Firm Organization in Complex and Uncertain Environments*

Jason Barr
Rutgers University, Newark
jmbarr@rutgers.edu

Nobuyuki Hananki
University of Tsukuba
hanaki@dpipe.tsukuba.ac.jp

Troy Tassier
Fordham University
tassier@fordham.edu

May 1, 2006
Verison 1

Abstract

We present a model of the firm as an information processing network. Firms are comprised of two kinds of agents: searchers and managers. Searchers seek to locate projects that will improve the firm's performance; managers process information and make decisions about which projects the firm should undertake. Firm networks are characterized by three variables: the depth and breadth of the organization, and the authority level of decision making. The environment (the set of potential projects) is modeled as a modified NK landscape. Environmental complexity is measured by K, the degree to which project components are related; uncertainty is measured by the probability with which project relationships change over time. In general, given the costs and benefits of search and information processing, we investigate what types of networks perform best given the degree of complexity and uncertainty in the environment.

*The paper has been prepared for the Third International Meeting "Complex Behavior in Economics: Modeling, Computing and Mastering Complexity," COMPLEXITY 2006, Aix en Provence, France, 17 - 21 May 2006. Comments are welcome.

1 Introduction

In his remarks before the Dallas Ambassadors Forum in April 1999, former Federal Reserve Chairman Alan Greenspan stated that the physical weight of U.S. output has been getting lighter and lighter over time. He notes that the changing nature of production and processing - from heavy, natural resource based activities to information production - is allowing workers to maintain their high rates of productivity growth. He calculates that though the economy has grown by a factor of five since the early 1950's, the actual, physical weight of the economy "is evidently only modestly higher today than it was fifty or one hundred years ago" (Greenspan, 1999).

In addition to the changing nature of business activity, large firms (with greater than 500 employees) continue to play an important role in the economy. According to the 2002 Census of U.S. Businesses, there are roughly 5.7 million firms in the economy. Though large firms only account for 0.3% of the total number of firms, they account for half of all firms' employment, and 61% of all firms' revenues. Quite tellingly, of the roughly 17,000 large firms in the U.S. economy, approximately 40% fall in the 2-digit NAIC category of "Management of Companies & Enterprises," classification number 55.

While the economy continues to evolve towards producing services and information, there has been a lag in the economic theory used to account for the behavior of firms in the new economy. The traditional production-function theory of the firm still holds sway as the main core of the theory of the firm and production, despite the fact that the U.S. economy has many large firms producing nothing more than information.

In this paper we model firms as networks of information processing agents. We build a computational model to explore the performance of these firm networks in uncertain and changing environments. Firms' organizational structures are determined by the number of agents, the number of levels in the firm's hierarchy, and the level of centralization of the firm. We explore how different organizational structures perform in different environments;

the environment is characterized by the level of complexity and uncertainty.

This paper extends the work of Barr and Hanaki (hereafter, BH)(2005) to include uncertainty and change in the environment in which the firm operates. We refer to a firm's *environment* as the set of technologies it uses to produce its output, and the set of opportunities available to it for earning profits. The environment includes both features internal and external to the firm. Internal features of the environment can include the technologies and expertise available for production, marketing, R&D, and distribution. In addition, the firm environment may include many items external to the firm such as the degree of competition and the nature of consumer demand. Because these internal and external factors depend upon the available technology and know-how of the firms in the industry, we refer to this environment as a *technology landscape*.

We assume that all of these environmental components interact in complex ways that change over time. For example, new technologies of production or distribution may become available to the firm, making new products more feasible and old products obsolete; competing firms may enter or exit the market and thereby create or eliminate product niches; consumer demand may shift thereby increasing or decreasing sales opportunities for the firm. The ability of a firm to react to these changes will be impacted by its internal organization. And, as Lawrence and Lorsch (1986) have illustrated, firm performance is determined by how well it is organized.

In general, we are interested in how the complexity and uncertainty of a firm's environment impacts its performance given the nature the organization's information processing structure. For instance, a complex environment (one that has many interactions between factors in the environment) may imply that the firm needs a holistic approach to decision making and information processing, where the firm considers a broad set of pertinent information when taking an action. Presumably, gathering a broad set of information is a time consuming process. In addition an environment that changes rapidly, and thereby creates large amounts of uncertainty, may imply that a firm needs to respond quickly in order to lower the costs from delaying a decision. Thus, there can be tension between the scope of firm decision making and the speed and cost with which the firm is able to react to change.

We incorporate these features by extending the model of BH, where the goal of a firm is to search a technology landscape in order to locate a set of projects that improve its profitability. Here, we explore the relationship between the complexity and uncertainty of a firm's environment and the structure of its information processing (IP) network. For simplicity, agents in the firm are arranged hierarchically, but information and decisions are made both in parallel and serially.[1]

As in BH, we model three features of the firm's organization: the number of vertical layers of the firm ("the depth"), the number of subordinates per manager ("the breadth"), and the level of the hierarchy at which decisions are made ("the authority level").

Also, as in BH, we model the environment in which a firm operates as a modified NK landscape. In BH, the complexity of the environment is measured by two parameters, K, the degree to which projects are connected (i.e., the "ruggedness" of the landscape) and $\beta \in [0, 1]$, a measure of the "orderliness" of the system, i.e., the degree to which project connections are global or local. When β equals zero, for example, all projects are interconnected locally (i.e., the system is nearly decomposable or nearly "modularizable" (Simon, 2002; Langlois, 2002); on the other hand, when β equals 1, the connections between projects are fully random, and thus the system cannot be decomposed.[2] Here in this paper, as will be explained below, instead of working with the parameter β, we introduce a new but related parameter that measures the frequency with which the location of interdependencies changes (thus, we introduce a measure of environmental instability).

[1] Future work will explore the evolution of network structure to investigate under what conditions hierarchical networks emerge. Other work has demonstrated the optimal nature of information processing hierarchies, for example Radner (1992).

[2] Note that the two measures of the complexity of the technology landscape are interesting when K is relatively small *vis a vis* N. As K grows relative to N, and the degree of interconnectedness grows, β becomes increasing irrelevant because almost all projects are connected.

Though this paper builds on the model of BH, it deviates in two key ways. As mentioned above, we introduce a measure of uncertainty in addition to complexity. Firms not only operate in environments that have differing degrees of technological and system-based complexity, but the environment also has differing degrees of uncertainty. As such, it is important to investigate firm performance in complex and uncertain environments to see which types of information processing networks perform best and are robust to environmental change.

Here we model uncertainty as changes to the connections between projects that occur over time. We refer to this type of uncertainty as *knowledge uncertainty* in that it encompasses changes to the structure of the environment in which the firm operates. In our model firms face two problems. First, the firm attempts to find a set of projects that maximizes payoffs in their environment. Second, the first problem implicitly requires the firm to uncover the underlying connections among the projects in its environment. Knowledge uncertainty incorporates changes in these connections that are unobserved by the firm and therefore must be learned.

As the environment becomes more fluid, organizations must put increased emphasis on finding good solutions to problems quickly. It does no good for a firm to find an optimal solution to a problem in a year if the problem it is trying to solve changes before it gets to the best solution. Thus there may be benefits to finding good, but non-optimal, solutions quickly as opposed to finding optimal (or near optimal) solutions but over a greater length of time.

Another feature of this paper that is different from BH is the explicit introduction of processing costs for organizations. In BH, the decision costs are embedded into the model as delays in the processing of information. In this paper we explore the relative costs to the firm for two explicit types of information processing activities: "decision making" and "information evaluation" to see how they affect net performance. The relative costs can reflect, for example, the efficiencies or comparative advantages of different firms–some may be more efficient at decision making, while others may be more efficient at evaluating information.

By explicit discussion of the costs we can focus on the advantages and disadvantages of firm size. In BH, in general, small firms sizes were optimal. Here we show that a larger firm size can increase the speed of information processing (and hence reduce the costs) since more agents are available to analyze information. We show under what circumstances a larger size firm performs better than a smaller one.

