models developed so far do not usually account for the "possibility that potential adopters might update [...] forecasts based on the number of adoptions they learned about through their networks" (Abrahamson, 1997: 307). Hence, they overlook one of the key properties of social networks, namely the fact that they enable actors to revise their estimates upwards or downwards, thus generating more or fewer adoptions of a given innovation (Chatterjee and Eliashberg 1989; Oren and Schwartz 1989).

Finally, most researches do not contemplate instances where the returns from the innovations adopted may increase over time as a function of the increased experience accumulated by their users. Again, although perfectly sound for consumer products, this assumption is not tenable when one considers the diffusion of a new technology. In these circumstances, the accumulation of experience with a system enables a user to realize increasing returns, thus automatically decreasing the user's willingness to abandon it in favor of a new one. Indeed, in spite of the extensive literature on the learning curve (Yelle 1979) and on learning in general (Huber 1990; Dodgson, 1993), most models that analyze the diffusion of complex technologies have systematically neglected these phenomena (see Levin (2000) for a notable exception).

This paper aims to complement and extend the research on networks and innovation diffusion along the lines suggested above. Our investigation begins from Abrahamnson and Rosenkopf's (1997) observation that "further research could explore, via computer simulation, the simultaneous diffusion of competing variants of an innovation across networks of varying structure" and verify under what circumstances one variant could "prevail over a competing, possibly technologically superior, variant" (Abrahamson, 1997: 307).

Building upon this observation, we consider an innovation adoption process that occurs through sequential choices by autonomous actors among multiple, competing technologies. We put forward a process model of innovation diffusion that offers a causal explanation for the time-varying adoption behavior of actors in a population by explicitly accounting for the effects of bounded rationality (Simon, 1969) and experiential learning (Levy, 1965).

Our objective is to understand whether the network structures that have been shown to maximize the exchange of information among agents (i.e. dense networks) are also beneficial when the network faces the diffusion of innovative technologies whose characteristics are fundamentally uncertain and, in particular, when the individual adopters are confronted with the classic "exploration-exploitation" tradeoff (March, 1991; March and Levinthal 1993). Towards this end, we simulate the behavior of three antithetical network configurations and we compare their relative performance in three different technological "regimes."

In the following three sections, we present the three fundamental building blocks that constitute the basis of our analysis. First we elaborate a model of technology adoption that describes in a precise fashion the mechanism through which individual choices determine the behavior of the network in which innovations diffuse. Then, we propose a framework, based upon the characteristics of the technologies eligible for adoption, that categorizes the technological environments where the diffusion process takes place. Finally, building upon existing literature on organizational learning, we derive a set of operational measures that are used to evaluate the effectiveness of the different network structures under analysis.

3. An agent-based model of technology adoption

Since the seminal work of Jim March and Richard Cyert (1963), organization theorists have recognized that firms are not monolithic profit-maximizing entities but, rather, complex and structured systems whose overall behavior is affected by the decisions made sequentially and independently by their members (Cyert and March, 1963; Cohen, March and Olsen, 1972; March and Olsen, 1975).

We posit an organization composed of, more or less, autonomous agents (actors) who utilize, possibly differing technologies to achieve some level of performance (Malone and Smith, 1988). The performance level of the overall organization is the sum of the many individual performance levels of the actors.[2] While in theory the actors might care about the overall performance of the organization, their immediate concern is their own performance within the organization.

Organization performance is determined by the individual technology choices of its actors. However, because these actors operate within a complex and uncertain technological environment, they are not aware of the complete set of technologies available to them, nor do they fully understand the performance potential of those technologies of which they are aware. Moreover, because the actors are boundedly rational, the actors in the network must rely on their own experience and the experience of adjacent actors in their social network. The repeated implementation of these decisions gives rise to a technology diffusion pattern that ultimately determines the learning process, and performance, of the organization. It is precisely the effect of network structure on this process that we aim to examine.

3.1 The organization as a social network

We model the organization as a network, that is, a set **G** of M actors (nodes) that interact among each other. For purposes of simplicity, one can assume that the relations among members are binary and symmetric—that is, either a connection exists or it does not. The whole organizational structure is thus easily represented by means of a symmetric $M \times M$ adjacency matrix **G**. The generic element $g_{n,m}$ is equal to 1 when there is interaction between actor n and actor m, and 0 otherwise. Actor m's "neighbors" are defined as the set G_m, where:

$$G_m = \{n \in G \mid g_{m,n} = 1\}.$$[3]

For the purposes of our investigation we assume that the network structure shapes information exchange and aspiration levels, but that it does not subsume the establishment of a hierarchical order in the organization. That is, the agents are assumed to be perfectly alike with respect to their level of authority and their operational capabilities[4].

The existence of a connection between actors m and n ($g_{m,n} = g_{n,m} = 1$) allows each actor to see the technology and performance level of the other. This has two implications. First, it allows each actor to become aware of new technologies and provides them with another data point on performance for known technologies. Second, it influences their aspiration level. Both of these will be explained in more detail below.

3.2 Technology, performance and the learning curve

Following the traditional agent-based view of the firm, we consider our focal organization as a network of actors who "perform tasks in order to achieve goals" (Malone and Smith, 1988). We assume that the target task is the same for all the agents and that a set $\boldsymbol{K}$ of technologies is available to complete these tasks. The performance of each actor is the result of both the intrinsic potential of the particular technology $k \in \boldsymbol{K}$ that they are using and the agent's ability to exploit it—that is, we posit a *learning curve* for each actor-technology pair (Levy, 1965; Yelle, 1979; Muth, 1986; Adler and Clark, 1991).

The technology potential is assumed to be predetermined and outside the control of the agent. However, it can be exploited more or less efficiently by the agents, depending on their *experience* with the system of interest. As an agent accumulates experience with a particular technology, they may obtain increasing returns from a technology because of the progressive acquisition of knowledge (Mukherjee, Lapre et al. 1998) and the discovery of appropriate routines to exploit the system (Zollo and Winter 1999). As a result, the performance realized by an agent m in a given period depends upon the type of technology adopted and on the experience that the agent has accumulated with that particular system. However, the fact that a technology is repeatedly adopted does not guarantee that its user's performance will increase with certainty. The phenomenon is also influenced by stochastic variability (Mazzola and McCardle, 1997).

We propose a learning curve that builds upon the well-known deterministic model first suggested by Levy (Levy, 1965). It models the increase in performance experienced by the user of a technology as a function of the maximum potential of the technology itself, the cumulative experience of the agent and a stochastic element that captures the effect of exogenous and unpredictable events on the process:

$$X_{m,t}^{k} = X_{\max}^{k}[1 - e^{-(a^k + b^k Q_{m,t}^{k} + x_t^k)}]$$

where:

$X^k_{m,t}$ = performance realized by agent m at time t through the exploitation of technology k;

---

[2] While we recognize that for some technologies there will exist network externalities, we do not include them here for two reasons: first, not all technologies exhibit network externalities; and second, we purposely chose to keep the model simple, leaving such extensions for future research.

[3] It will simplify notation if we include actor m in its own neighborhood.

