Complexity, uncertainty, and organizational congruency

Nobuyuki Hanaki, Hideo Owan

GREQAM
Groupement de Recherche en Economie
Quantitative d'Aix-Marseille - UMR-CNRS 6579
Ecole des Hautes études en Sciences Sociales
Universités d'Aix-Marseille II et III

Document de Travail
n°2010-08

Complexity, Uncertainty, and Organizational Congruency

Nobuyuki Hanaki

Hideo Owan

March 2010

Complexity, Uncertainty, and Organizational Congruency

Nobuyuki Hanaki
(GREQAN-Université de la Méditerranée)

Hideo Owan
(Institute of Social Science, University of Tokyo)

Many scholars in the fields of organization theory and management strategy have argued that there is a tension between the two types of organizational learning activities, exploration and exploitation. They appear to be substitutes: the greater the skill at one, the harder it is to do the other well. It is often argued that the two activities compete for scarce resources when firms need different capabilities and management policies to promote one over the other. We present another explanation that attributes the phenomenon to the dynamic interactions among the activities, search, knowledge sharing, evaluation, and alignment within organizations relying on the NK Landscape framework (Kauffman 1993). Our results show that successful organizations tend to bifurcate into two types: those that always promote individual initiatives and build organizational strengths on individual learning and those good at aligning the individual knowledge base and exploiting shared knowledge. Straddling between the two types often fails. The intuition is that an equal mixture of individual search and organizational alignment slows down individual learning compared to the first organization type while making it difficult to update institutionalized knowledge because individuals' knowledge base is not so sufficiently aligned as in the second type. In such "straddling" organizations, once individuals get stuck with locally-best solutions in an uncoordinated manner, they cannot agree on how to improve the organizational knowledge. Straddling is especially inefficient when the operation is sufficiently complex (in other words, the interdependency is high) or when the business environment is sufficiently uncertain.

Keywords: Congruency, Complexity, Uncertainty, NK Landscape

Complexity, Uncertainty, and Organizational Congruency

Nobuyuki Hanaki* Hideo Owan†

March 15, 2010

Abstract

Many scholars in the fields of organization theory and management strategy have argued that there is a tension between the two types of organizational learning activities, exploration and exploitation. They appear to be substitutes: the greater the skill at one, the harder it is to do the other well. It is often argued that the two activities compete for scarce resources when firms need different capabilities and management policies to promote one over the other. We present another explanation that attributes the phenomenon to the dynamic interactions among the activities, search, knowledge sharing, evaluation, and alignment within organizations relying on the NK Landscape framework (Kauffman 1993). Our results show that successful organizations tend to bifurcate into two types: those that always promote individual initiatives and build organizational strengths on individual learning and those good at aligning the individual knowledge base and exploiting shared knowledge. Straddling between the two types often fails. The intuition is that an equal mixture of individual search and organizational alignment slows down individual learning compared to the first organization type while making it difficult to update institutionalized knowledge because individuals' knowledge base is not so sufficiently aligned as in the second type. In such "straddling" organizations, once individuals get stuck with locally-best solutions in an uncoordinated manner, they cannot agree on how to improve the organizational knowledge. Straddling is especially inefficient when the operation is sufficiently complex (in other words, the interdependency is high) or when the business environment is sufficiently uncertain.

Keyword: Congruency, Complexity, Uncertainty, NK Landscape

*GREQAM, Université de la Méditerranée, and University of Tsukuba, nobuyuki.hanaki@univmed.fr and hanaki@dpipe.tsukuba.ac.jp

†Institute of Social Science, University of Tokyo, owan@iss.u-tokyo.ac.jp

1 Introduction

I think our growth will end at only two or three times the current scale once we undertake to increase routines. I think the biggest reason we don't see "mega ventures" emerging in Japan is that every firm starts organizing and routinizing their processes too early. When I worked with American firms such as Microsoft or Apple, I was always surprised to see disorder in many aspects of their operations. They wouldn't be called "firms" if they were in Japan. Their processes are not very routinized, but that is exactly why they can continue to grow. Their workplaces are full of chaos, especially compared to their Japanese counterparts. However, creative people who can make breakthroughs prefer working in such places. So it is with our company. We are paying the cost of disorganization and mistakes caused by the chaos...But, it is impossible to have both creative workers and routine workers in the optimal mix.

- Masayuki Makino, CEO, Works Applications, Inc.[1]

Tension between exploration and exploitation has been a central theme in the literature of organizational learning since March (1991). The issue is not a mere perception held only by academics, but a real problem faced by many business executives, including the above-quoted CEO of one of Japan's fast-growing business software companies. Nevertheless, too few theories have been offered to analyze the trade-offs that are creating this tension at the organizational level. In this article, we develop an agent-based model that formulates what we believe are four important steps of organizational learning: search, knowledge sharing, evaluation, and alignment, through which organizational knowledge is updated. The model is useful for examining the origin of the primary tension between exploration and exploitation.

We identify organizational congruency as a main driver of knowledge exploitation rather than its result. Organizational congruency is formulated as the degree of alignment in knowledge base imposed on individuals. Higher congruency facilitates the upgrading of organizational knowledge, which in turn accelerates exploitation in our model. Key to our theory is the idea that knowledge diversity makes exploration more effective even as it impedes the updating of organizational knowledge, thus creating

[1] " Works Applications, Inc. (A) " by Hideo Owan, Aoyama Business School case, November 2009.

an obstacle to exploitation. Although March (1991) has already argued that efforts to exploit existing knowledge eventually suppress exploration by reducing the diversity of knowledge, he does not illustrate the mechanism by which more exploration leads to less exploitation or the reason the diversity of knowledge can discourage efforts to exploit it. The mechanism we illustrate in this paper goes beyond the usual argument that exploration and exploitation compete for scarce resources (March, 1991; Roberts, 2004). Our theory suggests that the tension between exploration and exploitation comes not necessarily from resource constraint, but rather from the substitutability between initiatives and alignment.

Crossan, Lane, and White (1999) have expressed a view similar to ours: "This tension (between exploration and exploitation) is seen in the *feedforward* and *feedback* processes of learning across the individual, group, and organization levels." According to them, *feedforward* is the transference of learning from individuals and groups to the organizational levels where ideas are embedded in the form of systems, structures, strategies, and procedures. *Feedback* is the way in which this embedded or institutionalized learning affects individuals and groups. Note that the feedback process relates to exploitation where developing organizational knowledge assimilates learning and actions at the individual level. Although our view is similar to that of Crossan et al. (1999), the tension arises endogenously in our model, whereas it is assumed as one of the four key premises in their frameworks (Crossan et al. 1999: 523).

There are two important assumptions that drive the key results in our theory. First, we assume that updating organizational knowledge requires a certain level of consensus. In our model, organizational knowledge is successfully updated only when a majority of high performers agree with the proposal to change it. This means that the mere showing of new ideas to others is not enough for implementation at the organizational level. Many others must have a similar view of the world and see common meaning in the idea. This means that a certain level of organizational congruency is required to promote organizational learning.

Second, we assume that the organization's management influences individuals' time allocation between exploration (i.e. experimenting with new ideas) and exploitation (i.e. copying organizational knowledge) at the individual level. The famous story of 3M's mandatory rule that its technical people should be able to devote 15 percent of their time to any projects of their own choice rather than those they are officially assigned to is a classic example. Management uses its authority and techniques for control, such as process management, to affect the way people achieve their targets. Furthermore, individuals engaging in exploitation do not have the option of ignoring organizational knowledge. Each kind of monitoring and incentive mechanism encourages individuals to take in institutionalized knowledge as their own knowledge base.