We investigate firm performance (excluding costs) from two different angles. First we demonstrate some time series runs for several organization types. to show how different organizations perform over time. For example, in the short run, decentralized organizations have a speed advantage, but centralized organizations tend to catch up and maintain a lead. The reason is that centralized organizations, as we discuss below, tend to search locally but have a wider scope in decision making. Also over time, we show that a similar time tradeoff exists between large and small firms: large firms move fast, but slow firms are "steady" and over time they eventually catch up.

In addition to time series results, we look at average performance over time to see what types of organizations perform best. As we show below, different firm sizes and levels of centralization are best depending on the nature of the environment.

In term of the costs to information, we show that decentralized firms have a cost advantage, since they evaluate data and make decisions faster. We demonstrate how the "net profit" of the firm is determined by the relative costs and benefits from its organizational structure.

The remainder of this paper proceeds as follows: In section 2 we discuss the two components of the firm as an information processing network and a technology landscape. We then present, in section 3, our model which includes explicit specifications of the environment in which our firms compete, and, in section 4, the organizational features of our firms. In section 5, we discuss how a firm's search algorithms and organization interact to affect firm performance through a series of computational experiments. In section 6, we discuss the costs involved with processing information for each type of organization. In section 7 we present results of the net profits of organizations, where we combine the costs and benefits to search and organizational structure. Finally, we offer some concluding remarks in section 8.

2 The Firm as Complex System

For the purposes of this paper, we can think of the firm as two types of complex systems: a production (technology) system and an information processing (hereafter, IP) system. The production system is the set of activities and resources that the firm has at its disposal to generate its outputs. The IP system is a network of agents who process information and make decisions. The goal of a firm is to locate a set of projects that is relatively efficient, i.e., produces the best fit with the economic environment.

2.1 The Firm as an Information Processing and Managerial Network

As Coase and Williamson discuss, firms emerge as a way to reduce both the transaction and principal-agent costs involved with producing goods and services (Coase, 1937; Williamson, 1985). Thus much of the market is brought "in house" as a way to centralize and coordinate the tasks of employees; the firm can be thought of as a type of micro command economy (Chandler, 1977). Therefore, a firm can be thought of as a network of agents; the agents must process and communicate information to each other so that the system can function efficiently, in light of the fact that firms operate in competitive environments.

The firm's internal characteristics evolve and adapt according to its global performance, which, in turn, is a function of how well the internal systems of the firm are put together. The organization is more analogous to the human body, where the functioning of the different parts must be coordinated to operate as a whole. Here we take an agent-based approach to investigate the behavior of the firm in the sense that the firm is modeled as a network of agents. Because organizations are highly complex systems, we simplify the nature of the agent activities to investigate how information processing activities affect performance. We assume agents to be boundedly rational and to have limited computational abilities. Workers and managers within the firm must specialize in a limited set of activities. There are no incentive issues between agents. Managers must also engage in the act of information processing in order to make decisions for the firm. We focus on how the complexity and uncertainty of a firm's environment can affect the performance of an organizational structure.

Work in this vein includes Radner (1992), who models the firm as a type of "adding machine" to explore what types of organizations have the lowest delay. In Miller's (2001) model, organizations are comprised of randomly generated networks of agents that perform associative operations. Over time organizations evolve by using genetic algorithm type rules to mutate or combine suborganizations in order to improve their speed of computation. Rivkin and Siggelkow (2003) model the organization as a CEO and two subordinate managers, with decision making control for a department or firm unit. Firms differ in the decision making role of the CEO, the quantity of information received from the subordinates, and agents' incentives regarding consideration of other agents' information. Finally, the work of Chang and Harrington (2006) summarizes the state of research on agent-based models of organizations.

2.2 The Firm as Technology Network

Because the production system itself is a complex system, we model it as a type of rugged landscape (Kauffman, 1993). Over the past ten years, there has been an increasing body of literature that views the firm's activities as some form of an NKlandscape; an idea that was developed in the biological sciences by Kauffman (1993). Kauffman's model depicted biological entities with genetic codes that evolved over time. Economists have adapted the model for understanding the evolution of economic phenomena since biological and economic systems have many things in common. Most importantly, economic systems like biological ones are inherently evolutionary in the sense that these systems must change over time as the environment changes; as such, internal changes (such as in genes in animals or routines in firms) in the system are required to maintain a high level of performance. In our model firms must search the landscape in order to find a set of activities that improve its performance. The search mechanism can be thought of as a type of evolution of its genetic code. Firms can search locally by making small incremental changes in its activities or globally by making drastic changes (such as corporate restructuring where firm's may sell off whole product lines and acquire new ones.)

In this sense, rugged landscape models can be used to model different aspects of the firm. In our case the rugged landscape is seen as a set of projects that may be undertaken by the firm. And the individual projects engaged affect other potential projects. In this sense there are *intranalities* within firms that reflect such interdependencies (Kauffman, et al., 2000). Specifically in our model, intranalities refer to the inter-connected effect for each payoff; that is decision i is affected by K other decisions. Thus when we change any of these K other decisions we also change the payoff to decision i.

3 The Environment

In this paper, as in BH, when we refer to the *environment* we mean a particular *technology network*, which represents a nearly decomposable system of decisions that interact in a non-linear way. The firm does not necessarily know how these activities interact, and therefore, the firm engages in search along the landscape to find, if not the best set of projects, then the ones that produce relatively high payoffs. Since the space of projects is large and agents have bounded rationality and computation abilities, they cannot simply select the global best set of projects. In addition, even if a firm could locate a global optimum in a certain period, environmental uncertainty can render this optimum obsolete; thus the firm must keep searching to maintain competitiveness.

As discussed above, we model the environment as a modified NK Landscape (Kauffman, 1993), which is a useful way to depict environmental conditions facing the firm. In this paper, our measure of complexity is how tightly coupled the environmental system is (K). In addition, we have a measure of environmental uncertainty (α), which is the probability that project couplings change (discussed in more detail below.) Without loss of generality, in the rest of the paper, when we refer to "the environment" we refer to a particular network structure given by $NK\alpha$ (which is a particular class of networks.)

3.1 Payoffs

The landscape is a mapping from the set $\mathbf{X} = \{0,1\}^N$ to R_+. That is to say, an element from the environment $\mathbf{x} \in \mathbf{X}$ is a vector of binary digits of length N, and each $\mathbf{x}$ is associated with a payoff $\pi(\mathbf{x}) \in R_+$. How this mapping occurs also depends upon the value of K, which is a parameter that specifies the degree of interaction among the elements of $\mathbf{x}$. Thus K is an important measure of environmental complexity. A large value of K means that elements of $\mathbf{x}$ are highly interdependent; a change of one value of $\mathbf{x}$ can cause dramatic changes in the payoffs associated with the new vector $\mathbf{x}'$. In the simplest case, when $K = 0$, there are no interdependencies (or intranalities), and, as a result, a change in $\mathbf{x}$ of one bit will result in a relatively smooth change in payoffs.

In the $K = 0$ case, the value of each bit is independent of the value of the other bits for a particular $\mathbf{x}$. Thus we create the "landscape" as follows: for each bit $x_i, i = 1, ..., N$ we assign a payoff $\pi_i(x_i)$ that is a randomly generated number from a uniform $0-1$ distribution. So for example if $x_i = 0$, it assigned a particular randomly generated number,$\pi_i(0)$, and if $x_i = 1$ it is assigned another randomly generated number, $\pi_i(1)$. In this case then we generate $2N$ different random payoff values. The value of each $\mathbf{x}$ is given by

$$\pi(\mathbf{x}) = \frac{1}{N}\sum_{i=1}^{N} \pi_i(x_i), \tag{1}$$

the average of the payoffs for each bit.