[4] While we recognize the potential of structure equivalence (Burt, 1987) to influence aspiration levels in a social network, our focus is mainly on organizational structures as communication channels within a complex and uncertain technological environment: removing this assumption would introduce a second-order effect that may confuse the phenomena we wish to observe.

$X^k_{\max}$ = technology performance “asymptote”—that is, the maximal expected performance achievable with technology k;
a^k = initial level of performance for technology k;
b^k = rate of improvement (learning rate) for technology k;
$Q^k_{m,t}$ = number of previous periods during which actor m has used technology k (cumulative experience); and
$x^k_{m,t}$ = random disturbance in actor m’s performance utilizing technology k in period t.

An actor’s cumulative experience stems directly from their decision to use a particular technology in a particular period:

$$Q^k_{m,t} = \sum_{l=0}^{t-1} u^k_{m,l}$$

where $u^k_{m,l}$ is an indicator function that accounts for the basic decision of the agent. It is equal to 1 if actor m uses technology k at time l and 0 otherwise.

According to the model proposed, each agent that starts the process with technology k realizes a performance equal to:

$$X^k_1 = X^k_{\max}[1 - e^{-(a^k + x^k_1)}]$$

and it progresses asymptotically towards the asymptote $X^k_{\max}$ at rate b^k as the actor accumulates experience with the technology. Note that the influence of unpredictable events on performance is greater at the initial stages of adoption, when the innovative system is still largely unknown to the user, but that this effect is then progressively reduced over time as cumulative experience allows the user to get to know the system better and to implement appropriate adjustment actions that attenuate the impact of randomness.

The assumption of homogeneity across actors implies that both the maximum performance and the rate of improvement of an agent that uses technology k are only determined by the intrinsic characteristic of the technology, except for a small random disturbance. Also, the communication process that underlines the network structure permits information exchange, but it does not allow actors to transfer the set of routines and tacit knowledge necessary to fully exploit a technology. This knowledge can only be built individually by each member of the organization, by virtue of direct experience with the technology (Jovanovic and Nyarko, 1996; Szulanski, 1996).

3.3 The adoption process

Actors influence the innovation process of the organization by selecting technologies. The set of technologies available to actor m at time t, denote $K_{m,t}$, depends on the technologies adopted by its neighbors, $n \in G_m$. Each actor can use only one technology at a time—that is, experimentation can take place only through direct use of the technology. The basic technology decision is made sequentially, period after period. After observing both its own previous performance and that of its neighboring actors, each actor revises its technology choice in order to maximize its own performance. Let $a_{m,t}$ denote the technology choice of actor m at time t—that is, $a_{m,t}=k$ whenever actor m is utilizing technology k at time t.

We assume that actor m’s performance in period t, $R_{m,t}$, is directly proportional to technology performance, less the switching cost incurred when changing from one technology to another:

$$R_{m,t} = X^{a_{m,t}}_{m,t} - c\ I^{a_{m,t}}_{m,t}$$

where:

c = cost of switching from one technology to another;[5] and

$$I^k_{m,t} = \begin{cases} 1 & if\ a_{m,t} \neq a_{m,t-1} \\ 0 & if\ a_{m,t} = a_{m,t-1} \end{cases}$$ = indicator function of a switch in technology.

Rather than selecting *ex-ante* a sequence of actions $\{a_{m,t}: t=1,2, \ldots\}$, the perfectly informed and fully rational agent would determine a *policy* that specifies the actions to be taken given the information available at each period. For a fully rational agent, the computation of an optimal policy that solves this problem is a prohibitive task. The simplified case of a single agent that operates its choice among a fixed set of possibilities can be modeled as a multi-armed bandit problem (Gittins 1979; Whittle 1980). In this case dynamic programming and the Gittins index rule provide a viable

[5] Switching costs should not be confused with learning curve effects. Switching costs are incurred even when returning to a previously known technology, and are due mainly to performance losses caused by the interruption of the tasks to be executed.

solution.

However, this closed-form solution is dependent upon a number of assumptions (Ross, 1983) that do not hold in the more complex case described above. When several independent decision-makers exchange information, the performance realized by one agent is observed by its adjacent actors, who then may modify their decisions accordingly. Thus, the composition of the set of technologies available to an agent is non-stationary. It evolves over time as a result of the adoption decisions of its neighbors. Finally, in the traditional version of the multi-armed bandit problem, it is customary to assume that the stochastic processes governing the innovations are of Markovian type and that the agents know the transition probabilities of these processes.

Although convenient for computational purposes, the above conditions do not hold in the case we wish to examine. Furthermore, and more important for the purposes of the present work, the complexity of such a problem exceeds by far the computational capabilities of most decision-makers[6]. Empirical research has shown that in real settings, even in the case of simple problems that could be solved analytically, actual decisions are typically based on simpler rules that depart from the optimal policy: decision makers often adopt policies based on trial and error strategies (Meyer and Shi, 1995).

We model this behavior by means a decision rule that is most commonly used in the economic literature to represent the actions of boundedly rational decision makers who compute expectations of an unknown stochastic process. We assume that agents use a constant-gain expectations rules, in which the performance estimates are recursively updated each period t by combining (i.e. averaging) the previous estimate and the observed performance using a time-invariant rule (Sutton and Barto, 1998; Lettau and Van Zandt, 1999). Thus, in any decision epoch each agent updates its estimate of the performance of technology $k \in K_{m,t}$ at time t, $X^k_{m,t}$, by combining its previous estimate of performance with an increment:

$$X^k_{m,t} = (1-\delta)\ X^k_{m,t-1} + \delta\ \overline{X}^k_{m,t},\ 0<\delta<1,$$

where:

$$\overline{X}^k_{m,t} = \frac{\sum_{n \in \{\Gamma_m\}} X^k_{n,t}}{\sum_{n \in \{\Gamma_m\}} u^k_{n,t}}$$

and $u^k_{n,t}$ is the usual indicator function that accounts for the use of technology k by actor n. The magnitude of the step size δ determines the importance given to recent information with respect to previous estimates. Each agent selects the technology for which the performance estimate minus the switching cost has the highest value:

$$a_{m,t} = \underset{k \in K_{m,t}}{\arg\max}\{X^k_{m,t} - cI^k_{m,t}\}.$$

4. measures of organizational performance

Contrary to much of the previous research in this area, our analysis does not focus on the extent of diffusion of an innovation throughout a network. Rather, we aim to understand whether, in an environment characterized by high technological uncertainty where agents choose among competing technologies, the internal structure of an organization influences the processes of exploration and exploitation, and how this, in turn, affects the organization's ability to *learn*.

Huber suggests that "an entity learns if, through its processing of information, the range of its potential behaviors is changed" (Huber, 1990: p 89). Miner and Haunschild take a population level perspective and define learning as a "systematic change in the nature and mix of organizational action routines in a population of organizations, arising from experience." (Miner and Haunschild, 1995: p 118). Similarly, Zollo and Winter stress the importance of the systematic aspect of the process, by defining a dynamic capability as "a learned pattern of activity through which the organization *systematically* generates and modifies its operational routines in pursuit of improved effectiveness" (Zollo and Winter, 1999: p 10).

