

**Managing radical and incremental change:
Organizational search in changing environments**

Markus Becker, Thorbjørn Knudsen, Nils Stieglitz

Strategic Organization Design Unit (SOD)
Dep. of Marketing and Management, University of Southern Denmark
Campusvej 55, DK-5230 Odense M, Denmark
mab@sam.sdu.dk; Tel +45 6550 1000
www.sdu.dk/sod

Draft paper, 14 September 2007

Abstract: The paper uses a NK simulation model to study organizational search in firms confronted by both radical and incremental change. Our analysis suggests that organizations must be able to systematically adapt their proclivity to explore over time. They increase exploration in times of radical change and decrease it during times of tranquillity. A modicum of exploration is required at all times to provide the organization with a sensing capability that tracks discontinuous changes. The effectiveness of organizational adaptation may be improved by the firm-wide reallocation of resources to new promising task environments. Our results show that coordination could be a way to overcome the shadow of prior history, since it takes time until organizational members update their expectations about the performance characteristics of task environments.

Keywords: Search, exploration, exploitation, coordination, discontinuous change

1. Introduction

Organizations face a fundamental challenge in adapting to both discontinuous and continuous change. They need to sense discontinuous change and reallocate effort to new, promising capabilities, while continuously refining existing capabilities to realize their full potential. However, our understanding of these two issues is rather limited. In one way or the other, organizations have to dynamically balance the antagonistic demands of exploration and exploitation (March, 1991; Levinthal & March, 1993). An exclusive reliance on exploitation jeopardizes an organization's ability to sense and respond to discontinuous change (Cohen & Levinthal, 1990). In contrast to exploration, the refinement of existing capabilities provides immediate performance rewards that reinforce prior choices by positive feedback. Firms therefore fall prey to the success trap in capability development, since positive performance feedback traps the firm in a narrow domain. A focus on the more uncertain and distant benefits of exploration, however, may lead the firm into the failure trap. Initial negative feedback stimulates ever more exploration that again leads to negative feedback. The organization constantly tries and abandons new capabilities. By not committing itself to a capability, the organization foregoes the potential benefits of continuous refinement. Successful organizational adaptation clearly requires a mixture of exploration and exploitation.

Prior research has identified two principal approaches to dynamically balancing exploration and exploitation (Levinthal & March, 1993; Gupta et al., 2006; Greve, 2007). A temporal sequencing of exploration and exploitation suggests that an organization should dynamically switch between short bursts of exploration followed by extended periods of exploitation (e.g. Burgelman, 2002). The alternative is the organizational decoupling of

exploration and exploitation. The ambidexterity literature claims that firms should continuously house separate organizational units specialized to exploration and exploitation. Drawing heavily on the classic contingency approach in organization theory, the research proposes that an organic organizational form enables effective exploration, while a mechanistic one supports exploitation. For the overall organization, balancing exploration and exploitation crucially depends on having the 'right' mixture of specialized subunits for the two activities.

However, these two approaches to dynamically balancing exploration and exploitation only offer limited insight on how organizations adapt in changing environments. First, the conceptual distinction between the temporal and organizational decoupling of both activities is more apparent than real (Greve, 2007). For instance, Tushman & O'Reilly (1996) and Tushman (2002) suggest that an explorative business unit develops into an organization supporting exploitation after initial successes (see also Duncan, 1974). Burgelman (2002), on the other hand, proposes that the impetus for bursts of exploration emanates from business units that operate outside the traditional domain of a company.

Secondly, and more fundamentally, the organizational search processes underlying the discovery and refinement of capabilities have not been systematically analyzed. During the last decade, our theoretical understanding of organizational search within a given complex task environment has increased substantially (e.g. Levinthal, 1997; Rivkin & Siggelkow, 2007; Levinthal & Posen, 2007). Discontinuous change and organizational search among multiple task environments have received considerable less attention. Our theoretical understanding of how organizations search among as well as within changing task environments is therefore rather limited.

Thirdly, both approaches beg the question how to balance exploration and exploitation. That is, when should firms step up exploration to sense changes in their environments? Should the mixture between 'exploring' and 'exploiting' business units be fixed or flexible? If fixed, what is a right balance? How and when should efforts be redirected to new promising capabilities or product markets? These critical questions remain largely unaddressed. Prior research has been pointed to the role of senior management in dealing with these issues. Senior management needs to be "consistently inconsistent" (Tushman, 2004, p. 81; cf. Smith & Tushman, 2005) to deal with exploration and exploitation – but what is a sufficient level of inconsistency? Our existing theories do not provide an answer.