In the case of $K = 1$, the payoff of each bit x_i is also determined by the value of a bit x_j, $i \neq j$. Thus for each possible value of x_i and x_j we randomly generate a payoff value $\pi_i(x_i; x_j)$. (We discuss the relative locations of x_i and x_j below.) Since the payoff of x_i depends on the value of x_j, we have four possible payoffs associated with x_i : one when $x_i = 0$ and $x_j = 0$, one when $x_i = 1$ and $x_j = 0$, and so on. Thus to create a landscape, we generate $4N$ payoffs. The payoff of any particular vector is given by

$$\pi(\mathbf{x}) = \frac{1}{N}\sum_{i=1}^{N} \pi_i(x_i; x_j). \tag{2}$$

In the general case for any value of K, we have payoffs of x_i determined by the value of K other elements:

$$\pi(\mathbf{x}) = \frac{1}{N}\sum_{i=1}^{N} \pi_i\left(x_i; x_j^1, ..., x_j^k\right), i \neq j$$

To create a landscape we randomly generate $2^{k+1}N$ payoff values.

3.2 Uncertainty

We implement environmental uncertainty as changes in the relationships between projects over time. Let γ_i^t be the set of elements to which element i is connected in period t. We define a process where γ_i^{t+1} changes as a function of γ_i^t and an uncertainty parameter α. Suppose we have element i dependent on elements $j, k, ...$ With probability α, each connection is rewired to a new connection. i.e., with probability α the connection ij becomes connection

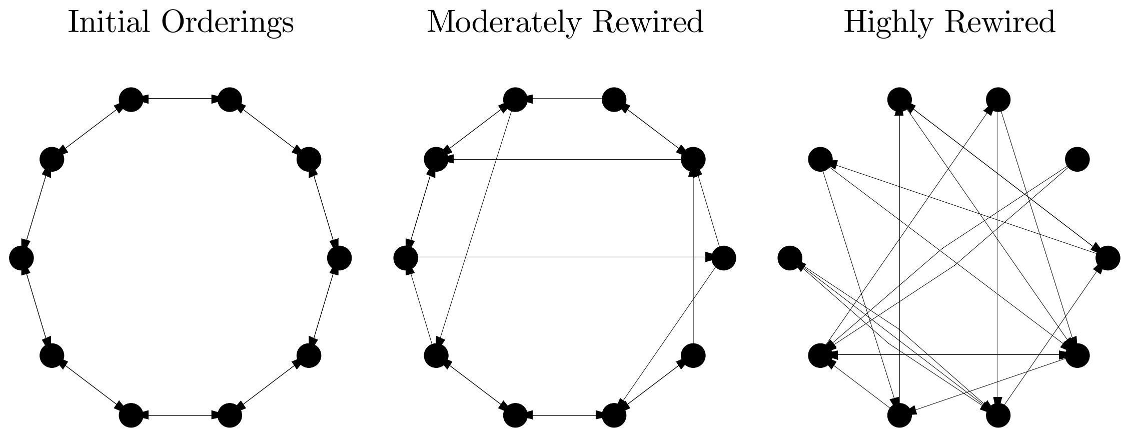

Figure 1: Example of evolution of connections over time, $k = 2, \alpha = 0.05$.

ij'; with probability α the connection ik, becomes connection ik', etc. Each of these potential re-wirings are independent. So, the probability that all of the K connections are re-wired is α^K. As these re-wirings occur over time the environment moves from an ordered to a more disordered system. Consider the example shown in Figure 1. In the first period each element is connected to its nearest neighbor. But as the re-wirings occur the connections between elements become less ordered and less localized. Thus the firms must learn to deal with change and with a more disordered and difficult system of interactions.

Note that when the re-wirings of connections occur we keep the payoffs the same but just change the interactions. In this formulation the firm tries to learn the new connections, and it leaves the set of potential payoffs unchanged.

4 The Organization

As discussed above, the firm is modeled as an information processing network. We assume, for simplicity, that the structure of the network is inherently hierarchical (and leave for future work the endogenous, dynamic formation of network structure). Here, our aim is to ask: how do different network structures perform given the environmental characteristics?

In this paper, we model the organization as a type of directed graph. There are two

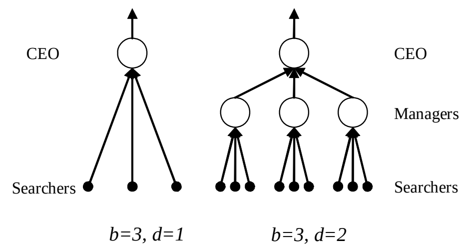

Figure 2: Two types of organizational structures

types of agents, *searchers* and *managers*. Searchers are associated with a particular location on the NK landscape. They search and report the location to the managers, who then transmit information up the network to the final or terminal node (CEO). An important parameter for the organization, apart from those that determine its graphical structure, is the locus of decision making authority, which will be discussed below following a description of organizational structure and the role of searchers.[3]

The Organizational Structure Information flows from the field agents to the managers and to a final node (CEO). We generate organizations based on two parameters, $b \in \{2, 3, \ldots\}$ and $d \in \{1, 2\}$. The value of b, the "branching ratio," (or breadth) is the number of subordinates per node. "The depth," d, is the number of vertical layers. The size of an organization, including CEO, with branching ratio b and depth d is $\sum_{j=0}^{d} b^j$ where b^d of them are searchers.[4] Two examples of organizations are given in Figure 2.

The organization divides up the landscape as follows. First each of the b managers are given "control" over the greatest integer number of bits less than N/b (i.e., each managers is assigned $\lfloor N/b \rfloor$ bits). If $b \lfloor N/b \rfloor \neq N$ then each manager is assigned additional bits in turn, until the landscape is fully partitioned. Given that each manager has b subordinates, the set of bits assigned to a manager is further divided up into smaller sets, and assigned to his

[3]Note that the description of the firm and it's activities close parallels that of BH (2005).

[4]This way of generating organizations is, of course, rather limited in the sense that it does not generate all the possible structures for a given organizational size. But this we leave for future research.

subordinates in a similar manner as described for the managers. Therefore, the landscape of size N is fully partitioned among the searchers.

The Searchers For example, suppose there are S searchers. For a given landscape of size N, we partition the searchers (also referred to as "field agents") so that each searcher evaluates a particular subvector of $\mathbf{x}$ such that at a given time agent s evaluates ω_s bits. Mathematically, we refer to the bits under consideration of agent s as χ_s, such that $\cup_{s=1}^{S} \chi_s = \mathbf{x}$ and $\cap_{s=1}^{S} \chi_s = \emptyset$; we call a particular χ_s at a given time the *proposal* of agent s.

Each period each field agent flips a randomly chosen bit under her control. If the searcher has decision making authority, she calculates the payoff of the new proposal, which is the average payoff of the payoffs of the bits under her control:

$$\pi_s(\chi_s) = \frac{1}{\omega_s} \sum_{i \in \chi_s} \pi_s(x_i). \tag{3}$$

Notice that the searcher only observes a payoff for each bit, but does not have any knowledge about how the bits are interconnected. If the change improves her payoff (calculated by using equation 3) she proposes that change; if not, she proposes the status quo. If she does not have decision making authority she simply passes the proposal up to her manager.

Authority and Decision Making Another important parameter is $a \in \{0, ..., d\}$, the "authority level," which designates the layer at which the final decisions are made. The level of authority determines how centralized the decision making is. If $a = 0$, the authority is given to the highest level, i.e., to the CEO. Thus the organization is *fully centralized*; in the case of $a = d$, decisions are made by the searchers themselves, and the organization is *fully decentralized*; for $0 < a < d$ decisions are made by the middle managers if they exist.

In the fully decentralized case, each period a searcher randomly flips one of the ω_s bits under her control and evaluates the payoff according to equation (3). Then she compares it to the previous proposal. If the new payoff is greater than the current one, she selects the new one. If not, she keeps the current proposal. Next she passes the proposal (i.e., the subvector) up the hierarchy to the next level. The manager above the searcher then takes the proposal fed to him and "joins" it with proposals from the other searchers under him. Notice that when authority resides in the lowest level, the agents above the decision makers only act as information processors – joining the proposals and passing them up the hierarchy to the CEO, who then calculates the final payoff for the entire proposal.