For the purposes of this paper, a more context-specific definition of learning is necessary. One contingent to the mechanisms through which this process takes place. We define learning as the process resulting from *individual* actions of the organization members, which enables the organization to *systematically* identify technologies, procedures or processes that are superior to the ones currently in use. The process may occur through the realization of two complementary mechanisms, namely:

Exploitation: The application of and improvement to an existing technology through repeated use and consequent experience accumulation; and

Exploration: The identification and adoption of a potentially superior alternative.

We suggest that organizational learning can be evaluated with respect to three main dimensions; (i) the cumulative performance realized by the organization; (ii) the retention of superior technologies; and (iii) the establishment of an

---

[6] For instance, Watkins and DeCanio (1998) show that even for a problem of much simpler nature finding the optimal organizational structure that maximizes the diffusion of a single innovation would be NP hard.

organizational order.

Cumulative performance measures the overall ability of the organization to utilize technologies to maximize performance over time. It is measured as the cumulative sum of the individual performances of its members:

$$\Pi_t = \sum_{j=1}^{t} \sum_{m=1}^{M} R_{m,t}$$

Although this indicator reflects the ability of the agents to perform their tasks, it does not provide information on the characteristics of the adoption process that has generated it, nor does it shed light on the ability of the firm to generate further profits in the future (e.g. by adopting the technologies that guarantee superior performances in the long run).

The second dimension considers the ability of the organization to retain *superior* technologies within its organization by virtue of having at least one actor exploiting a superior technology. By 'superior,' we mean technologies with the greatest performance asymptote, $X^k_{\max}$. It is measured by means of the probability, or proportion of simulations, that the superior technology survives the selection process. This measure simply evaluates the ability of the superior technology to survive in a competitive environment. It does not take into account the number of adopters that exploit it.

The third dimension, organizational order, reflects the *"systematic"* aspect of learning. It examines whether, by virtue of the interaction among members, the organization systematically (i.e. repeatedly) develops a common code of conduct—that is, to what extent and at what rate does a particular network configuration favor the emergence of "order" in the organization by facilitating the adoption of a common set of technologies (Levinthal, 1997).

In our analogy order is closely related to two other concepts, namely: *convergence* and *equilibrium*. First, we say that the adoption process converges when all the agents in the organization adopt the same technology. Second, we say that the organization is in equilibrium if each agent selects a particular technology (not necessarily the same for all the agents) and does not deviate from its choice in any of the following epochs.

Achieving convergence and equilibrium (all the agents adopt the same technology and never deviate from that choice) correspond to a situation of maximal order in the organization (i.e. a situation of maximal technological uniformity and maximal stability). Organizational order is thus measured as the probability, or proportion of simulations that converge to a common technology.

5. design of the simulation experiments

Contrary to the belief that experimentation – and, hence, the structures that favor this behavior - is "efficacious, perhaps even required, for survival in fast changing and unpredictable environments" (Huber 1990: p. 93), we claim that the efficacy of a particular network structure is not independent of the technology environment within which the firm operates. Indeed, it is not unlikely that network structures that perform one way in one particular environment may perform quite differently in another.

For this reason, we need to analyze our focal organizations under a number of different technological environments. In its most general sense, an "environment" is defined as a system that is outside the organization and influences its behavior and its properties (Ackoff, 1981). We define the *technological environment* as the set of characteristics that portray the technologies available to the focal organization. Specifically, we use the dimensions of complexity and dynamism (Dess and Beard, 1984), where *complexity* reflects the diversity across the performance asymptotes (i.e. differences in $X^k_{\max}$) of the technologies that compete for diffusion, and *dynamism* the change in performance possible (i.e. size of max β^k) in exploiting a technology over time. The larger the diversity across technologies, the greater the complexity of the environment. Similarly, the larger the performance improvement achievable through the accumulation of experience, the more dynamic the environment.

The combination of the above two dimensions naturally identifies four technological environments (see Figure 1). Three of them are relevant to our analysis. In the *mature* regime (low dynamism, high complexity) the network experiences the diffusion of a set of technologies that have different performance potentials, but that are not "dynamic" in that there are no learning curve effects from exploiting the technology. The performance of an agent is affected by stochastic variability, but it does not depend on experience: technologies with average superior performances remain superior even when competing systems are adopted more often. This situation corresponds to an environment where the competing technologies have already reached their maturity and are well-understood. Exogenous factors may affect performance, but these are independent of the actions undertaken by the adopters.

Figure 1: Four possible technology environments

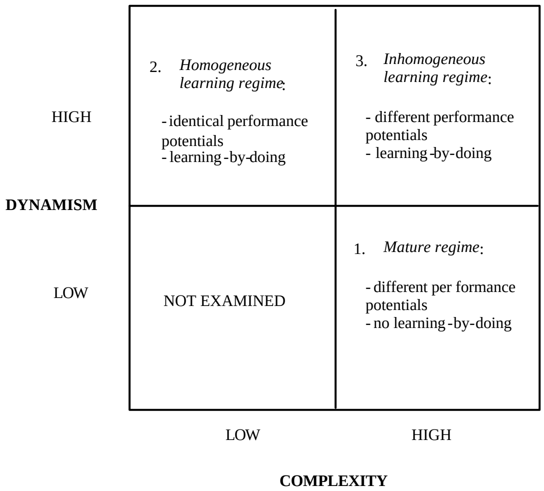

The *homogeneous learning* regime (low complexity, high dynamism) considers a set of technologies that have the same performance potential (i.e. $X_{\max}^{k} = X_{\max}$ for all k) and show identical learning curve effects (i.e. $\beta^k = \beta > 0$ for all k). This implies that both the performance potential and learning rates are identical across all technologies. Finally, the *inhomogeneous learning* regime (high complexity, high dynamism) considers a set of technologies that have different performance potentials (i.e. $X_{\max}^{k} \quad X_{\max}^{j}$ for some k,j), and show different learning curve effects (i.e. $\beta^k \neq \beta^j$ for some k,j). Such a setting is useful to represent the typical situation of radical and incremental innovations that compete for diffusion. An incremental innovation usually hinges upon a number of practices and routines already well established in the adopting organization. Thus, when exploited it does not need a long "warm up" period: it may generate high profits relatively quickly already after few adoption cycles. However, it is also the system with the lowest plateau, i.e. the one with the minimum long term potential.

Conversely, the radical innovation often necessitates of a longer warm up, as the adopting organization requires some time to get to know the new system and to implement the appropriate procedures and routines to exploit it. Furthermore, as a result of its newness, it is more affected by stochastic variability than the incremental innovation. However, by hinging upon a superior technology, the radical innovation has also the highest long-term potential. That is, it is the system that offers to the adopter the best guarantees to achieve long-term advantages over its competitors.

6. types of Network configuration

There has been a long tradition of attempting to classify the various forms of organizational structure that are most commonly encountered in the business world. Needless to say, each taxonomy takes a particular viewpoint. For instance, with respect to the functionality of the different departments inside the organization researchers have identified the functional structure, the product structure and the matrix structures (Daft, 1988). Conversely, when the focal point was the flexibility of the hierarchical relations among agents, organizations have been classified into organic and mechanistic (Burns and Stalker, 1961).