One important implication of our model is non-concavity in the optimization problem. As argued, for example, by Levinthal and March (1993) and Ghemawat and Ricart i Costa (1993) among others, firms tend to end up with the "extremes" in the presence of such non-concavities. According to our theory, successful firms tend to bifurcate into two types: those that always promote individual initiatives and build organizational strengths on individual learning and those good at aligning the individual knowledge base and exploiting shared knowledge. Let us call the former *high-initiative* organizations and the latter *high-alignment* organizations.

Straddling between the two types often fails. This bifurcation arises when the operation is sufficiently complex (in other words, the interdependency is high enough) or the business environment is sufficiently uncertain. The intuition is that an equal mixture of individual search and knowledge alignment slows down learning through individual search compared to the high-initiative organization while making it difficult to update institutionalized knowledge because individuals' knowledge base is not sufficiently aligned. In such organizations, once members get stuck with locally best solutions at the individual level, they cannot agree on how to improve the organizational knowledge. We find that the resultant inefficiency is especially large when tasks are interdependent or when the environment continuously changes.

More formally, we develop a model of organizations that undertake tasks of varying degrees of complexity. The complexity of a task is formulated as interdependency among tasks in determining functional performance following the NK landscape model (Kauffman, 1993). We also introduce environmental uncertainties by allowing the performance function to be redefined randomly from time to time.

Each organization consists of several members who search independently for better practices and also learn from the organizational knowledge. The organizational knowledge evolves over time as proposals from the members of the organization to modify it are constantly evaluated by high-performing members. Specifically, in each period each member of an organization either conducts an individual search for a better configuration (with probability λ) or adopts the configuration from the organizational knowledge (with probability $1-\lambda$). After all the agents have either conducted individual searches or learned from the organizational knowledge, each agent proposes, with probability p, to modify the organizational knowledge. Each proposal will be evaluated by the members of organization whose performance at the time of evaluation is higher than the average one in the organization. If the majority of high-performing members agree on a proposed modification, the organizational knowledge is changed.[2]

The two parameters of the model, λ and p, represent the degree of organizational congruency and the frequency of knowledge sharing within the organization, respectively. The lower λ is, the higher is the degree of organizational congruency. In other words, when λ is low, it is more likely that members of the organization follow practices or procedures that embody the organizational knowledge, regardless of whether or not adopting them results in higher performance. When p is high, every agent attempts to change the organizational knowledge more frequently.

Through extensive simulations of the model, we find that (1) frequent knowledge sharing within an organization has positive influences on its performance and (2) there exist non-monotonic relationships between the degree of organizational congruency and the organizational performance, especially when tasks are either complex or the environment is uncertain. Namely, the performance of an organization is high when there is a substantial, but not overly strong, degree of congruency ($\lambda \approx 0.2$) or no

[2]In case of a tie, a proposal is approved with probability 1/2.

congruency at all ($\lambda \approx 1.0$).

It is rather intuitive that too much congruency is counter-productive, since agents do not search for better practices and since without such individual searches it is not possible to improve the organizational knowledge from which agents learn. However, more detailed analyses may be required in order to understand the non-monotonisity of the relationship between congruency and performance. Our analyses reveal that the performance of an organization with no congruency (i.e. $\lambda = 1.0$) is always better than that of any other organization in early periods. When the task is sufficiently complex, however, on average, the performance of an organization with a substantial degree of congruency and frequent knowledge sharing catches up with and exceeds that of the former organization in later periods. This reversal of performance comes from the difference in the rate at which the organizational knowledge improves. When a task is complex, each member of an organization may arrive at distinct but locally best practices and may fail to agree on how to change the organizational knowledge. As long as agents rarely follow the organizational knowledge, this is not a problem (just as in the case of an organization without congruency). However, an organization with a moderate degree of congruency suffers from such disagreements because they result in its agents adopting practices that are not proven to generate high value.

The same phenomenon arises when environmental uncertainty is high. In such an environment, an organization needs to keep up with the changing environment by continuously modifying the organizational knowledge. However, when agents are equally likely to be conducting individual searches and adopting practices from the organizational knowledge, they fail to agree on how to adjust the organizational knowledge and, as a consequence, organizational performance suffers.

When congruency is high but not overly strong ($\lambda \approx 0.2$), such disagreements about how to modify the organizational knowledge do not arise as frequently. Most members of the organization share the same ideas and search in similar directions. Hence, organizational knowledge will improve continuously as long as there is sufficient knowledge sharing within the organization. And, the performance of such organizations eventually exceeds that of organizations with a low or an intermediate level of congruency.

The rest of the paper is structured as follows: in Section 2 we present the model studied in this paper, in Section 3 we present and discuss the results of simulation, and in Section 4 we conclude the discussion.

2 Model

Consider an organization that consists of M agents. Each agent undertakes an identical task having N dimensions. The value that an agent generates depends on how the agent configures each dimension of the task. The performance of the organization depends on the values generated by the agents therein.

Let $x_j^i(t) \in \{0,1\}$ be agent i's configuration of dimension j in period t,[3] and $X^i(t) \in \{0,1\}^N = \{x_1^i(t), x_2^i(t), ..., x_N^i(t)\}$ be i's configuration of the task in period t. The corresponding value agent i generates, or the performance of agent i in period t, is $\Pi^i(t) = \pi(X^i(t))$. The performance of the organization, $\Pi(t)$, is defined simply as the mean performance of its members, i.e., $\Pi(t) = \frac{1}{M}\sum_i \Pi^i(t)$.

We formulate the performance function $\pi(\cdot)$ based on the NK Landscape model (Kauffman, 1993), which allows us to parameterize the interdependencies among N dimensions with a parameter K. Namely, the ideal configuration for dimension j depends on i's configurations of K other dimensions.[4] Let $\Gamma_j = \{l_1, l_2, ..., l_K\}$ be the set of these K dimensions that affect the effectiveness of $x_j^i(t)$, and $X_j^i(t) = \{x_{l_1}^i(t), x_{l_2}^i(t), ..., x_{l_K}^i(t)\} \in \{0,1\}^K$ be i's configurations of these K dimensions in period t. Then,

$$\pi(X^i(t)) = \frac{1}{N}\sum_{j=1}^{N} q_j(x_j^i(t), X_j^i(t))$$

[3] The assumption that each dimension can be either zero or one is made for the sake of simplicity. One can easily extend the model so that a dimension can be configured in many ways. If the number of possible configurations of a dimension is c, then the number of possible configurations of the task becomes c^N.

[4] We assume that these K dimensions are chosen randomly from $N-1$ other dimensions. One can also consider many other possible structures of interdependencies. Rivkin and Siggelkow (2007) demonstrate that the structure of interdependencies, even controlling for K, affects the complexity of the environment and the effectiveness of various search strategies.

where $q_j(x_j, X_j)$ is defined by assigning values drawn randomly from $U[0, 1]$ to each possible x_j and X_j. Every agent faces the same performance function $\pi(\cdot)$ or $\{q_j\}$.

When modeled in this way, the number of possible configurations of the task is 2^N, a potentially large space in which agents must search for better configurations. The parameter K captures the complexity of the task. When $K = 0$, there is no interdependency among dimensions; thus, changing the configuration of one dimension results in smooth changes in performance. The larger K is, the more interdependencies there are among different dimensions. When K is large, changing the configuration of one dimension has a non-additive effect on the value generated by the agent.

To capture the uncertainty of the environment in which organizations operate, we introduce a parameter $\mu \in [0, 1]$ such that, in each period and for each possible configuration of the task X, $q_j(x_j, X_j)$ is redefined randomly to a value drawn from $U[0, 1]$ with probability μ. When μ is zero, the values associated with each possible configuration of the task remain constant over time. When the value of μ is higher, these values can change quite drastically from period to period.