If the authority resides with the middle managers then the decision making works as follows. Each searcher randomly flips one of the ω_s bits under his control and evaluates the payoff as above. The field agent passes the new proposal, if any, up to the manager above her. Each manager "takes in" the proposal from his subordinates and compares the proposals one by one holding everything else constant. For example, if a manager has 3 subordinates, and all the subordinates have passed up a new proposal, the manager first evaluates the new proposal of subordinate 1, while keeping constant the old subvectors of subordinate 2 and 3 . Next, he evaluates the new proposal of subordinate 2, keeping constant the old subvectors assigned to subordinate 1 and 3, and so on. Thus if each manager has b subordinates, then each manager only evaluates at most b new proposals, and passes the best one up in the hierarchy to the CEO, who joins the b proposals from the managers and calculates the final payoff. Note that each manager evaluates at most b new proposals, and then takes the maximum value of the $b+1$ proposals (the b new ones and the old/current one). The assumption about managers evaluating one proposal at time and choosing the best one, instead of considering all the possible combination of received proposals, is made based on the assumption that managers have limited time and processing capability, and thus a natural simplification is that they only evaluate one possible change at a time, and that the number of these one-bit changes are limited.[5]

[5]In future work, we plan to explore the effect have having managers with greater capabilities that can explore many or all possible project combinations.

A similar procedure is applied when the authority rests with the CEO. Now the middle managers, if they exist, act as screeners, comparing the b proposals offered to them by their subordinates. Then they pass up the best one to the CEO. The CEO evaluates the b new proposals from his subordinates and selects the best one.

Notice that the higher is the authority level, the more information is used to make a decision in the sense that agents see a larger proportion of the total number of bits in each proposal. This becomes particularly relevant as we increase K, for example, since, *cet. par.*, an increased value of K means more interdependencies among the environmental bits, and greater centralization confers a greater ability to evaluate all the proposals, but changing b, d, and a affect the distance of search, the quantity of information processed and the scope of project evaluations. After a brief example of how the search algorithm works, we discuss these different facets of firm search below.

4.1 Organizational Structure and Search Distance

Here we investigate the relationship between the Hamming distance and organizational structure. Take two project vectors, $\mathbf{x}, \mathbf{x}' \in \{0,1\}^N$. Define the *Hamming distance* (HD) as

$$HD(\mathbf{x}, \mathbf{x}') = \sum_{i=1}^{N} |x_i - x_i'|,$$

which is just the number of bits that differ in each vector. The HD can be used as a measure of how far the organization is searching each period, if $\mathbf{x}$ is the firm's current set of projects and $\mathbf{x}'$ is a new proposal under consideration.

In the $d = 1$ case, we have two types of organizations with regard to authority level: fully centralized, $a = 0$, or fully decentralized, $a = 1$. Let's begin with the decentralized organization with b agents. Each period, each of the b agents flips 1 bit, thus $HD = b$. In other words, the firm can change up to b bits each period. As we discuss below, however, greater b may have other effects as well.

For comparison, consider the centralized organization ($a = 0$). In the first stage, we have $HD = b$ since each of the agents flip one bit. But the organization filters out information (also discussed below) as the CEO then chooses only the best of these proposals to be implemented. After filtering and evaluation the organization has $HD = 1$. So the centralized organization has localized search in that it only changes one bit every period. The firm can only move a small distance from its current location each period.

Now consider organizations with more depth, $d = 2$. Again begin with the decentralized case, ($d = 2$, $a = 2$), with breadth of b. This gives b^2 searchers. Thus we have $HD = b^2$, so these types of organizations can search at large distances from their current location. If $a = 1$, then each manager evaluates b proposals with $HD = 1$, but there are b managers each making independent decisions. Thus we have $HD = b$. In this case the organization has an intermediate search distance. Finally, with full centralization, $a = 0$, the organization filters down until the CEO selects one of b proposals that each have $HD = 1$. Again, full centralization has only localized search.

4.2 Granularity and Filtering

As decision making becomes more centralized, for example, the organization *filters out* information that otherwise might be useful. That is, since managers evaluate a limited number of proposals from their subordinates, they are in effect reducing the number of projects evaluated (in addition to the distance of search). Thus filtering means a reduction in the *number of projects* evaluated. Filtering increases with greater levels of centralization. But greater centralization also allows for the decision maker to have more information available to her for the decision. We refer to this as the *scope* of the decision maker. Thus centralization introduces a tradeoff: a reduction in search distance, but an increase in the ability to account for the intranalities.

Another issue in regard to organizational structure and performance relates to what we call the *granulatriy effect:* as we increase b we increase the number of searchers and thus agents have fewer bits under their control. Their decisions are more likely to work at cross purposes with other agents, since they do not take into consideration the decisions that other agents will make. That is, the fewer projects that agents evaluate (larger b means fewer projects evaluated per agent), the more likely their decisions will have negative consequences for the firm as a whole (though it may be good for the individual agent) since agents will be working at cross purposes.

To illustrate, let's take two extreme cases. First, let's say there is only one searcher, and thus she evaluates the entire project proposal and thus is able to measure the effect of proposal changes to the entire firm. On the other hand, let's say $b = N$, in this case, each searcher's choice is made irrespective of the effect of any other project choice on the firm. Thus, unless $K = 0$, the agent will not evaluate the entire effect of her proposal since it affects other parts of the firm that are beyond her knowledge; there will be some amount of granularity effect. Clearly the granularity effect increases as we increase either K or b or both. The most dramatic granularity effect occurs for large b and K values. However, larger values of b introduces a tradeoff: the greater the value of b means that the firm searches across a larger distance from its current location (as measured by the Hamming distance) but a large b also implies a large granularity effect and limited scope for the firm.

5 Simulation Experiments for Organizational Structure and Search

We now use the model described in the previous sections to illustrate the effects of search distance and granularity in various environments. We first illustrate the search distance and scope effects of authority level as we vary the level of complexity (k) and uncertainty

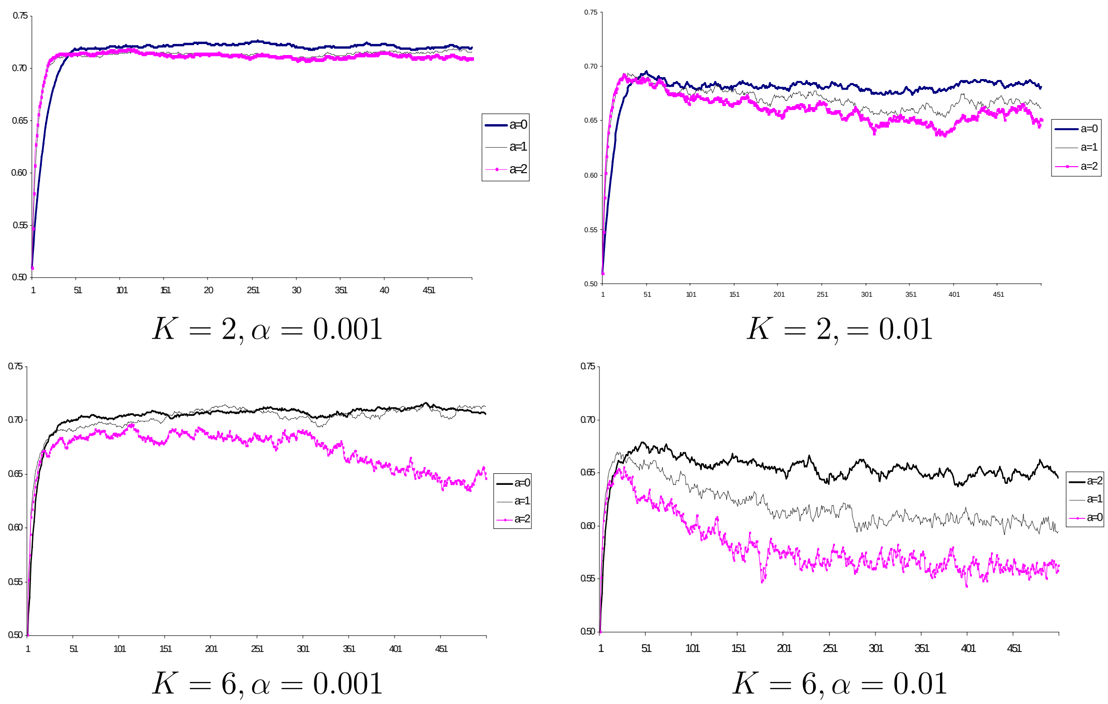

Figure 3: Performance over time of $d = 2, b = 4$ organizations in different environments.