Three basic network structures have been extensively studied with the objective to identify the communication system that is most efficient. These are: the *dense* network (also referred to as the "all channel system", or the "completely connected pattern"), the centralized network with a central *broker* (also referred to as the "wheel pattern"), and the *linear* network (that is named "circle pattern" when a connection is established between the first and the last node of the graph) (Hall, 1991).

These structures are somehow stylized archetypes of three fundamental designs commonly encountered in organizations. The first configuration (dense network) refers to the somehow ideal case of a perfect organization where all the agents are connected among each other. In this setting there is an instantaneous transfer of information across all members. As the maximum distance between two agents is always equal to zero, at any point in time, each actor observes the adoption decisions and the rewards received by *all* the other agents. Although this structure is seldom encountered in real organizations because of its cost, it is often regarded as a sort of theoretical optimum with respect to information transfer.

In the centralized network with a *broker* all the actors are connected to a central broker but not among each other. This configuration reflects the structure of organizations that give emphasis to rapid information exchange across members, but that cannot "afford" the implementation of many connections.

Finally, the third organizational design consists of a perfectly *linear* structure, where each agent is connected to its predecessor and to its follower, except for the two terminal actors who have only one connection. This structure obviously does not facilitate information exchange: albeit it has as many links as the centralized case, the distance between two members of the organization may be as large as N-2.

6.1 Design of the simulation

The effect of network topologies on the learning process of a set of adopters is analyzed by means of a simulation. Eighteen experiments have been designed to simulate the behavior of three basic network structures in three different *technological regimes* for two different values of the switching costs. The adoption model described in the previous sections has been implemented to simulate the diffusion of 7 technologies in a network of 7 actors. The simulation is initialized by assigning a unique technology to each agent. In each technological environment, the parameter α^k is adjusted so that all technologies yield the same expected performance at the time of the first adoption.

Random perturbation	$\xi_t^k \sim N(m = 0,\ s = 0.2)$
Step size	$\delta_t^k = \delta = 0.3$, for all k and t.
Switching cost	*Case i) c = 0;* *Case ii) c = 2s = 0.4;*

Learning curve parameters

Innovation

Experiment 1: Mature technological regime		1	2	3	4	5	6	7
	Plateau	4	5	6	7	6	5	4
	α	0.4	0.4	0.4	0.4	0.4	0.4	0.4
	β	0	0	0	0	0	0	0
Experiment 2: Homogeneous learning regime		1	2	3	4	5	6	7
	Plateau	10	10	10	10	10	10	10
	α	0.2	0.2	0.2	0.2	0.2	0.2	0.2
	β	0.01	0.01	0.01	0.01	0.01	0.01	0.01
Experiment 3: Inhomogeneous learning regime		1	2	3	4	5	6	7
	Plateau	7	3.8	3.8	3.5	3.8	3.8	3.8
	α	0.2	0.4	0.4	0.45	0.4	0.4	0.4
	β	0.01	0.03	0.03	0.04	0.03	0.03	0.03

Table 1: *Parameters of the simulation*

The simulation is terminated after 100 decision epochs[7]. Each simulation run in any of the eighteen experiments is repeated 10 times. The results presented are the average of these 10 runs. The random perturbation $x_{m,t}^k$ that affects the behavior of the learning curves has been drawn from a normal distribution with mean $\mu = 0$ and standard deviation $\sigma = 0.2$. The value of the variance guarantees that the random disturbance – albeit significant - does not overwhelm the learning effect determined by cumulative adoption[8]. The system parameters used in the experiments are summarized in Table 1.

The technique of common random numbers (CRN) has been adopted to reduce the impact of stochastic variability on the final results of the simulation (Law and Kelton, 1991). For this purpose, in each experiment the random perturbations $x_{m,t}^k$ have been drawn from an identical sequence of random numbers. The use of common random numbers guarantees that the different performances of the various network configurations are entirely due to a different intrinsic behavior of the systems and not to pure randomness.

7 This choice is a compromise between two exigencies. On the one hand we need to observe the behavior of the network for a number of decision epochs sufficiently long to allow the occurrence of a significant number of cycles. On the other hand, we have the exigency to consider a time horizon consistent with that of real business processes (i.e. not excessively long).

8 The robustness of the results has been examined by repeating the experiments for different values of the coefficient α. The latter influences the extent to which the random perturbation contributes to determine the performance of the innovations. These additional experiments have confirmed the trend observed and are not reported here.

7. Results

The three dimensions of diffusion (organizational performance, survival of the superior system, and emergence of organizational culture) are used to compare the behavior of the three network structures in each of the three possible technology environments. In order to highlight the differences among configurations we have reported a relative measure of organizational performance, using the fully connected, or dense structure as a benchmark. That is, for each technological regime included in the analysis, we have computed the ratio of the performance generated by a given configuration to the performance generated by the completely connected one.

7.1 The mature technology environment

In the "mature" technology environment, the innovations that compete for diffusion have different plateau performance levels but there are no learning curve effects. In these circumstances an optimal policy should initially emphasize exploration rather than exploitation, so as to rapidly identify the technology that guarantees superior performance.

Figures 2 shows that the centralized, or broker configuration yields, on average, the greatest organizational performance. This, of course, makes sense in that the broker is in a good position to quickly identify the best technology, and given its central position, to make this technology immediately available to the rest of the organization. The linear network performs consistently worse than the broker network, regardless of switching costs. Here, individual exploration is hindered by a lack of connections in the network, and organizational exploration is hindered by the structure of the network (i.e. lack of centrality). Thus, too many actors exploit inferior technologies for too long.
Without switching costs, the fully connected network does eventually approach the performance of the broker configuration. However, when we introduce switching costs, the relative performance of the connected network decreases. The introduction of switching costs can have two possible effects in this case. First, an increase in switching costs will reduce net organizational performance in networks that rely on individual exploration. Continuously trying out new technologies might allow actors to eventually identify optimal solutions, but it also implies the occurrence of higher search costs that reduce profit. Second, in the presence of switching costs, individual exploration may be hindered, resulting in too many actors exploiting inferior technologies for too long.

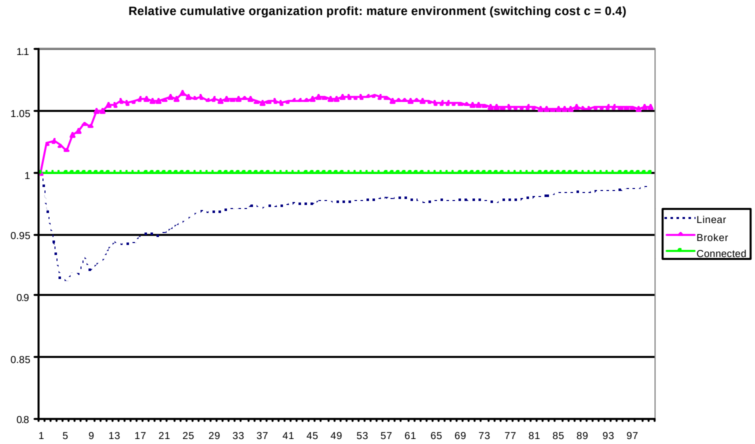

Figure 2: Relative cumulative performance in the mature technology environment.