We develop a formal modelling structure to address these three issues. We draw on and clarify the conceptual distinction between the temporal and the organizational decoupling of exploration and exploitation. We make a critical distinction between organizations that relates to the proclivity to explore of the individual organizational members. An implicit assumption of temporal switching is that an organization may more or less effortlessly change its proclivity to engage in exploration and exploitation. The organizational members supporting temporal switching are assumed to be equally skilled at exploration and exploitation. In a general sense, the organization comprises of generalists, which are capable of moving back and forth between the conflicting demands of exploration and exploitation. On an organizational level, this translates into a structure with flexible task assignment, since each member may choose its own proclivity to explore. This contrasts with the organizational decoupling that assumes that organizational members specialize in either exploration or exploitation. The organization consists of specialists with different, fixed proclivities to explore. This corresponds to an organization with rigid task assignments and

a structure that supports either pure exploration or exploitation. The characteristics and compositions of specialists within the organization determine how the firm searches its environment. Over time, the organization may change the composition of specialists to adapt to environmental changes. Changes in the composition represent the hiring and firing of organizational members or (uncertain) changes in the organizational structure by reassigning tasks.

Our simulation model places these organizations into complex and changing task environments. The organizations face several task environments that differ in performance (exploration). A given task environment is very complex, leaving ample room for refinement search (exploitation). The organizations therefore face the twin challenge of exploration and exploitation. The task environments are subject to discontinuous changes that lead to sharp reversals in their relative performance. Task environments emerge, decline, and re-emerge over time. This captures industry or technology life cycles in a stylized manner.

Our main results suggest that organizations need to be able to systematically adapt its proclivity to explore over time. They need to increase exploration in times of discontinuous change and decrease it during times of tranquillity. However, a modicum of exploration is required at all times to provide the organization with a sensing capability that tracks discontinuous changes. The effectiveness of organizational adaptation may be improved by the firm-wide reallocation of resources to new promising task environments. Coordination helps organizational members in breaking free from the past. Our results show that coordination could be a way to overcome the shadow of prior history, since it takes time until organizational members update their expectations about the relative performance advantages of task environments.

Our paper is structured as follows. In section 2, we present the nuts and bolts of our simulation model. Section 3 presents our result. In section 4, we discuss the results and point to avenues for further research.

3. The model

Kauffman's (1993) NK model has been widely used in the study of organizational search processes (e.g. Levinthal, 1997; Sørensen, 2000; Gavetti, 2005; Knudsen & Levinthal, 2007). We use a variant of the model and extend it to study ambidextrous and non-ambidextrous organizations in changing task environments. Broadly, an organization comprises of a fixed number of agents with differing proclivities for exploration and exploitation. In line with the essential features of ambidextrous organizations as described in the previous section, the organization's degree of ambidexterity is reflected in the composition of agents. The organization faces the challenge of identifying a promising task environment (exploration) and finding a high-performing configuration of activities within a given task environment (exploitation). Thus, an organization searches among as well as within task environments. Since a core proposition of the ambidexterity literature is the relative superiority of ambidextrous organization in the face of environmental change, we model changing task environments. Over time, the organization may adapt its composition of agents (with fixed or adjusting proclivities for exploration and exploitation), thereby influencing the balance between exploration and exploitation and its degree of ambidexterity. In addition, an organization may reassign agents to task environments deemed to be the most promising through organization-wide coordination.

Managing radical and incremental change: Organizational search in changing environments

Our model consists of four major building blocks. The first establishes the nature and change of task environments. The second building block specifies how agents search among and within task environments; the third establishes and modifies the composition of agents within organizations. The fourth building block models how organizations may reassign agents to task environments.

The nature and change of task environments

We generate a set of performance landscapes that represent separate task environments. Task environments differ in average performance and the average performance of task environments is subject to change over time.

In each task environment, a choice set consists of N binary attributes. These may represent activities within a task or the policy attributes of a business unit relating to sourcing, production, sales, support function, etc. Each attribute can take on two states, so there are 2^N different task configurations in each landscape. The performance landscape created by the model is a mapping of the set of attributes onto performance values. The performance values of each of the N attributes are determined by random draws from a uniform distribution over the unit interval.¹ The overall performance of a configuration in one landscape is the average of the values assigned to each of the N attributes.