(α) in the environment. Following which we examine the effects of the branching ratio and granularity in these same environments.

We begin with a $d = 2$ and $b = 4$ firm structure that can have authority levels of $a = 0$ (centralization, CEO level decision making), $a = 1$ (intermediate, manager level decision making), and $a = 2$ (decentralization, searcher level decision making.) We consider two levels of complexity, $k = 2$ (simple environment) and $k = 6$ (complex environment), and we consider two levels of uncertainty, $\alpha = 0.001$ (low uncertainty or "stable") and $\alpha = 0.01$ (high uncertainty or "unstable").[6]

We display the firm performance of each authority level for these four possible combinations of complexity and uncertainty in Figure 3.[7] Beginning with the simple and stable environment ($k = 2$, $\alpha = 0.001$) we see that in period 500 all firms achieve roughly the

[6]We find that for levels of uncertainty slightly larger than 0.01 that all firm structures perform equally poorly due to there being too much noise in the environment.

[7]Each plot is the average performance of twenty runs, with each run having the same initial starting vector and payoff.

same level of payoffs. The more decentralized firms achieve slightly worse payoffs but the difference is small. However the most interesting thing in this panel is the initial speed with which the firms begin their search. The decentralized firms more quickly achieve their level of payoffs. They achieve a payoff of 0.70 almost immediately and shortly thereafter they level out to a flat payoff level. The moderately centralized and centralized firms behave in a very similar manner except that they take longer to reach their payoff plateau. The moderately centralized firms reach a plateau around period 30 and the centralized firms reach a plateau around period 40. This observation captures the speed effects that go along with the search distance of the firms. The decentralized firm is able to move more quickly across the environment to find a good solution because of its greater per period search distance. The other firms move more slowly. But in the end, all of the firms reach nearly the same level of payoffs on average because the problem that they are trying to solve is relatively simple.

In the top right panel of Figure 3 we increase the uncertainty of the environment $\alpha = 0.01$ but keep the complexity at a relatively simple level $k = 2$. Again we view that the centralized firm is the slowest moving organization type. But this time the centralized firm consistently out-performs the other firm types by the end of the simulations. To understand this result recall that the level of uncertainty measures the probability that each bit interaction is rewired in each period. So as the system progresses we are moving from an ordered system to a more disordered system (the interaction bits are less likely to be near each other in period 500 than in period 1.) Thus, as disorder increases, the scope of the firm becomes increasingly important. The centralized firm is able to view these long range interactions but the more decentralized firms are not.

In the bottom left panel of Figure 3 we return to a low uncertainty environment ($\alpha = 0.001$) but increase the level of complexity ($k = 6$). Here we observe that the decentralized firm does very poorly relative to the other firm types. The decentralized firm is unable to account for the level of complexity that is contained in the environment. In addition, the performance of the decentralized firm significantly decreases over time as the system becomes more disordered. The decentralized firm is unable to deal with this level of complexity. However, the centralized firms and the moderately centralized firms both perform at levels comparable to the simple environment. The increased level of complexity does not appear to effect the ending payoffs achieved by these firm types. But notice that it takes longer for the $a = 0$ and $a = 1$ firms to achieve these payoffs relative to the top left panel. Because the environment is more complex these firms are solving a more complex problem and as one would expect, it takes them a longer time to do so. In the end the $a = 0$ and $a = 1$ firms are able to solve the problem, but the decentralized firm is not (relative to the centralized and moderately centralized firms.)

In the final (bottom right) panel of Figure 3, we consider a highly complex ($k = 6$) and highly uncertain ($\alpha = 0.01$) environment. Recall that there are two effects of the authority level that come into play: First as we centralize the firm we increase the scope of the firms ability to deal with the intranalities. Thus as we saw in the previous paragraph, decentralized firms do not deal well with highly complex environments. Second, more decentralized firms respond more quickly because of their larger per period search distance. (This is best shown in the first panel of Figure 3.) As uncertainty increases we need firms to respond more quickly to their environment. Else, the payoff they are working toward may be gone by the time they get there. Thus one can imagine that decentralized firms may perform well as uncertainty increases. But, if the environment is too complex, the decentralized firm may not be able to find adequate solutions even if they are able to move quickly, as we saw in the third panel of Figure 3. We observe in the highly complex and highly uncertain environment of our model that the scope effects win out in this case. In this version of our model uncertainty also creates disorder over time and thus a centralized firm structure is most beneficial.[8]

[8]We also have another version of our model in progress where only the payoffs, and not the re-wirings, change. This model contains uncertainty because of the changing payoffs but it will hold constant the orderings of the interactions across time. We expect this model to more clearly demonstrate the speed effects of decentralization. However, the results are not yet complete.

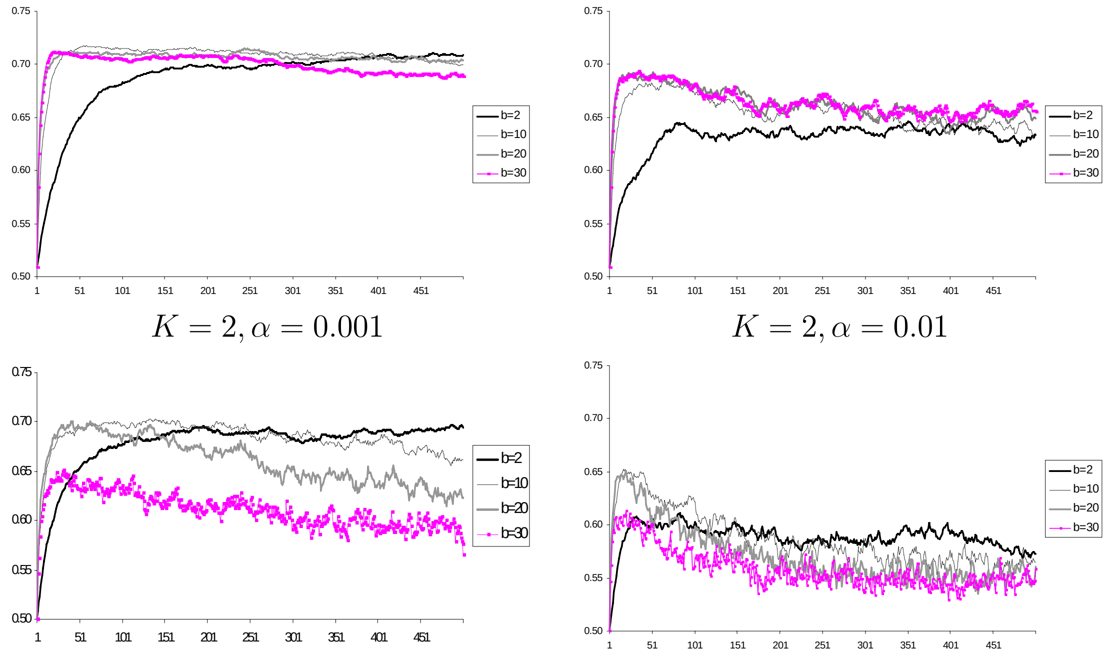

$K = 2, \alpha = 0.001$

$K = 2, \alpha = 0.01$

$K = 6, \alpha = 0.001$

$K = 6, \alpha = 0.01$

Figure 4: Performance over time of $d = 1, a = 1$ organizations in different environments.