The NK landscape has been applied in the literature on organization theory. In their series of papers, Rivkin and Siggelkow have considered how to design a decision-making process in hierarchical (or multi-level) organizations undertaking complex projects in uncertain environments (Rivkin and Siggelkow, 2003, 2007; Siggelkow and Rivkin, 2005, 2006). In particular, one of their foci has been to understand the relationship between the levels of decision making when lower layers of hierarchy have narrower scope. Other applications of NK landscape in management literature include, for example, Gavetti and Levinthal (2000). Gavetti and Levinthal (2000) consider the interplays between a forward-looking "off-line" search and a backward-looking "on-line" search and their implications to organizational performance. Our analysis departs from the existing literature in the following two respects: (1) we do not consider hierarchical organization, and (2) the focus of our analysis is on organizational congruency, which has not been considered formally in the literature. We now turn to how the behaviors of organizations are modeled.

2.1 Search, learning, and organizational congruency

We assume that each agent receives a randomly configured task initially; i.e., $x_j^i(0) \in \{0, 1\}$ are randomly set for all i and j. Agents modify their configurations over time by conducting individual searches and by learning from the organizational knowledge. Let $\Omega(t) = \{o_1(t), o_2(t), ..., o_N(t)\}$ where $o_j(t) \in \{0, 1\}$ represents the organizational knowledge at period t. We assume that $\Omega(0)$ is randomly set. We allow the organizational knowledge itself to evolve over time as agents in the organization contribute their knowledge in the manner described below. Note that knowledge sharing improves organizational performance because agents face the same $\{q_j\}$.

In each period, each agent either conducts an individual search for a better configuration (with probability λ) or learns from the organizational knowledge (with probability $1 - \lambda$). The individual searches are conducted as follows: an agent chooses one of the N dimensions randomly and ascertains whether changing its configuration generates a greater value. If it does, he adopts the change. Otherwise, his configuration remains as before.

When an agent learns from the organizational knowledge, he randomly chooses a dimension such that his current configuration differs from that of the organizational knowledge and adopts the organizational knowledge for the chosen dimension. Copying a part of the organizational knowledge in this manner gradually assimilates the individual knowledge base to the organizational knowledge if λ is sufficiently low. We call this process *alignment* because the assimilated knowledge base facilitates upgrading of the organizational knowledge as explained later. When adopting the configuration from the organizational knowledge, the agent does not check whether doing so generates a higher value. We assume that the organization's management institutes a monitoring and incentive mechanism to enforce the alignment of the organizational base.

On the one hand, when λ is low, agents tend to follow the configurations of the organizational knowledge; thus, their practices are all aligned with each other. When λ is high, on the other hand, since agents pursue individual searches, their configurations may remain diverse. Therefore, we can interpret the parameter λ to represent the level of organizational congruency (i.e., high λ indicates low organizational congruency). It is of our interest to discover the relationship between the degree of organizational congruency represented by λ and organizational performance under various levels of complexity and uncertainty.

At the end of each period, with probability p, each agent proposes to modify the organizational knowledge. When an agent proposes to modify the organizational knowledge, he randomly chooses one dimension such that his configuration differs from that of organizational knowledge. And, he proposes to change the configuration of the chosen dimension in the organizational knowledge so that it will be the same as his current configuration. Whether the proposal is accepted or not depends on voting by agents in the organization. It is accepted if the majority of the agents whose performance is no less than the performance of the organization[5] configure the dimension as proposed.[6] The parameter p can be interpreted as the frequency of knowledge sharing within the organization. The higher the value of p is, the more quickly organizational knowledge improves.

The model presented here is similar to that of March (1991), which studies the relationship between exploitation of existing knowledge and exploration of new knowledge in an abstract model. In particular, both our model and that of March (1991) consider interactions between individual search and learning within an organization through communications among its members. Our model differs, however, from that of March (1991) in its assumption of the underlying configuration space. While the latter assumes the existence of a unique configuration that an organization needs to discover, our model employs the NK landscape so that there can be many distinct locally-best configurations when a task is complex. We are not aware of existing research that formally attempts to understand the effect of organizational congruency in the face of such multiplicities of locally-best configurations.

[5]That is, we consider the majority of those agent k such that $\Pi^k(t) \geq \Pi(t)$.
[6]In the case of a tie, a proposal is approved with probability 1/2.

1. Set parameter values, and give a random seed
2. Initialize NK landscape, agents, and organizational knowledge
3. for $1 \leq t \leq 1000$
    (a) for each agent $1 \leq i \leq M$
        (i) Search (λ) or Copy ($1 - \lambda$)
        (ii) update the agent's configuration and performance
    (b) Update performance of organization
    (c) for each agent $1 \leq i \leq M$ (in random orders)
        (i) Propose (with prob. p) or not.
        (ii) If proposes, evaluate the proposal
        (iii) update organizational knowledge
    (c) if $\mu > 0$, update NK landscape, the performances of agents and the organization
    (d) Back to (a)

Table 1: Pseudo Code

3 Results

There are six parameters in our model. The task is defined by its size N and its complexity K. The environmental uncertainty is captured by μ. The organization is characterized by the number of agents, M, the frequency of knowledge sharing among them, p, and the level of organizational congruency, λ. In all the simulations, we fix N and M to be 100 and 20, respectively, and vary other parameters to investigate the effect of congruency and communications on the performance of organizations under various degrees of complexity and uncertainty.[7]

For each set of parameter values, the payoff function, $\pi(\cdot)$, initial configuration for individual task $X^i(0)$, and organizational knowledge, $\Omega(0)$, are generated. Then, we allow organizations to operate in the manner described in the Model section for 1000 periods. Table 1 shows the pseudo code of the simulations.

To summarize the performance of an organization over time, we mainly focus on present discounted values (PDVs) of organizational performance over these 1000 periods: namely, $\Pi = \sum_{t=1}^{1000} \delta^t \Pi(t)$ where δ is a discount factor. Below, we will set the discount factor to be $\delta = 0.999$.[8] This discount factor gives the payoff for the last

[7] The appendix shows the results from $M = 10$ and $M = 30$, which are similar to the results from $M = 20$.

[8] If we set the discount factor too low, the payoff at $t = 1000$ will be given a negligible weight compared to those for earlier periods. For example, if $\delta = 0.95$, the weight for the payoff at $t = 1000$ will be less than