Table 2, which summarizes each network's ability to retain the superior technology, sheds some light on this phenomenon. The broker network performs best in terms of its ability to retain the superior technology: It was able to retain the superior technology in 90% of the cases, regardless of switching costs. Without switching costs, the connected network performs equally well (90%). However, with switching costs, the connected network retains the superior technology only 60% of the time, indicating that the presence of switching costs is hindering individual exploration for the superior technology. Thus, actors are too quite to settle on an inferior technology rather than switch to a potentially superior one.

The linear network performs the worst in this regard, retaining the superior technology only 60% of the time. This is increased to 70% in the presence of switching costs, due mostly to the fact that switching costs reduce convergence to a common technology from 60% to 40% (see Table 3). Thus, the fact that there is more than one technology present at the end of the decision horizon increases the probability that one of these will be the superior technology.

Convergence to a common technology is achieved in all the experiments for the connected network, regardless of switching costs. The broker network achieves convergence in 90% of the cases both with and without switching costs. Thus, while the broker network outperforms the connected network in terms of organizational performance, the connected network consistently generates higher order in the organization by favoring the retention of a common technology.

Proportion of times that superior technology survives within the organization						
	Connected		Broker		Linear	
	c = 0	*c = 0.4*	*c = 0*	*c = 0.4*	*c = 0*	*c = 0.4*
Complex & Static (Mature)	**0.90**	0.60	**0.90**	**0.90**	0.60	0.70
Homogeneous & Dynamic[*]	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.
Complex & Dynamic (Inhomogeneous)	0.30	0.20	0.10	0.20	**0.40**	**0.40**

[*] In the homogeneous case, all technologies have the same long-run potential; thus, there is no one superior technology.

Table 2: *Organizational retention of the superior technology.*

7.2 The homogeneous technology environment

In this setting a number of technologies with identical performance potential, in terms of their asymptotic performance, compete for diffusion. Given that all the systems guarantee the same returns to adoption, the optimal policy to secure high organizational performance would be to maximize the exploitation of the technology used in the first decision epoch without undertaking any form of exploration. Of course, this policy would also maximize the "disorder" within the system in terms of the diversity of technologies retained. An ideal compromise between the two exigencies would be for the organization to converge as rapidly as possible towards a unique technology and keep exploiting it.

Proportion of times convergence is attained before the 100th decision epoch						
	Connected		Broker		Linear	
	c = 0	*c = 0.4*	*c = 0*	*c = 0.4*	*c = 0*	*c = 0.4*
Complex & Static (Mature)	**1.00**	**1.00**	0.90	0.90	0.60	0.40
Homogeneous & Dynamic	**1.00**	**1.00**	0.90	0.60	0.00	0.00
Complex & Dynamic (Inhomogeneous)	**1.00**	**0.90**	0.50	0.10	0.00	0.00

Table 3: *Emergence of organizational order.*

Among the three configurations, only the completely dense network favors such a behavior. In all the experiments performed, such a configuration enables actors to converge towards a unique technology in a limited amount of time. Although the central network does not perform significantly worse with respect to the proportion of time convergence is attained (see Table 3), it always needs a larger amount of time to achieve it. Finally, in no case did the linear network converge to a common technology, regardless of switching costs.

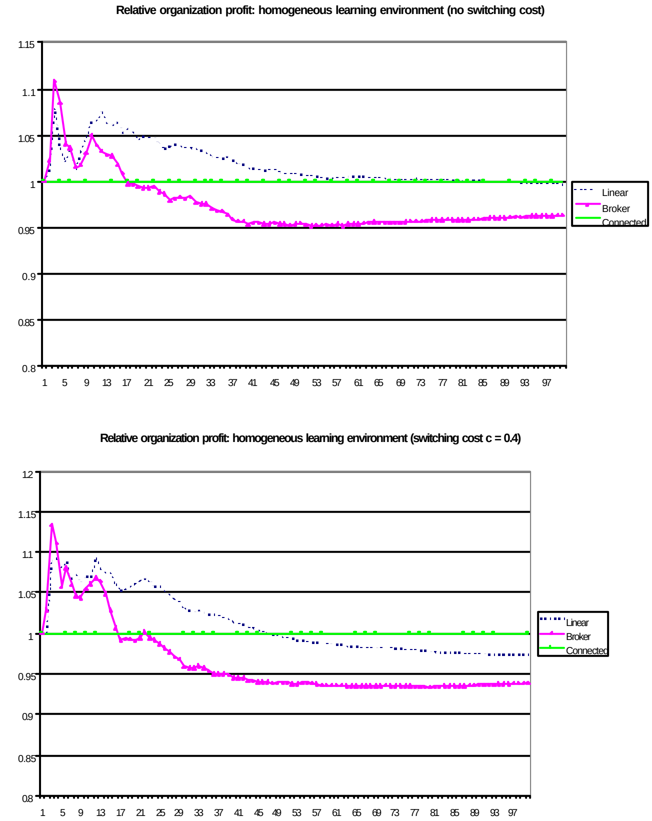

Figure 3: Relative cumulative performance in the homogeneous technology environment.

However, as expected in this technological regime, failing to converge does not necessarily have severe consequences on the organizational performance generated. Figure 3 shows that the fully connected network, after a relatively slow start, eventually performs the best of the three networks. The linear configuration performs reasonably well relative to the other two networks, though the presence of switching costs presents a clear advantage to the connected network. The centralized structure is consistently inferior, except in the very first epochs when the connected configuration "pays" for the high level of individual exploration that it induces. There is a sound explanation for this apparently contradictory effect.

Given the fact that all the innovations retained have the same behavior, an organization would suffer from the lack of order only if the achievement of the latter is impeded by the continuous shift of any single adopter from one technology to another. (This behavior favors exploration, which is useless with identical systems and severely penalizes exploitation). However, the organization does not suffer from "anarchy" if this is caused by the rapid attainment of multiple local equilibria. These occur when all the actors quickly adopt different systems and then keep exploiting them without further switching. Since all the innovations have the same performance, it does not matter on which system exploitation focuses, provided that it occurs quickly.

Indeed Figure 4, which depicts the evolution over time of the average number of different innovations retained in each of the three networks throughout the decision process, suggests the occurrence of a phenomenon of this sort. The connected network rapidly reduces the number of technologies in the network, settling down to less than 2 technologies on average within the first couple of periods. The linear configuration operates the initial selection as efficiently as the broker network, but only reduces the number of technologies in the network to two units. However, the linear network does not explore further once this 'equilibrium' is reached, while the broker network continues to explore almost twice as long as the linear network on its way to converging to a common technology.