5.1 $d = 1, a = 1$ Organizations

In Figure 4 we complete the same analysis but this time we vary the branching ratio, b. Here we hold the depth and authority level of the organization constant at $d = 1$ and $a = 1$. Recall that there will be two effects of an increased branching ratio. First a large b implies that each searcher controls and searches over fewer bits. Thus a large b implies a small scope. Second, a large b also implies that the firm will have a large search distance. Thus large b firms will be able to react more quickly.

In the simple and stable environment (the top left panel of Figure 4) we most easily see the search distance (speed) effects of the firms. A larger b implies that the firms should reach its solution more quickly. Each of the firms reaches their plateau level of payoffs and stays there with the largest b firms reaching their plateau more quickly but each plateau is also successively lower as we increase b. Essentially the slow firm with the large scope finds a better solution on average but it takes much longer to get there.

If we increase the uncertainty in the environment (top right panel of Figure 4) we see that speed becomes more important. As uncertainty increases the firms need to locate good solutions quickly (even if they are not optimal solutions) before the environment changes again. We see that the firms with a larger branching ratio are better able to cope with the uncertainty.

If we return to the low uncertainty case but increase the level of complexity ($\alpha = 0.001$, $k = 6$) we see that the large b firms do very poorly and their performance decreases as the amount of disorder in the system increases. Similar to the decentralized authority level, the large b firms cannot handle high levels of complexity. The $b = 2$ organization again reacts very slowly and it takes it a long time to find a good solution but with enough time, the scope of the $b = 2$ firm allows it to perform best in highly complex environments.

Finally, when we consider an environment with high complexity and high levels of uncertainty we see the slow speed effects of the small b firms. The overall performance of the small b firm drops to being just above the other firms. Essentially the slow speed disadvantage is magnified in the high uncertainty case to make it only marginally better than the high b firms.

Overall, this section has provided examples to demonstrate how authority level and firm size characteristics affect the search performance of the firm. In summary, centralization provides large scope for firm decision making but also slows the response of firms to changes in the environment. Large firms (those with large b) provide fast reactions but do so at the cost of limited scope.

6 The Costs of Processing Information

In the section above, we explored the performance of different networks and how they balance the tradeoff between speed and scope. Here we turn to the question of how network structure affects the costs of the firm. We focus on two particular costs that are related to the speed with which a firm can process information: the number of decisions made each period and the number of evaluations each period. In short, our focus is on the algorithmic "delay" involved with selecting a set of firm activities. The main costs to IP are the time costs involved with generating a decision (Radner, 1992). We therefore count the number of steps involved in processing information and apply parameters that convert the number of steps to a monetary time cost.

A *decision* occurs each time an agent considers a new proposal. Note that if a subordinate passes up the same proposal as the last period then the manager does not have to make a decision in the present period with respect to that proposal because it is already part of the status quo proposal.

An *evaluation* occurs when an agent evaluates a new proposal or combines proposals together. That is to say an evaluation is the process of calculating a payoff for a particular proposal or subproposal; on the other hand, a decision involves comparing different proposals or subproposals and selecting the one with the maximum value. Thus an evaluation is the act of adding all of the individual payoffs together to prepare for a decision on a given proposal.

The Decision Costs Define $\overline{b}_t$ as the number of decisions that need to be made by a manager each period and $\overline{b}_t'$ as the number of decisions made by the CEO if she has the authority. We then calculate per period *decision cost* as follows.

$$DC_t = c_1 \left\{ \begin{array}{c} 0 \text{ if } (d-a)=0 \\ \overline{b}_t \text{ if } (d-a)=1 \\ \overline{b}_t + \overline{b}_t' \text{ if } (d-a)=2 \end{array} \right\},$$

where c_1 is a parameter which converts the number of decisions made to a monetary value.

In the decentralized cases, we normalize the number of decisions to zero. The reason is that searchers make the same number of decisions, cet. par, regardless of the authority level and firm size because of the parallel nature of decision making at the searcher level. Thus we can compare the costs of medium centralization and full centralization relative to the decentralized cases.

Note that $\overline{b}_t$ is related to b in the sense that the more agents there are the more decisions that will have to be made. But in general $\overline{b}_t \neq b$. The actual number of decisions may be less than the maximum (which is given by b) in any given period if searchers are not finding new proposals to evaluate. As well, the actual number of decisions will be a function of complexity, uncertainty and time. In short, decisions costs are a simple linear function of the total number of decisions made by an organization. Notice also that, cet. par., increasing centralization increases the costs with respect to decision making since more levels are involved in the decision making process.

The Evaluation Costs The following equations describe the *evaluation costs* for our model. For ease of exposition, we explicitly state the cost function separately for $d = 1$ and $d = 2$ organizations.

$$AC_t = c_2 \left\{ \begin{array}{c} N/b^d + \delta N \text{ , if } a = 2 \\ N/b^d + \overline{b}_t N/b + \delta N \text{, if } a = 1 \\ N/b^d + \overline{b}_t N/b + \overline{b}'_t N \text{, if } a = 0 \end{array} \right\} \text{ if } d = 2 \tag{4}$$

and

$$AC_t = c_2 \left\{ \begin{array}{c} N/b^d + \delta N, \text{ if } a = 1 \\ N/b^d + \overline{b}_t N, \text{ if } a = 0 \end{array} \right\} \text{ if } d = 1, \tag{5}$$

where c_2 is a parameter. δ is an indicator function that equals 1 if a new decision is made, and 0 otherwise.[9]

Notice a few things about the cost functions. Every firm has to perform at least N/b^d

[9] Note that we ignore wage costs on the assumption that for the size networks we work with node costs are fixed.

evaluations because this is the number of evaluations made per agent at the searcher level.[10]

The greater the size of the organization, as given by b, the lower the evaluation costs, cet. par. This is simply due to the fact that each agent has a responsibility for a smaller part of the landscape. Thus, all else equal larger organizations should have lower costs to evaluating information. However, if $a < d$, then there is some degree of centralization and this increases the number of evaluations (and decisions) being made. As a result, increased centralization will offset any reduction in costs due to increased b.

Lastly, in general, as a firm searches over time, the average number of decisions that it makes will tend to stabilize. As a result, in the simulations below, the average value of δ across runs becomes an empirical probability that a firm will choose a new proposal at each time period. For "easy" environments, δ will converge to zero; increasing complexity and/or uncertainty will mean that δ will not converge to zero, in general.

In summary we can conclude the following from the cost functions. The determination of costs for information processing is a function of a, b, and d. In the long run, decentralization will always have a lower cost than centralization. The difference in the long run costs between centralization and decentralization will always be a function of the nature of the environment. In complex environments, where there are many more decisions to be made each period, on average, the centralized firm will have a much greater cost than in an easier environment.

The tension between performance and cost from centralization introduces an important tradeoff: on one hand, as we saw above in section 5, the centralized organization tends to have better performance as we increase the complexity and instability. However, we see that with the cost functions, decentralization will have a long run cost advantage. Which type of organization is best suited for each environment, will, of course, depend on the cost parameters, and the firm's structure. We investigate this below by looking at average "net

[10] Note that we don't have managers and/or CEO's simply add the payoffs of the searchers' choices, but rather they evaluate the project proposals passed up to them. The reason for this is that when $K > 0$, the payoffs from joining subproposals will not equal the sum of the payoffs of the subproposals.

profit" versus different organizational structures in different environments.

7 Costs and Performance

The aim of this section is to compare different organizations to see which ones have the highest average net profits. In addition, we look at all organization types on the same graph to get a better sense of the relative performance of these organizations. To compare organizations we look at the total number of agents in each organization. For the $d = 1$ case, the number of agents equals $b + 1$; while for the $d = 2$ case, the number of agents equals $b^2 + b + 1$.[11]

7.1 No Cost

In the no-cost case we can simply look at average performance versus size for the five different organization types ($d = 1, a = 1,\ a = 0; d = 2, a = 0, a = 1, a = 2$). Figure 5 shows the results of this.