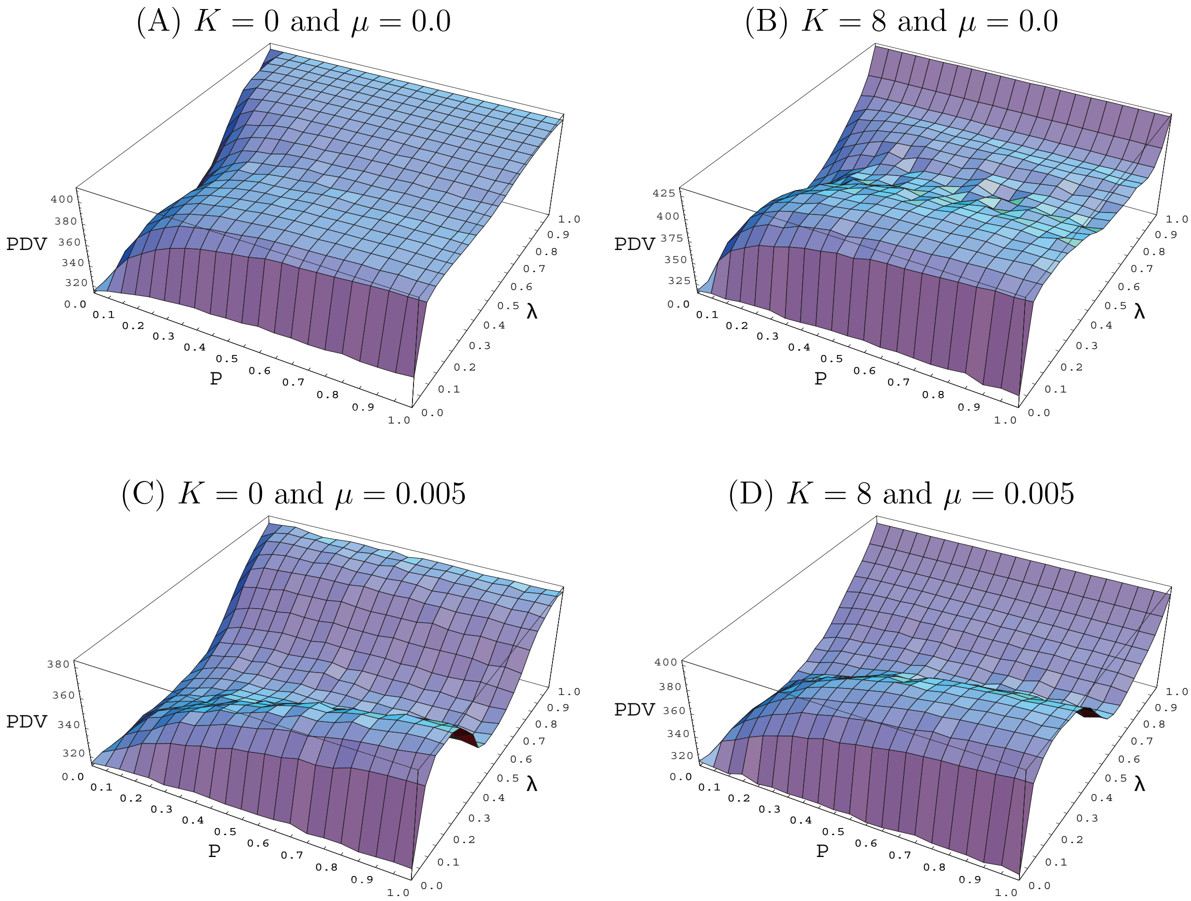

Figure 1: Average PDVs of organizational performances for various p and λ in four combinations of complexities ($K = 0$ (left) and $K = 8$ (right)) and uncertainties ($\mu = 0.0$ (top) and $\mu = 0.005$ (bottom)). Discount factor is $\delta = 0.999$. Data are generated by taking average over 100 simulation runs for each set of parameter values.

period a weight that is 0.37 of that for the first period. We also take the average over 100 simulation runs based on varying random seeds.

3.1 Effect of p and λ

Figure 1 shows the PDV (shown by height) for various degrees of activeness of communication p (the horizontal axis) and of control λ (the vertical axis) under four combinations of complexities, $K \in \{0, 8\}$, and uncertainties, $\mu \in \{0.0, 0.005\}$. Recall that a higher λ represents weaker organizational control (or congruency).

0.01% of the payoff at the period $t = 1$. Since it is possible for performance to take a long time to level off, we would like the payoff at $t = 1000$ to have enough weight compared to the initial payoffs.

In all four panels in Figure 1, one can see that organizational performance measured by PDV becomes higher as we move along the horizontal axis from left to right. That is, the more frequently individuals engage in knowledge sharing, the higher their organizational performance, regardless of task complexity and environmental uncertainty.

While the figure exhibits a monotonic relationship between PDV and frequency of knowledge sharing, p, the relationship between PDV and level of organizational congruency, λ, is not monotonic in three out of four cases shown in Figure 1. Except for cases in which the task is simple and there is no environmental uncertainty (Panel (A)) or the level of knowledge sharing in the organization is very low (low p), the performance of the organization has two peaks: (1) when organizational congruency is very low, i.e, $\lambda \approx 1.0$; and when there is a substantial, but not overly strong, degree of congruency, $\lambda \approx 0.2$. We call the former, a *high-initiative* organization and the latter a *high-alignment* organization. When λ is very small (so that individual searches are seldom conducted) or has intermediate values of around 0.5 or 0.6, performance becomes low.

The non-monotonic relationship between λ and organizational performance can be better seen in Figure 2. For the same four combinations of K and μ as in Figure 1, Figure 2 plots PDVs against various values of λ for three values of p: $p = 0.0$ (solid black), $p = 0.5$ (solid gray), and $p = 1.0$ (dashed black). One can see from the figure that for $\lambda \approx 0.2$ to generate as high a level of performance as $\lambda \approx 1.0$, the efforts to share knowledge must be persistent (i.e., high p). Another point indicated by Figure 2 is that in the case of complex tasks ($K = 8$), organizational performance declines sharply as one moves away from the state of full individual search, i.e., $\lambda = 1.0$, while there is a wider range of λ with values around $\lambda = 0.2$ that maintains high performance. This suggests that when knowledge sharing activities are sufficient, any shift from full exploration (i.e., $\lambda = 1.0$) may require a substantial leap toward a state of high congruency.[9]

[9] Figure 2 also shows that when the task is simple ($K = 0$), the introduction of a small degree of congruency can improve performance compared to the case of no congruency.

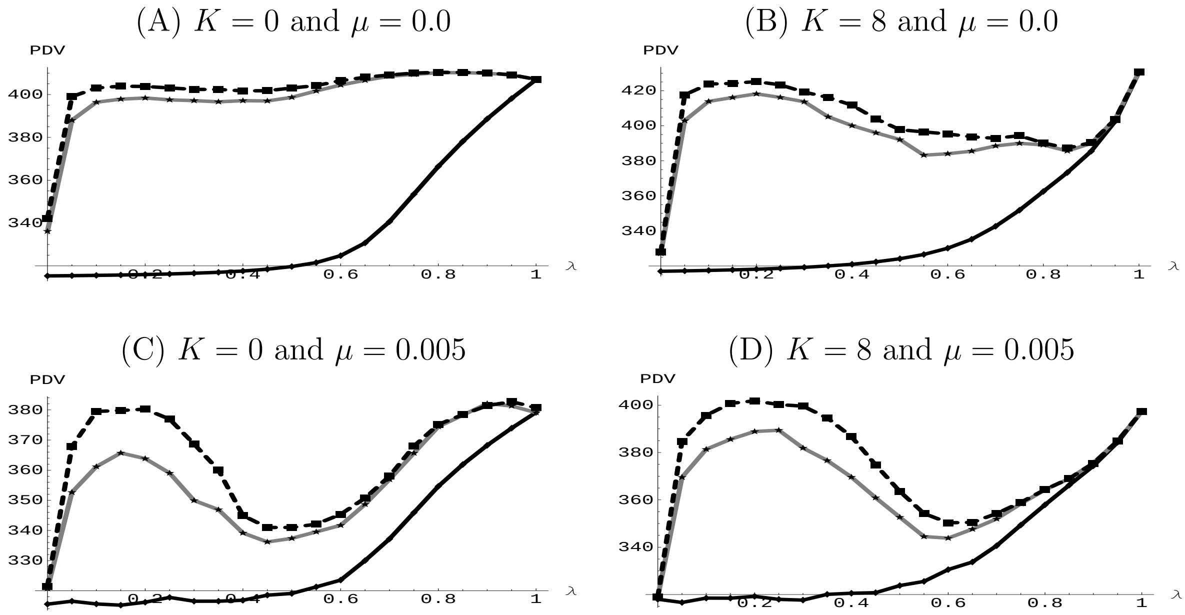

Figure 2: Average PDVs of organizational performances for various λ in four combinations of complexities ($K = 0$ (left) and $K = 8$ (right)) and uncertainties ($\mu = 0.0$ (top) and $\mu = 0.005$ (bottom)). In each figure, results for three distinct ps are reported: $p = 0.0$ (solid black), $p = 0.5$ (solid gray), and $p = 1.0$ (dashed black). Discount factor is $\delta = 0.999$. Data are generated by taking average over 100 simulation runs.

Why do successful organizations, in the face of complex tasks and uncertain environments, bifurcate into two types: high-initiative ($\lambda = 1.0$) vs. high-alignment ($\lambda \approx 0.2$)? In other words, why do we obtain non-monotonic relationships between λ and organizational performance? To answer this question, we turn to dynamics of performance over time.