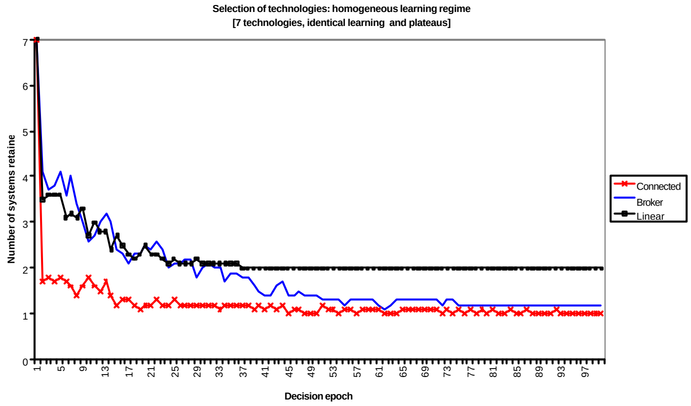

Figure 4: Number of technologies in each of the three networks over time.

By virtue of the experience effect, the agents exploiting a particular technology rapidly accumulate experience and receive high rewards already after few decision epochs. Switching to a new technology after having already exploited an alternative would imply starting over at the bottom of the learning curve. Since this is identical to the technology just abandoned, on average the change would not generate higher profits.

The performance charts in Figure 3 summarize this behavior. The completely connected structure guarantees order, but in order to do so, it wastes some decision epochs to explore all the technologies before it eventually converges. This is reflected by the cumulative profit that is always the lowest in the first decision epochs. The linear configuration does not induce order, but by virtue of the fact that individual exploration is not excessive, it guarantees high profit levels. Conversely, the centralized structure suffers from both the disadvantages of the other two configurations. It spends quite some time in exploration, but it does not do so to an extent that is sufficient to guarantee rapid convergence. As a result, it consistently secures lower profit levels than the other architectures.[9]

7.3 The inhomogeneous technology environment

In the inhomogeneous environment, technologies are characterized by both high complexity and high dynamism. That is, not only do the technologies competing for diffusion differ with respect to their plateau performance ($X^k_{\max}$), but they also have different rates of learning-by-doing (β_k) at the individual level. Based on the results of the previous analysis, one would naturally suspect that even when the competing technologies have different learning rates, the configurations that facilitate the exchange of information perform better than the linear structure. Indeed this is not necessarily the case, at least not with respect to all the dimensions considered.

On the one hand, consistent with previous results, the connected network facilitates the emergence of organizational order (see Table 3). The broker network converges only 50% of the time without switching costs and only 10% of the time with switching costs. Not surprisingly, the linear network failed to achieve convergence in any of the simulations.

On the other hand, the situation is radically different when one examines the performance of the three configurations with respect to their ability to retain the superior technology (see Table 2). Although the linear network does not enable the organization to converge to a common technology, at least within the time horizon considered in the experiments, it is the one that guarantees the highest rate of survival for the superior technologies. In 40% of the runs (both with and without switching costs) at least one agent of the linear network still exploits the best long-term optimal system after the 100th decision epoch, thus preventing its premature extinction. The survival rate is slightly lower for the dense network (30% without switching costs and 20% with switching costs) but it is significantly inferior for the configuration with a broker (10% and 20% for the two cases). In the broker network, individual exploration yields order as members adopt those technologies with the highest exploitation rates.

The phenomenon has important implications and deserves further discussion. To fully capture the reason for this efficiency gap, let us focus on the linear and the centralized networks. From a resource use perspective, the two configurations require the same level of investment for actual implementation, as the number of connections between actors is identical[10] in both cases. However, *ceteris paribus*, the configuration with a central broker facilitates the exchange of information, as the maximum distance between two different actors is always inferior to the one of the linear case. Furthermore, the broker has proportionally more influence on the organization than any other actor in the network: at any time *all* the adjacent actors observe both its adoption decision and the rewards achieved and then use this information to update their sample estimates.

---

[9] Measures related to the identification of the optimal systems are not relevant in this case, given that all the innovations guarantee identical long-term performance.

[10] In a network with N actors the number of edges is equal to N-1 for both configurations.

Let us suppose that the initial adoption of the broker consists of an incremental innovation – i.e. a technology that allows its users to quickly increase the initial performance but that has a lower plateau - whereas one or more of its neighbors adopt a radical innovation. By virtue of the fact that the broker initially realizes the highest profits in the organization, all the other agents are induced to abandon their initial choices in favor of this short-term optimal system, before they have the time to exploit the 'true' optimal system and discover its long term potential.

The organization thus suffers from a sort of self-reinforcing kind of competency trap (Levinthal and March, 1993). As a result of the specialized capability of one of its members (i.e. its central location in the network), the broker network favors the exploitation of the inferior technologies and creates a lock-in situation. The inferior technologies survive the selection process and cause the extinction of the competing alternatives, even though these are potentially superior. The organization is not protected against this phenomenon, precisely because the particular network configuration enables the broker to communicate instantaneously with the totality of the agents.

Conversely, in the linear configuration the information exchange occurs at a slower pace. It may take several decision epochs before an agent at the end of the network discovers the existence of an innovation initially adopted at the other end, even if this is more profitable in the short run. Thus, by virtue of this slower exchange, the linear configuration favors the establishment of *technological niches*, where the radically innovative system is protected from competition at the initial stages of its evolution. The niche allows actors to exploit the radical innovation to an extent sufficiently large to enable it to outperform the incremental innovation (which is concurrently under exploitation in a different "zone" of the organization). By the time the members of the niche discover the existence and the performance of the incremental system, they have already exploited the radical innovation above the break-even point and do not consider shifting convenient any more.

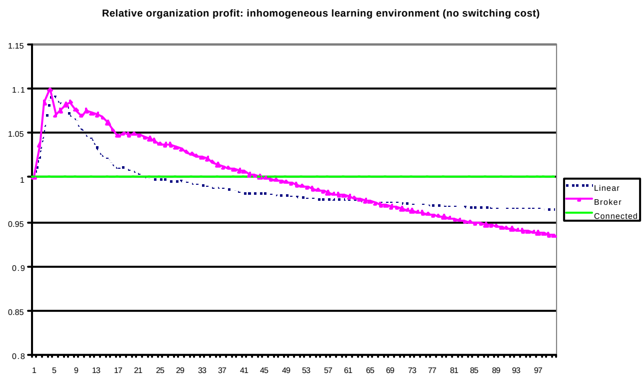

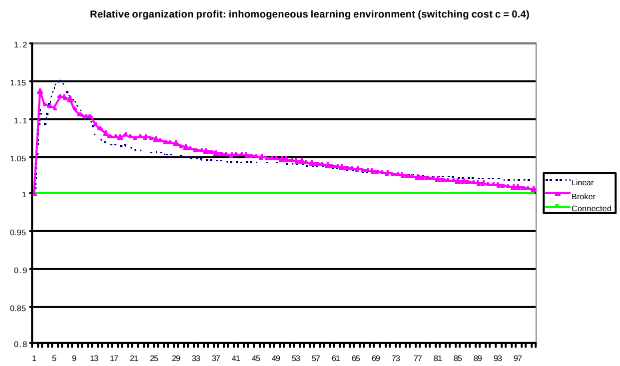

Figure 5: Relative cumulative performance in the homogeneous technology environment

The profit charts reflect this behavior (see Figure 5). Initially, the fully connected network pays a price for its early individual exploration of competing alternatives. Early on, exploitation has a large impact on performance due to learning curve effects. However, the high rate of individual exploration allows the connected network to achieve a balance between retaining 'good' technologies, though not always the superior technology, and exploiting those technologies consistently across the network. When we introduce switching costs, the connected network's reliance on individual exploration comes at too great a cost for it to recoup the initial cost. By the end of the simulation, it is still performing no better than the other two networks.