In the easy environment (upper left), all the firms' performances are relatively close, though the medium-size centralized firms seem to have a slight edge. The reason is that in a simple environment, as shown in Figure 3, centralized firms have the best long run performance due to their having the right balance between search distance and lack of granularity. Size does not appear to add an advantage since in a relatively easy environment, all firms achieve a good position relatively quickly.

The graphs on the right (unstable environments) show that the centralized and large organizations have a clear edge. This is because increasing the size increases the number of proposals that CEOs view and the CEO can take advantage of centralization, which is necessary due to the difficulty of the environment.

[11] Average performance is determined by taking $\frac{1}{500}\sum_{i=1}^{500} \pi_t - (AC_t + DC_t)$ for each organization. Note that π_t, AC_t, DC_t are average performance and costs over 20 runs (to smooth out per-run fluctuations).

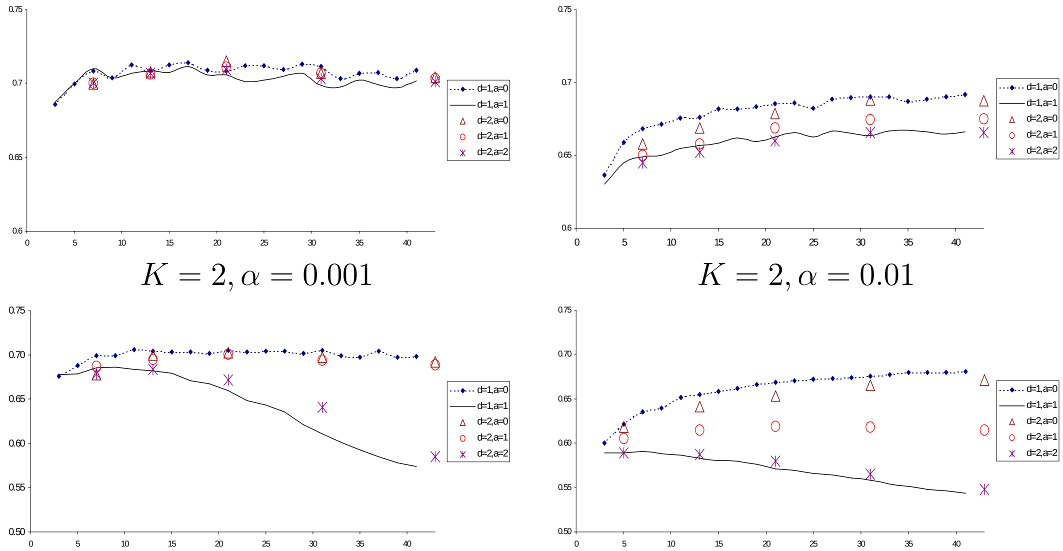

$K = 2, \alpha = 0.001$ $K = 2, \alpha = 0.01$

$K = 6, \alpha = 0.001$ $K = 6, \alpha = 0.01$

Figure 5: Average performance versus size for different organizations with no IP costs.

7.2 Low Costs

By imposing costs (i.e., c_1, $c_2 > 0$), we penalize firms with large b in the centralized case, though large b can be an advantage in the decentralized cases. In figure 6, we show the results when we impose a "small" cost for processing information. Based on this we can see that decentralized organizations improve their relative status. In the easy case (upper left) the decentralized firms seem to have a relative edge *vis a vis* the no-cost case. In the no-cost cases, large firms seem to be favored. However, when we impose a cost, this result is not always true. In the $s - u$ case (upper right), for example, large decentralized and medium decentralized organizations perform best, while in the $c - s$ case (lower left), there are several types of organizations that perform equally as well: small decentralized, medium centralized and medium size with medium decentralization. In the $c - u$ case (lower right), large centralized firms are still the most profitable. This is because the small cost imposed does not exceed the benefits to centralization, which are relatively large with a difficult environment.

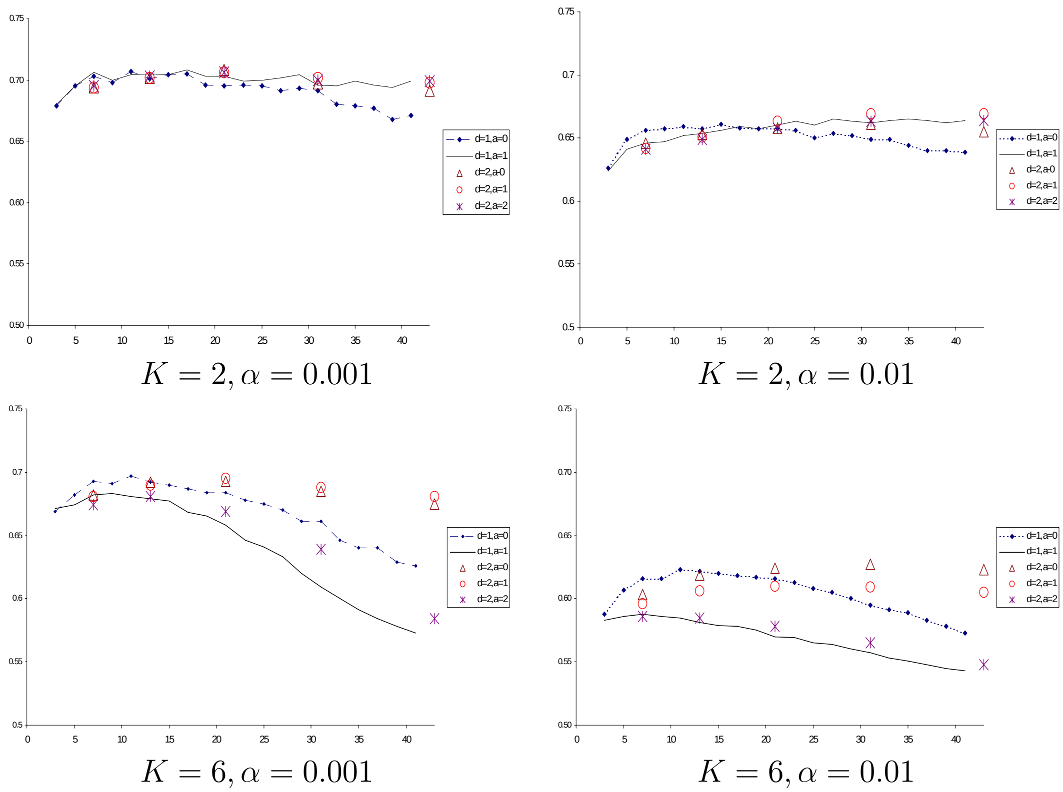

Figure 6: Average performance versus size for different organizations with small IP costs ($c_1 = 0.001, c_2 = 0.0001$).

7.3 High Costs

In Figure 7, we can see what types of organizations have the highest net profit when we apply a higher cost to information processing. In these cases we see an even further relative advantage given to decentralized organizations. In the easy case (upper left) we see