3.2 Dynamics

Figure 3 shows the dynamics of organizational performance for four values of λ, $\lambda \in \{0.0, 0.2, 0.6, 1.0\}$ (shown in solid black, solid gray, dashed black and dashed gray, respectively), for the four combinations of $K \in \{0, 8\}$ and $\mu \in \{0.0, 0.005\}$ considered in Figures 1 and 2. These four λs are chosen because they correspond approximately to the local maxima and minima of organizational performance measured by PDV. Knowledge sharing within an organization is assumed to be very frequent, $p = 1.0$,

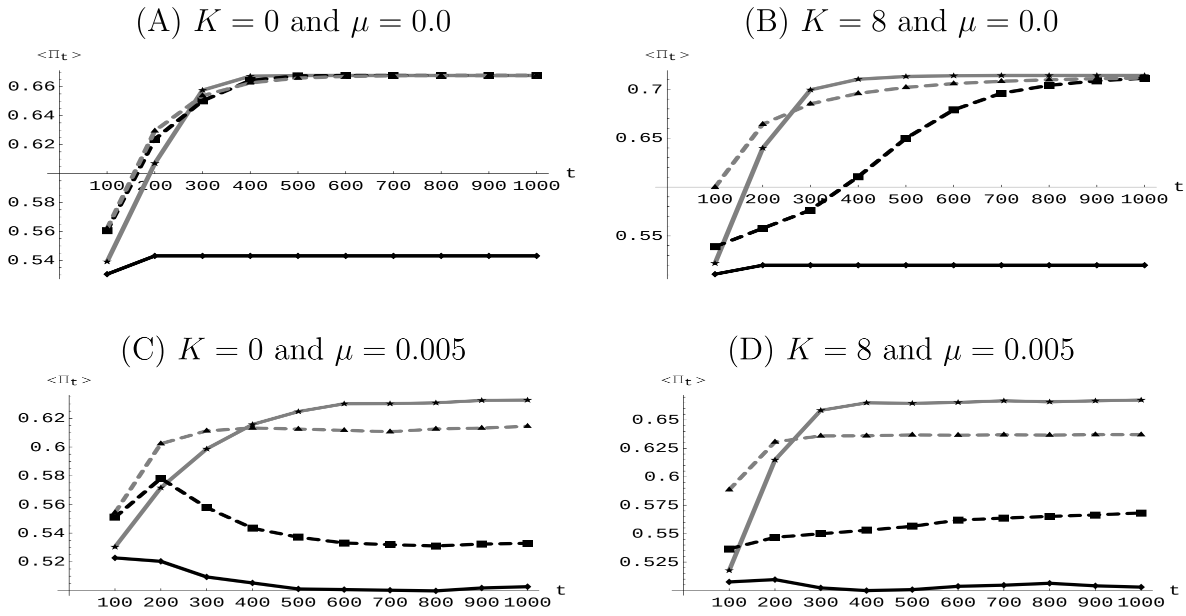

Figure 3: Average performance over time for two levels of complexity, $K = 0$ (left) and $K = 8$ (right), and two degrees of environmental uncertainty, $\mu = 0.0$ (top) and $\mu = 0.005$ (bottom). Four values of λ are considered: $\lambda = 0.0$ (solid black) ,$\lambda = 0.2$ (solid gray), $\lambda = 0.6$ (dashed black), $\lambda = 1.0$ (dashed gray). The communication is active ($p = 1.0$). The data is generated by taking the average per period organizational performance for each block of 100 periods. The average from 100 simulation runs is reported.

which corresponds to the dashed black curves in Figure 2.[10] Figure 3 reports the average per-period organizational performance for each block of 100 periods, i.e., for $t \in [1, 100]$, $t \in [101, 200]$ and so on.

When the task is simple and there is no environmental uncertainty, organizations with lower degrees of congruency (i.e., higher λs) demonstrate faster improvements in performance (Panel (A)). When $K = 0$, the best configuration for the task can be found through a local search procedure, such as the individual searches considered in this paper. Therefore, all the agents in an organization will eventually find the unique best configuration as long as $\lambda > 0.0$. (Note that when $\lambda = 0.0$, there is no individual search, so that agents never find the best configuration). Once everyone in an organization has found the best configuration, there is no further improvement in organizational performance. When all the agents in an organization only conduct individual searches and do not copy from organizational knowledge (i.e, $\lambda = 1.0$), each agent rapidly finds the best configuration. Organizations that employ full individual search ($\lambda = 1.0$) are not especially superior to organizations with higher levels of congruency ($\lambda < 1.0$) because organizations can achieve very good performance provided knowledge sharing takes place frequently enough and organizational knowledge is improved rapidly. As knowledge sharing becomes less frequent (i.e., lower p), however, higher congruency imposed on individuals (i.e, lower λ) would instead disturb individual searches and retard organizational performance.

When the task is complex, the results are quite different, even in the complete absence of environmental uncertainty. Panel (B) of the figure shows that a higher value of λ corresponds to a higher performance only at early periods, namely the first 100 periods. In later periods, for example for $t \in [101, 200]$, while $\lambda = 1.0$ (dashed gray) exhibits the highest performance, the performance of $\lambda = 0.6$ (dashed black) is lower than that of $\lambda = 0.2$ (solid gray). Eventually, the performance of $\lambda = 0.2$ exceeds that of $\lambda = 1.0$. This case demonstrates the possible trade-off between the performances in

[10] As we have seen above, a lower value of p results in lower performance for an organization with λ smaller than one.

earlier periods and those in later periods.[11] It should also be noted that the average performance in all three cases except for $\lambda = 0.0$ converge to the same level in the final 100 periods as shown in the Panel (B) of the figure.

In the presence of uncertainty (Panels (C) and (D)), the performance of the organizations with $\lambda = 0.2$ and $\lambda = 1.0$ exhibit a similar pattern, as seen in the case shown in Panel (B). Namely, initially, the organization with $\lambda = 1.0$ demonstrates higher performance than the one with $\lambda = 0.2$, but in the later periods the performance of the latter exceeds that of the former. The dynamics of the performance for $\lambda = 0.6$ (solid black) is quite different in the face of environmental uncertainty than in its absence. That is, when there is sufficient uncertainty, performance does not improve much over time (in the case of a complex task, Panel (D)) or can even deteriorate (in the case of a simple task, Panel (C)). Note that due to the environmental uncertainty organizational performance in an individual simulation run demonstrates ups and downs over time. Such volatilities are hidden, however, in Figure 3, where we plot the averaged performance in each block of 100 periods and further take averages across 100 simulation runs. Nonetheless, deterioration of average performance appears for $\lambda = 0.6$ after period 200. This puzzling outcome deserves more careful analyses.

Why does the reversal of relative performance between $\lambda = 0.2$ and $\lambda = 1.0$ take place during the course of organizational learning? And, why does performance for $\lambda = 0.6$ in an uncertain environment with a complex task deteriorate after an initial improvement?

Looking at the dynamics of diversity of individual configurations and the dynamics of the value of organizational knowledge, $\pi(O(t))$, helps us to answer these questions. The former is plotted in Figure 4 and the latter in Figure 5.

The diversity of the individual configurations, $D(t)$, is measured by the average distance between individual configurations and their means. The mean configuration of dimension j at period t is $\bar{x}_j(t) = \sum_i x_j^i(t)/M$. The distance between configurations of

[11]Therefore, if we set the discount factor δ too low, we pick up only performance in earlier periods and fail to capture this trade-off.