To summarize, in technological environments that are both complex and dynamic, there does not appear to exist a universally superior network. The choice of the optimal form depends on the objectives of the organization. If the establishment of a common code of conduct is the main goal, then a completely connected structure should be preferred. Conversely, if one is more concerned about the retention of potentially superior technologies, the linear configuration seems to be more appropriate. Only the centralized structure seems to present no advantages in either case.

8. Conclusions

This paper has examined the effect of network structure on the processes of technology adoption and diffusion that occur through individual choices among ambiguous alternatives. More specifically, its major objective is to understand whether the same network structures that have been shown to maximize the exchange of information among agents are also beneficial when a firm faces the diffusion of innovative technologies with uncertain characteristics.

Towards this end, we have first proposed a model of technology adoption and diffusion within a network. Our model is characterized by a number of distinctive features that attempt to address some of the limitations of previous studies. First, we consider the diffusion of multiple, competing innovations. Second, we explicitly account for the fact that innovation diffusion is the result of a sequence of adoption decisions made independently by a population of boundedly-rational agents (Rusmevichientong and Van Roy, 2000). Third, we account for learning-curve effects in the diffusion of technologies where individual agents *learn* as experience with a technology accumulates. Finally, we notice that the effectiveness of a particular network structure on the process of innovation adoption should be evaluated through a multidimensional measure that accounts for the complexity of the phenomenon.

The results of our simulation experiments offer several interesting insights into the processes of exploration and exploitation within various network structures and technology environments. Dense networks, such as the fully connected network examined in our study, facilitate *individual exploration* of technologies early in the diffusion process, while rapidly driving conversion to a common technology across the network. Network configurations that facilitate information exchanges help an organization learn through the development of a common code of conduct. This occurs because information exchange accelerates first hand exploration and forces the organization to adopt the systems that provide the highest immediate rewards.

On the other hand, the analysis also suggests that, in uncertain environments where innovations with different learning rates compete for diffusion, network configurations that do not favor information transfer may perform better. By quickly identifying the systems that offer the highest immediate rewards, centralized structures hamper the development of long term optimal technologies, thus favoring their extinction. Conversely, configurations that do not transfer information quickly engender the development of technological niches inside the organization, which protect innovations from fierce competition and help their long-term growth.

Thus, our results confirm findings from previous studies that no structure is universally superior. The effectiveness of a particular configuration is contingent to the specific technological environment where the firm operates. In complex but scarcely dynamic environments centralized structures appear to be the optimal compromise between cost and efficiency. Conversely, in environments where complexity is low and dynamism is high the above configuration is inferior both to the completely connected one and to the linear structure. Finally, in environments that are both complex and dynamic, the choice of the optimal configuration depends on the firm's objectives. Completely connected structures favor the establishment of organizational order, whereas linear configurations guarantee the survival of superior technologies.

Not surprisingly, the occurrence of switching costs tends to reduce the propensity of the organization to explore. That is, it induces a sort of "band of inaction", that prevents each agent from trying new technologies whose performance estimate is only marginally superior to the one currently in use. As a consequence, switching costs reduce the effectiveness of network configurations that benefit from individual exploration.

This result confirms that the two types of network structures (dense and central on one side, linear on the other) function in an antithetic fashion and exploit two different leverages. On the one hand, the dense and the central structure base their effectiveness upon the rapid search guaranteed by the information exchange process. As a consequence, any effect that attenuates this feature (such as the occurrence of switching costs) proportionally reduces the effectiveness of the two configurations.

On the other hand, the linear configuration owes its effectiveness to the fact that it "protects" radical innovations from external competition in the early stages of their development. In this case the occurrence of switching costs enhances this property, and it further helps the linear configuration retain and exploit optimal technologies.This phenomenon is mainly due to the excess search induced by the inherent structure of the connected network that guarantees high rates of information exchange. In such a configuration, each agent is literally overwhelmed with new information that can be processed effectively. This is consistent with results observed by Abrahamson and Rosenkopf (1997) and DeCanio and Watkins (1998) for the diffusion of one technology. Interestingly, Huberman (1997) suggests that a similar phenomenon is the cause of some famous exceptions to the learning curve phenomenon. In his model, the introduction of a too large number of new "search procedures" may significantly decelerate the learning process of a manufacturing unit and cause cost to increase even when cumulative production increases. As a consequence, in technological regimes that are characterized by high complexity and low dynamism, the centralized configuration suggests itself as the one that offers the best compromise between effectiveness and implementation cost. Indeed, the relatively small improvements achievable in some measures when passing from this configuration to the completely connected one would not seem to justify the investment necessary to establish all the additional connections required to operate the change[11].

References

Abrahamson, E. a. L. R. (1997). "Social network effects on the extent of innovation diffusion: A computer simulation." *Organization Science* 8(3).

Ackoff, R. L. (1981). *Creating the corporate future*. New York, Wiley&Sons.

Adler, P. S. and K. B. Clark (1991). "*Behind the Learning Curve: A Sketch of the Learning Process*." Management Science 37(3): 267-281.

Barley, S. R. (1986). "Technology as an Occasion for Structuring: Evidence from Observations of CT Scanners and the Social Order of Radiology Departments." *Administrative Science Quarterly* 31(1).

[11] For instance, in the case of the 7 actor network considered in the simulation, transforming a central configuration into a dense one would require the establishment of 15 new connections.

Bass, F. (1969). “A new product growth model for consumer durables.” *Management Science* 15 (2).

Blackman, A. W. J. (1972). “A Mathematical Model for Trend Forecast.and Technological Change.” *Technology Forecasting and Social Change*(3).

Bonacich, P. (1972). “Factoring and weighting approaches to status scores and clique identifica-tion.” *Journal of Mathematical Sociology* 2: 113-120.

Burns, T. and G. M. Stalker (1961). T*he management of innovations*. London, Tavistock Publications.

Burt, R. (1987). “Social contagion and innovation: Cohesion versus structural equivalence.” *American Journal of Sociology* 92: 1287-1335.

Chatterjee, R. and J. Eliashberg (1989). “The Innovation Diffusion Process in a Heterogeneous Population: A Micromodeling approach.” *Management Science* 36(9).

Christensen, C. (1997). “The Innovator's Dilemma”. Cambridge, MA, Harvard Business School Press.

Cohen, M. D., J. G. March, et al. (1972). “*A Garbage Can Model of Organizational Choice*.” Administrative Science Quarterly 17: 1-25.

Coleman, J., E. Katz, et al. (1966). “Medical Innovation.”. New York, Bobbs-Merill.