The attributes of an organizational configuration may be more or less interdependent. Attributes are interdependent if the value of each of the N individual attributes depend on both the state of that attribute itself and the states of K other attributes. If $K = 0$, attributes are

¹ Since our studies include rather large performance landscapes ($N= 50$), we report raw values from the NK runs rather than normalized results.

independent. As K increases, more and more attributes of a configuration become interdependent, with $K = N - 1$ being the case of interdependence among all attributes. The number of interdependencies given by K determines the surface of the performance landscape. With higher values of K , there are more local peaks, and performance differences among neighboring configurations differing only in a single attribute become relatively more pronounced.

We do not simulate changes in the task environments by respecifying the task environment (e.g. Ethiraj & Levinthal, 2004). Rather, we reset the average performance of the landscape by multiplying the performance of a configuration by a set parameter. The advantage of our approach is that we capture the evolution of task environments in a stylized manner. Task environments become relatively more or less attractive during the simulation run, reflecting the emergence and decline of products, industries, or technologies, a key concern in the ambidexterity literature (e.g. Tushman, 2002). Table 1 shows the evolution of task environments over time. In the beginning, task environment 1 is superior by a large margin. Then, after the simulation has run for 50 time steps, the performance in the task environment is heavily discounted (by a factor of 0.01), while performance in task environment 2 is multiplied by a factor of 1, and so on. Note that task environment 4 is inferior in each and every time step, so agents should stay away from it at all times.²

Task environment	1	2	3	4

² We set the number of time steps between the emergence and the decline of task environment to 50. This corresponds to earlier treatments in the literature (e.g. Levinthal, 1997), since it usually takes less than 50 time steps for myopic agents to find a local peak, with no additional performance increases in subsequent time steps. Note also that limiting the ‘golden age’ of a landscape to 50 time steps gives a relative advantage to agents that settle in a superior landscape early on, since they profit longer than agents still exploring other landscapes.

Up to time step...				
50	1	0.01	0.01	0.01
100	0.01	1	0.01	0.01
150	0.01	0.01	1	0.01
175	1	0.01	0.01	0.01
200	1	1	0.01	0.01
225	0.01	1	1	0.01
250	0.01	0.01	1	0.01

Table 1: Evolution of task environments over time

Search among and within task environments

We characterize agents as boundedly rational, yet capable of some level of forward deliberation by forming expectations or beliefs about the different task environments (e.g. Gavetti & Levinthal, 2000).

Within task environments, agents refine their configurations by engaging in local search. We follow prior research in our conception of organizational search within task environments. Agents revise policy attributes by flipping a randomly chosen single bit and examining the outcome. The revision is implemented only if the result increases performance, otherwise, it gets rejected.³ Search within task environments corresponds to exploitation, as agents refine the configuration. Boundary-spanning search (Narr & Rosenkopf, 2001), or exploration, is captured by organizational search among task environments.

Agents form expectations or beliefs about the relative performance of task environments by engaging in search among task environments. An agent forms beliefs on the basis of knowledge generated through prior search processes. To model this relationship between past search efforts and the formation of beliefs, we turn to research on reinforcement

³ We limit our study to perfect evaluation and do not consider the elaborations introduced by Knudsen & Levinthal (2007) relating to imperfect evaluation of proposals.

learning. The Softmax algorithm, attributed to Luce (1959), provides a straight-forward way to model the formation of beliefs of agents (Vermorel & Mohri, 2004). The Softmax algorithm makes the choice of a performance landscape at time step t dependent on the observed mean performance \bar{x}_i of the task environments d and on the proclivity for exploration τ :

$$(1) p(t) = e^{\bar{x}_i / \tau} / \sum_{i=1}^d e^{\bar{x}_i / \tau}$$

The parameter τ , called the “temperature”, influences the degree to which an agent adheres to prior beliefs. A lower value of τ increases the probability that an agent remains within the performance landscape chosen before, as the most attractive one. Hence, it places a higher emphasis on exploitation. A higher τ downplays the role of past search processes and increases the probability that an agent samples and explores a different task environment. An agent samples a task environment by engaging in local search within the landscape. Note that agents continually update the estimates of mean performances. The more an agent explores task environments, the better the estimates of mean performances. The temperature τ therefore encapsulates the balance between exploration among landscapes and exploitation within landscapes.