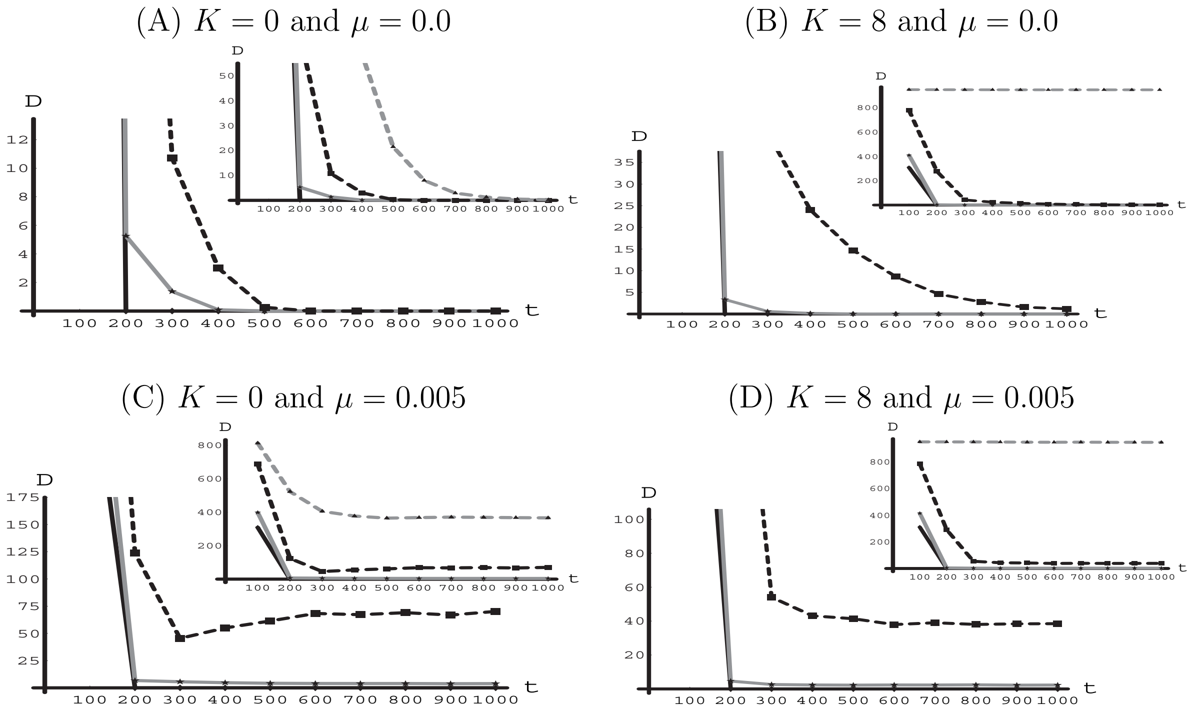

Figure 4: The diversity within an organization for two levels of complexity, $K = 0$ (left) and $K = 8$ (right), and two degrees of environmental uncertainty, $\mu = 0.0$ (top) and $\mu = 0.005$ (bottom). The communication is active ($p = 1.0$). Four values of λ are considered: $\lambda = 0.0$ (solid black), $\lambda = 0.2$ (solid gray), $\lambda = 0.6$ (dashed black), $\lambda = 1.0$ (dashed gray). For the clarity of exposition, only $\lambda \in \{0.0, 0.2, 0.6\}$ are shown in the main figure. See the insets for all the four λs. The extent of diversity is measured based on the discrepancy between individual configurations and their means. The data is generated by taking the average per period diversity for each block of 100 periods. The average from 100 simulation runs is reported.

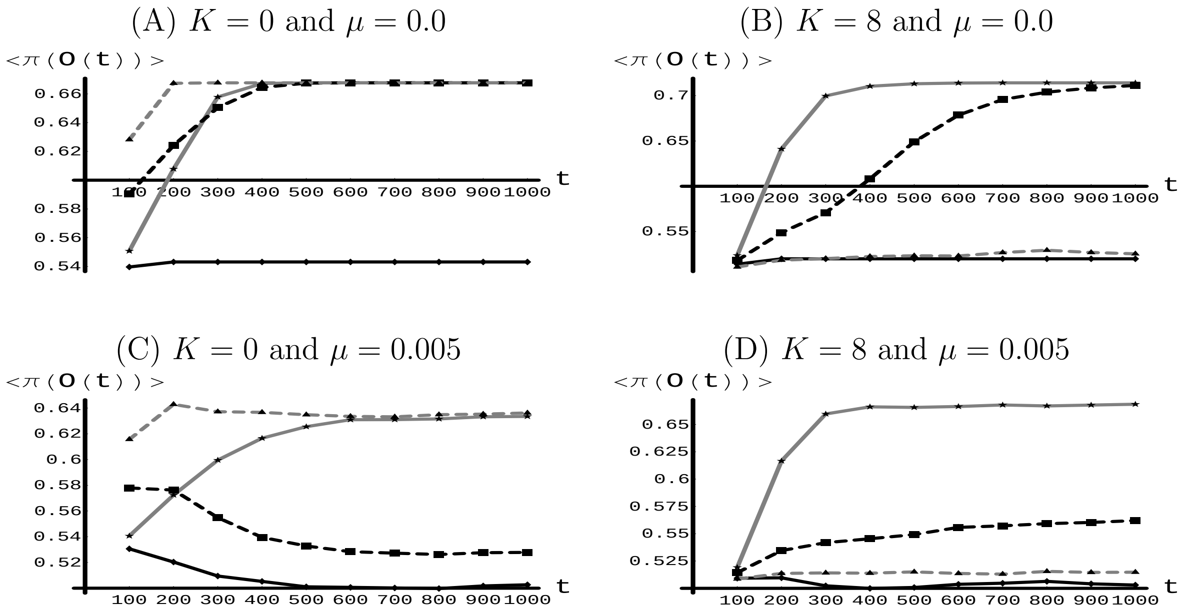

Figure 5: The dynamics of the average value of organizational knowledge for two levels of complexity, $K = 0$ (left) and $K = 8$ (right), and two degrees of environmental uncertainty, $\mu = 0.0$ (top) and $\mu = 0.005$ (bottom). Four values of λ are considered: $\lambda = 0.0$ (solid black) ,$\lambda = 0.2$ (solid gray), $\lambda = 0.6$ (dashed black), $\lambda = 1.0$ (dashed gray). The communication is active ($p = 1.0$). The data is generated by taking the average per period value of organizational knowledge for each block of 100 periods. The average from 100 simulation runs is reported.

individual i and the mean configurations at period t is therefore $D^i(t) = \sum_{j=1}^{N} |x_j^i(t) - \bar{x}_j(t)|/N$. The diversity for an organization in period t is $D(t) = \sum_{i=1}^{M} D^i(t)$.

As in the case of organizational performance shown in Figure 3, we took the average over each block of 100 periods.[12] The results for four combinations of complexities $K \in \{0, 8\}$ and uncertainties $\mu \in \{0.0, 0.005\}$ are plotted. Four values of λ, $\lambda \in \{0.0, 0.2, 0.6, 1.0\}$ (shown in solid black, solid gray, dashed black, and dashed gray, respectively) are shown. $p = 1.0$ is assumed.

When the task is simple and there is no environmental uncertainty, the diversity measure converges to zero for all types of organizations plotted in Figure 4 (Panel (A)). The lower λ is, the faster is convergence. The reason even an organization with $\lambda = 1.0$ (dashed gray, shown only in the in-sets) demonstrates zero diversity in the later period

[12]Figure 3 reports the results obtained from taking the average of these averaged per-period diversities across 100 simulations.

is that, as discussed above, when $K = 0$ and $\mu = 0.0$, every agent in the organization eventually finds the unique best configuration through individual searches unless λ is zero. And, since everyone eventually agrees on the best configuration, organizational knowledge will also be configured accordingly. Therefore, the value of organizational knowledge shows rapid improvement, as shown in Figure 5 (Panel (A)). When $\lambda = 0.0$ (shown in solid black), this is not the case. Although the diversity measure goes to zero as well, the value of organizational knowledge shows little improvement. Diversity disappears because everyone adopts configurations from the same organizational knowledge. However, because agents never search for better configurations, they do not have new information necessary to improve organizational knowledge.