Cyert, R. M. and J. G. March (1963). *A behavioral theory of the firm*, Prentice Hall.

Daft, R., L (1988). *Organization Theory and Design*. St. Paul, West Publishing Company.

Davis, G. (1991). “Agents without principles? The spread of the poison pill through the inter-corporate network.” *Administrative Science Quarterly* 36: 583-613.

DeCanio, S. J., C. Dibble, et al. (2000). “The importance of organizational structure for the adoption of innovations;.” *Management Science* 46(10).

Dess, G. G. and D. W. Beard (1984). “*Dimensions of organizational environment*.” Administrative Science Quarterly 29: 52-73.

Di Maggio, P. J. and W. W. Powell (1983). “The Iron Cage Revisited: Institution Isomorphism And Collective Rationality in Organizational Fields.” *American Sociological Review* Vol. 48.

Dierickx, I., K. Cool, et al. (1997). “Diffusion of innovations within organizations : electronic switching in the Bell System, 1971-1982.” *Organizations Science* 8(5).

Dodgson, M. (1993) "Organizational learning: A review of some literatures." *Organizational Studies* 14: 375-394.

Doering, D. S. and R. Parayre (2000). “Identification and Assessment of Emerging Technologies.”. In: “Wharton on Emerging Technologies.”. G. S. Day, P. J. Schoemaker and R. Gunther. New York, John Wiley &Sons.

Fisher, J. C. and Pry R.H. (1971). “A Simple Substitution Model of Technological Change.” *Technology Forecasting and Social Change*(3).

Freeman, L. (1978). “Centrality in social networks: Conceptual clarification.” *Social Networks* 1: 215-239.

Galaskiewicz, J. and S. Wasserman (1989). “Mimetic processes within an interorganizational field: An empirical test.” *Administrative Science Quarterly* 34: 454-479.

Gittins, J. C. (1979). “Bandit Processes and Dynamic Allocation Indexes.” *Journal of Royal Statistical Society Ser. B*. 14: 148-177.

Hall, R. H. (1991). *Organizations: structures, processes and outcomes*, Prentice-Hall International editions.

Huber, G. P. (1990). “Organizational Learning: The contributing Processes and the Literature.” *Organization Science* 2(1).

Huberman, B. (1997). “*The Dynamics of Organizational Learning*.” Working Paper: Xerox Palo Alto Research Center.

Jovanovic, B. and Y. Nyarko (1996). “*Learning by doing and the choice of technology*.” Econometrica 64(6): 1299-1310.

Knoke, D. (1982). “The spread of municipal reform: Temporal, spatial, and social dynamics.” *American Journal of Sociology* 87: 1314-1339.

Labson, S. and P. Gooday (1994). “Factors influencing the diffusion of electric arc furnace steelmaking technology.” Applied Economics 96: 917-925.

Law, A.M. and D.W.Kelton (1991). *Simulation Modeling and Analysis*, McGraw Hill.

Masini and Pich

Levin, D. (2000). "Organizational learning and the transfer of knowledge: An investigation of quality improvement." *Organization Science* 11: 630-647.

Levinthal, D. A. (1997). "A*daptation on Rugged Landscapes*." Management Science 43(7).

Levy, F. K. (1965). "*Adaptation in the production process*." Management Science 11(6).

Mahajan, V., E. Muller, et al. (1990). "New Product Diffusion Models in Marketing: A Review and Direction for Research." *Journal of Marketing* 54(1).

Malone, T. W. and S. A. Smith (1988). "*Modeling the performance of organizational structure*." Operations Research 36(3): 421-425.

Mansfield, E. (1961). "Technical change and the rate of imitation." *Econometrica* 29: 741-766.

March, J. G. and J. P. Olsen (1975). "*The Uncertainty of the Past: Organizational Learning under Ambiguity*." European journal of Political Research 3: 147-171.

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

March, J. G. and D. A. Levinthal (1993). " The Myopia of learning." *Strategic management journal. Special issue. Winter 1993* 14.

Mazzola, J. B. and K. F. McCardle (1996). "A Bayesian Approach to Managing Learning-Curve Uncertainty." *Management Science* 42(5).

Meyer, R. J. and Y. Shi (1995). "*Sequential choice Under Ambiguity: Intuitive Solutions to the Armed-Bandit Problem*." Management Science 41(5).

Midgley, D. F., P. D. Morrison, et al. (1992). "The Effect of Network Structure in Industrial Diffusion Processes." *Research Policy* 21(6).

Milliken, F. J. (1987). "Three types of uncertainty about the environment: state, effect and response uncertainty." *The Academy Of Management Review* 12(1).

Miner, A. S. and P. R. Hainschild (1995). "*Population Level Learning*." Research in Organizational Behavior 17.

Mukherjee, A. S., M. A. Lapre, et al. (1998). "Knowledge driven quality improvement." *Management Science* 44(11, part 2).

Muth, J. F. (1986). "Search Theory and the Manufacturing Progress Function." Management Science 32(8).

Oren, S. S. and R. G. Schwartz (1989). "Diffusion Of New Products In Risk-Sensitive Markets." *Journal Of Forecasting* 7(3).

Robertson, T. S. and H. Gatignon (1986). "Competitive effects on technology diffusion." Journal of Marketing 50: 1-12.

Rogers, E. (1976). "New product adoption and diffusion." *Journal of Consumer Research* 2: 290-301.

Rosenkopf, L. (2000). "Managing Dynamic Knowledge Networks.". In: "Wharton on Emerging Technologies.". G. S. Day, P. J. Schoemaker and R. Gunther. New York, John Wiley &Sons.

Rusmevichientong, P. and B. Van Roy (2000). "Decentralized decision-making in a large team with local information." Stanford working paper.

Simon, H. 1965. *The Science of the Artificial*. Boston: MIT Press, 2nd edition.

Strang, D. and N. Tuma (1993). "Spatial and temporal heterogeneity in diffusion." *American Journal of Sociology* 99: 614-639.

Szulanski, G. (1996). "Exploring internal stickiness : impediments to the transfer of best practice within the firm." *Strategic Management Journal* 17 (Special issue, winter 1996).

Teece, D. (1980). "The diffusion of administrative innovations." *Management Science* 26: 464-470.

Tolbert, P. a. L. Z. (1983). "Institutional sources of change in the formal structure of organizations: The diffusion of civil service reform, 1880-1935." *Administrative Science Quarterly* 28: 22-39.

Watkins, W. E. and S. J. DeCanio (1998). "Information processing and organizational structure." *Journal of economic behavior & organization* 36(3).

Whittle, P. (1980). "Multi-armed Bandits and the Gittins Index." *Journal of Royal Statistical Society B.* 42: 143-149.

Wholey, D. R. and J. Brittain (1989). " Characterizing Environmental Variation;." *Academy Of Management Journal* 32(4).

Yelle, L. E. (1979). "the Learning curve: Historical Review and Comprehensive Survey." *Decision Science* 10: 302-328.

Zollo, M. and S. Winter (1999). “From Organizational Routines to Organizational Capabilities.” *Insead Working Paper* 99/48/SM.