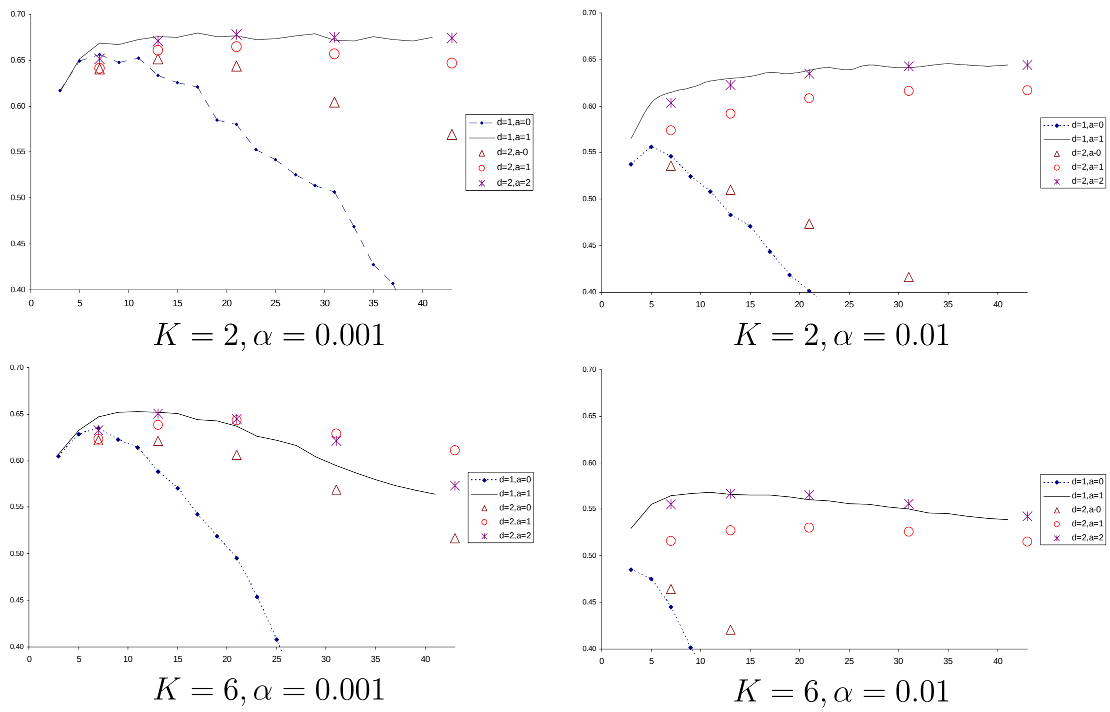

Figure 7: Average performance versus size for different organizations with large IP costs ($c_1 = 0.01, c_2 = 0.001$).

that larger decentralized organizations perform the best. In effect, the small advantage to centralization becomes "too expensive." A similar effect holds for the simple-uncertain case (upper right), where having more agents speeds up IP and decision making. However, in the most difficult environment, the small decentralized firms are best. This occurs because the smaller the firm the more "holistic" is its search (i.e., lower granularity). Thus complexity in general requires a relatively small firm size for good performance.

8 Conclusion

In conclusion, we would like to discuss our results within a broader perspective and provide a sense of possible future work. First, our results show that there can be a tension between the speed with which an organization is able to adapt and the ability to have a holistic perspective. In this paper these ideas play out in terms of the ability of decentralized organizations to search at large distances from its current location on the landscape, but at a certain search-related opportunity cost: in order for the firm to search a great distance the firm had to forgo having a large scope in its decision making. If the firm grants too much authority to its ground-level workers it cannot also ensure coordination of those workers with each other.

This tension plays out in different ways depending on the environment in which the organization operates. We saw that in simple and stable environments it did not matter to a large degree how the firm was organized, all firms performed at similar levels. But as the environment becomes more complex the firm needs to centralize its decision making in order to account for the interconnectedness of its projects. In addition, we investigated the role that firm size played in performance. Increasing uncertainty favors larger, centralized organization, while increasing complexity appears to be "size neutral." As the environment becomes increasingly rewired having more agents confers a speed advantage upon the firm since more agents means more information available for the manager to make decisions.

We also saw a tradeoff on the cost side. We saw that the cost of centralization had largely to do with the number of decisions and evaluations that the organization made over time. Of course the number of decisions and evaluations is partly a function of how well the organization is performing. If the organization has reached a satisfactory level of performance where new proposals are unlikely to beat the *status quo* then centralization tends to have very low costs.

On the other hand, centralization can be very expensive in terms of decision making and evaluation if the firm is finding new proposals frequently. There are two points to be made here: First, if the organization's level of performance is poor there will be many new proposals that can improve performance. Second, evaluating all of these proposals and deciding between them is costly in a centralized environment. Thus the results from the present paper can lead one to believe that there may be relationships between firm performance and the optimal level of centralization of the firm. Essentially, if the firm is performing poorly, perhaps it is best to allow the lower levels of the firm to make the decisions in order to attempt to find something that allows better performance. On the other hand, if the organization is already performing at a high level, the organization will have minimal costs from centralization because the lower level agents will not be finding many new performance improving solutions. And the ones that are found may be likely to conflict with the current state of the firm (if the environment is complex).

If we think about the above comments in a dynamic context they lead to some interesting ideas for future research involving the dynamic adaptation of firm structure. Presumably when an organization (or an industry is young) the organizations will not be performing at a high level. Thus a decentralized approach may benefit the firm. But over time, the firm should improve and move into a realm where it may want to slowly move toward a more centralized structure.

Further, if one considers the search process of the organization, if a firm is searching for its initial performance enhancing projects it most likely prefers a large search distance. The organization wants to try many new ideas to see which ones are potentially performance enhancing. But once the firm finds a sufficient level of performance it probably wants to search locally for other opportunities near the current location. One can think of this in terms of the research on exploration versus exploitation (March, 2001). The decentralized structure may allow an organization to explore a wide search area initially but then eventually settle down into a more refined and intricate local search that begins to take account of the various intranalities in the local environment. One can think of this second stage in somewhat the same context as exploiting the local area near the current solution. Further, these thoughts also resemble the process of simulated annealing (Kirkpatrick et al., 1983) in that the organization would begin with a large search distance associated with decentralization, but “cool” that search distance over time by moving toward centralization. Future work will explore the dynamic adaptation of firms structures in a context similar to that employed in this paper.

References

[1] Barr, J. and Hanaki, N. (2005). "Firm Structure, Search and Environmental Complexity." Rutgers University Newark working paper #2005-007.

[2] Chandler, A. D. (1977). *The Visible Hand: The Managerial Revolution in American Business.* Cambridge: Belknap Press of Harvard University Press.

[3] Chang, M.-H., and Harrington, J.E. (2006). "Agent-Based Models of Organizations." *Handbook of Computational Economics,* Vol 2. Eds K.L. Judd and L. Tesfatsion.

[4] Coase, R. H. (1937). "The Nature of the Firm." *Economica N. S.*, 4, 386-405.

[5] Greenspan, A. (1999). "Remarks by Chairman Alan Greenspan Before the Dallas Ambassadors Forum." Dallas, Texas. April, 16. http://www.federalreserve.gov/boarddocs/speeches/1999/19990416.htm

[6] Kauffman, S. (1993). *The Origins of Order: Self-Organization and Selection in Evolution.* Oxford University Press.

[7] Kauffman, S. Lobo, J. and Macready, W. G. (2000). "Optimal Search on a Technology Landscape." *Journal of Economic Behavior and Organization*, 43, 141-166.

[8] Kirkpatrick, S., Gelatt, Jr., C.D., and Vecchi, M.P. (1983) "Optimization by Simulated Annealing," *Science*, 220 (13 May), 671-680.

[9] Langlois, R. N. (2002). "Modularity in Technology and Organization." *Journal of Economic Behavior and Organization*, 49, 19-37.

[10] Lawrence, P. and Lorsch, J. (1986). *Organization and Environment: Managing Differentiation and Integration,* revised edition. Harvard Business School Press: Boston.

[11] March, J. G. (1991). "Exploration and Exploitation in Organizational Learning." *Organization Science*, 2(1), 71-87.

[12] Miller, J. H. (2001). "Evolving Information Processing Organizations." *Dynamics of Organizations: Computational Modeling and Organization Theories.* Lomi, Alessandro Larsen, Erik R., eds. MIT Press: Cambridge.

[13] Radner, R. (1992). "Hierarchy: The Economics of Managing." *Journal of Economic Literature*, 30, 1382-1415.

[14] Rivkin, J. W. and Siggelkow, N. (2003). "Balancing Search and Stability: Interdependencies Among Elements Of Organizational Design." *Management Science*, 49(3), 290-311.

[15] Simon, H. A. (2002). "Near Decomposability and the Speed of Evolution." *Industrial and Corporate Change*, 11(3), 587-599.

[16] Williamson, O. E. (1985). *The Economic Institutions of Capitalism.* New York: Free Press.