When there is no uncertainty but the task is complex ($K = 8$, $\mu = 0.0$), as shown in Figure 4 (Panel (B)), diversity remains high for organizations with $\lambda = 1.0$ (dashed gray, shown only in in-set) for long periods of time, while the diversity measures in the organizations with $\lambda = 0.0$ and $\lambda = 0.2$ quickly converge to zero. In the case of $\lambda = 0.6$, the diversity measure declines much more slowly than in the cases of smaller values of λ.

The high diversity for organizations with $\lambda = 1.0$ under the complex task is due to the existence of many distinct but locally-best configurations. And, since such locally-best configurations vary from one another, diversity among the individual configurations remains high and agents in the organizations do not agree on how to improve organizational knowledge. As a result of the inability to aggregate information and choose the best among many local optima, the value of organizational knowledge remains low (Figure 5, Panel (B)). Although the value of organizational knowledge for the $\lambda = 0.0$ organization also shows little improvement, this is due to the complete absence of individual searches.

Panel (B) of Figure 5 also shows that an organization with $\lambda = 0.6$ is much slower in improving organizational knowledge than one with $\lambda = 0.2$. Recall that, in our model, in order to modify organizational knowledge, not only must new proposals be submitted, they must also be approved. In the model considered here, a proposal is approved when a majority of better-performing agents configure the dimension as proposed. Also, as discussed above for the case of $\lambda = 1.0$, when a task is complex, individual searches may lead to various distinct configurations that are local optima. If many agents search in different directions, it is more difficult for them to agree on how to modify organizational knowledge. The insufficient alignment of the individual knowledge base resulting from lower congruency creates a bottleneck to improvement of organizational knowledge, which in turn slows down the improvement of organizational performance.

The inability to improve organizational knowledge can be detrimental when the environment is changing. In such an unstable environment, the quality of organizational knowledge itself can deteriorate over time, as one can see from Panel (C) of Figure 5. Note that an organization with $\lambda = 1.0$ is not capable of improving organizational knowledge when the task is complex, but it does not suffer from the bad organizational knowledge because agents never adopt bad configurations of the organizational knowledge. However, when $\lambda = 0.6$, an intermediate level, agents can still have sufficiently different knowledge configurations, which in turn hinders organizational knowledge from improving, as shown in Panel (C) of Figure 5. This explains why the average performance in organizations with $\lambda = 0.6$ deteriorates, as can be seen in Panel (C) of Figure 3.

3.3 Optimal degree of congruency

So far we have considered four combinations of complexities of task ($K \in \{0, 8\}$) and uncertainties of environment ($\mu \in \{0.0, 0.005\}$). We have seen that when the environment is uncertain or the task is complex, successful organizations bifurcate into two types: *high-initiative* vs. *high-alignment* organizations. We have also noted that for the latter to be successful, very frequent knowledge sharing within the organization is required to improve and maintain the quality of organizational knowledge. We would also like to find out the conditions under which a high-alignment organization outperforms a high-initiative one. Figure 6 illustrates which of the organizations with

(A) $\lambda \in \{0.2, 0.6, 1.0\}$ or those with (B) $\lambda \in \{0.2, 0.6, 0.95\}$ performs best under various levels of complexity of the task (K) and various degrees of environmental uncertainty (μ). We assume very frequent knowledge sharing $p = 1.0$. We introduce $\lambda = 0.95$ instead of $\lambda = 1.0$ in Panel (B) of the figure to test robustness of the results from the comparison. As we have seen above, in the face of complex tasks, organizational performance declines very sharply as one moves away from $\lambda = 1.0$, while there exists a wider range of λ around 0.2 that generates high performance.

In Panel (A) of the figure, the white (black) cells in the figure show the combinations of K and μ where an organization with $\lambda = 1.0$ ($\lambda = 0.2$) performs the best. For Panel (B), the white cells represent the combinations where $\lambda = 0.95$ performs the best.[13] The gray cells without the indication "ND" show the cases where $\lambda = 1.0$ (or $\lambda = 0.95$ in the case of Panel (B)) and $\lambda = 0.2$ perform equally well and are both better than $\lambda = 0.6$. The gray cells with the indication "ND" are the region where all three λ perform equally well.

One can see from Panel (A) that when uncertainty is low ($\mu < 0.001$), a high-initiative organization ($\lambda = 1.0$) exhibits superior performance even when the task is quite complex.[14] The advantage of a high-initiative organization disappears once a slight degree of alignment is introduced (Panel (B)), unless the task is simple. Overall, our results are not robust for $\mu < 0.005$. In other words, there is not much significant difference between high-initiative and high-alignment organizations in a relatively stable environment.

As the environment becomes more uncertain but not overly so ($\mu = 0.005$ or 0.01), a high-alignment organization ($\lambda = 0.2$) demonstrates its advantage assuming that the task is reasonably complex. When the environment becomes extremely unstable

[13] We have performed two-sample t-test based on PDVs generated by 100 simulations for each set of parameter values. Depending on the results of a variance comparison test, unequal variance or equal variance is assumed in performing t-test. One organization is said to outperform the other if the mean PDV is significantly greater at 5% significance level in a one-tailed test.

[14] Note that we are measuring performance with PDVs that put higher weights on earlier periods than later ones. As seen in the previous section, if we compare the average performance in later periods (or place more weight on the performance of later periods by using a higher discounting factor δ), in cases of very frequent knowledge sharing, $\lambda = 0.2$ is better than $\lambda = 1.0$.

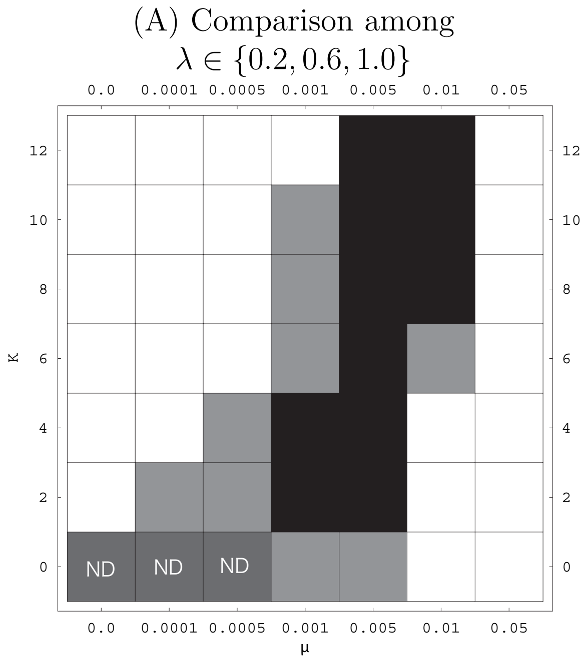

Figure 6: Comparison of PDV of (A) $\lambda \in \{0.2, 0.6, 1.0\}$ and (B) $\lambda \in \{0.2, 0.6, 0.95\}$ for various degree of complexity K and uncertainty μ. $p = 1.0$ and discount factor is 0.999 in calculating PDVs. White and black areas indicate $\lambda = 1.0$ ($\lambda = 0.95$ in panel B) and $\lambda = 0.2$ demonstrate the highest performance, respectively. Gray areas indicate no difference between $\lambda = 1.0$ ($\lambda = 0.95$ in panel B) and $\lambda = 0.2$, and both are better than $\lambda = 0.6$. Gray areas with "ND" indicate that all three λs perform equally well.

($\mu = 0.05$), though, a high-alignment organization loses its relative advantage and a high-initiative organization tends to perform better.

This exercise reveals that institutionalizing knowledge in the form of routines, procedures, strategies, and systems generally creates advantage when the business is reasonably complex and uncertain. The cost of institutionalizing knowledge is the delay caused by the aggregation of individual knowledge. Furthermore, the assimilation of the individual knowledge base could also hinder the improvement of organizational knowledge by reducing the range of individual searches when the environment changes drastically. Potential costs discussed above make it suboptimal to choose a high-alignment organization in an extremely unstable environment.

4 Conclusion

This work demonstrates the origin of the primary tension between exploration and exploitation and conditions in which choosing an "extreme" type of organizational learning is optimal. We identify important roles played by organizational congruency in facilitaing organizational learning. Our results show non-concavity in the optimization problem of organization design and imply that two types of organization with distinct natures could emerge. A high-initiative organization promotes individual initiatives to experiment with new ideas and build its strength on individual learning. A high-alignment organization, in contrast, assimilates the individual knowledge base and accelerates organizational learning through frequent knowledge sharing among individuals. An organization with a more equal mix of individual initiatives and knowledge alignment tends to perform worse, and especially so when the operation is reasonably complex (in other words, interdependency is high enough) and/or the business environment is reasonably uncertain. Although a high-alignment organization tends to be favored in a moderately uncertain environment, a high-initiative organization tends to be favored when the environment becomes extremely unstable.

There are two remaining issues to be considered. First, how robust are the results we have illustrated in this paper to changes in the assumptions? Specific rules and processes of updating organizational knowledge in our model must be relaxed to examine whether our bifurcation results continue to hold under different assumptions. Second, we did not consider the possibility that the degree of organizational congruency changes as time goes by. Such a question is especially formidable if parameter values, such as complexity of operation K, were to change in the course of a firm's growth. We believe that these questions can be successfully addressed by extending the basic framework we employed in this paper, but will leave them to future research.

References

CROSSAN, M. M., H. W. LANE, AND R. E. WHITE (1999): "An Organizational Learning Framework: From Intuition to Institution," *Academy of Management Review*, 24(3), 522–537.

GAVETTI, G., AND D. LEVINTHAL (2000): "Looking Forward and Looking Backward: Cognitive and Experimental Search," *Administrative Science Quarterly*, 45, 113–137.

GHEMAWAT, P., AND J. E. RICART I COSTA (1993): "The Organizational Tension between Static and Dynamic Efficiency," *Strategic Management Journal*, 14, 59–73.

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

LEVINTHAL, D. A., AND J. G. MARCH (1993): "The Myopia of Learning," *Strategic Management Journal*, 14, 95–112.

MARCH, J. G. (1991): "Exploration and Exploitation in Organizational Learning," *Organization Science*, 2(1), 71–87.

RIVKIN, J. W., AND N. SIGGELKOW (2003): "Balancing Search and Stability: Interdependencies Among Elements Of Organizational Design," *Management Science*, 49, 290–311.

——— (2007): "Patterned Interaction in Complex Systems: Implications for Exploration," *Management Science*, 53, 1068–1085.

ROBERTS, D. J. (2004): *The Modern Firm: Organizational Design for Performance and Growth*. Oxford University Press.

SIGGELKOW, N., AND J. W. RIVKIN (2005): "Speed and Search: Designing Organizations for Turbulence and Complexity," *Organization Science*, 16, 101–122.

——— (2006): "When Exploration Backfires: Unintended Consequences of Multi-Level Organizational Search," *Academy of Management Journal*, 49, 779–795.

A Results from $M = 10$ and $M = 30$

The main text concentrated on discussion of the results for $M = 20$. This appendix presents the simulation results from $M = 10$ and $M = 30$, which are qualitatively the same as the results in the main text. Figure 7 shows, in the same format as in Figure 1, PDV with discount factor $\delta = 0.999$ for various values of p and λ for four combinations of K and μ. $M = 10$ is on the top row and $M = 30$ is on the bottom row. As one can see from comparing Figure 7 and Figure 1, the main results do not change as we change the number of agents in the organization. What does change is the range of values of $\lambda < 1.0$ that maximizes performance.

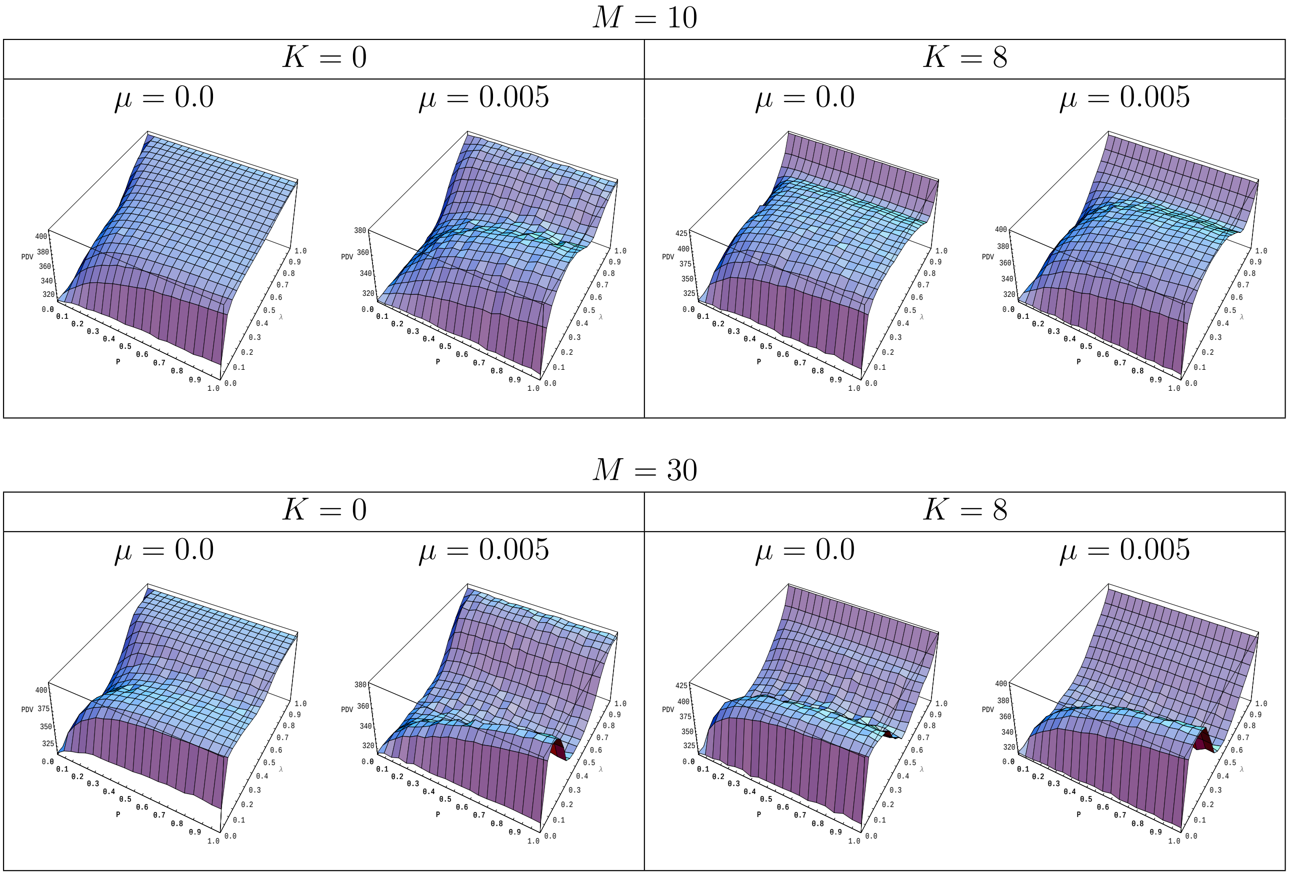

Figure 7: Present Discounted Values of organizational performance for $M = 10$ (top) and $M = 30$ (bottom) for four combinations of $K \in \{0, 8\}$ and $\mu \in \{0.0, 0.005\}$