The tendency of an organization to engage in exploration and exploitation crucially depends on how much emphasis agents place on prior knowledge. Ambidextrous and non-ambidextrous organizations thus differ in the composition of agents with different proclivities to engage in both types of search. Ambidextrous organizations that are equally good at exploration and exploitation will be characterized by more agents with both very high τ and very low τ , while non-ambidextrous organizations tend to have a majority of

agents with medium τ . Thus, the composition of agents within the organization matters for the level of ambidexterity.

Ambidexterity: The composition of agents within an organization

An organization consists of 100 agents. We model structural ambidexterity as an organization populated by specialists with fixed properties that are either very good at exploitation or exploration. The agents therefore have clearly defined roles. The non-ambidextrous organization consists of agents that have a medium proclivity to explore. The agents therefore have clearly defined roles that they may not change on their own. In contrast, contextual ambidexterity is represented by generalists that can easily switch between exploration and exploitation. That is, their roles are much more flexible.

Structural ambidexterity. At the start of the simulation, each agent's temperature is randomly determined by a draw from a continuous Beta distribution defined on the interval [0.1, 10]. Ambidextrous and non-ambidextrous organizations draw from differently shaped Beta distribution. Ambidextrous organizations tend to be populated by pure explorers (very high temperature) or pure exploiters (very low temperature), while non-ambidextrous organizations have a higher probability of drawing average temperatures. The two non-negative shape parameters, α and β , determine the shape of the probability distribution. Table 2 shows the different values for α , β and its interpretation. Generally, higher α , β make the organizations less ambidextrous, while $\alpha > \beta$ skews the distribution towards exploration, and vice versa.

Beta distribution on the interval [0.1, 10]	α	β	Mean temperature	Interpretation
Ambidextrous organization	0.5	0.5	5	Organization attracts agents with very high or very low temperature equally.
Ambidextrous organization skewed towards exploitation	0.5	5	1	Organization attracts more agents with very low temperature and few with very high temperature.
Ambidextrous organization skewed towards exploration	5	0.5	9	Organization attracts more agents with very high temperature and some with very low temperature.
Non-ambidextrous organization	5	5	5	Organization mainly attracts agents with average temperature.

Table 2: Beta distribution and forms of structural ambidexterity

As a baseline, we keep the composition of the organizations fixed over time. That is, the organization does not attempt to adapt the composition of agents to the environment. However, more realistically, we model an organizational policy to changing the proclivity to explore (“temperature”) of agents over time. The organization may fire the least-performing agents and immediately replace them with new agents drawn from the respective probability distribution. The replacements get assigned to the position formerly occupied by the fired agent. In essence, the replacement takes over where the old agent has stopped. This could also represent the retraining of the agent to change the proclivity to explore. We thus assume that an organization sticks to a fixed personnel policy when it comes to hiring or the training of employees. The fixed policy is represented by the probability distribution the organization draws from. In each time step, the fraction of agents an organization replaces is endogenously determined. An agent gets replaced if the current performance falls below 0.5, the average expected performance in an attractive industry.

Contextual ambidexterity. The model is contrasted with the endogenous change of the individual agent's proclivity to explore. That is, an agent may freely adapt its proclivity to explore to performance feedback. This corresponds to contextual ambidexterity, in which individuals choose the level of exploration and exploitation on their own. They are generalists in the sense that they may seamlessly switch from pure exploration to pure exploitation, and vice versa. Inspired by the literature on aspiration-level organizational learning (e.g. March, 1988; Greve, 2003), we depict the temperature τ of an agent as determined by the agent's current performance $\pi(t)$ and an exogenous parameter δ :

$$(2) \tau_{t+1} = \delta/\pi(t).$$

An increase in performance results in a decline of the temperature τ and a toning down of the exploration. Performance decreases, on the other hand, again amplify the tendency to explore. Thus, the aspiration-level model captures the dynamic adjustment of an agent's temperature and makes the organization as a whole more adaptable as agents respond to performance feedback. Throughout the simulation, we set δ equal to 0.05. This gives a reasonable range for the agents' proclivity for exploration. With an average performance of 0.5 in an attractive landscape, the temperature goes down to 0.1, allowing only for a very modest level of exploitation. In a declining task environment the average performance drops to 0.005, increasing the temperature to 10. Thus, the average lower and upper bounds of the temperature in the dynamic adjustment model corresponds to the chosen interval of the Beta distribution for structural ambidexterity.

Coordination and reassignment

Agents choose task environments based on the individual assessments of their relative performance and the proclivity to engage in exploration. However, organizations are characterized by the transfer of knowledge between organizational members and the coordination of their behavior. To capture this critical property of an organization, we also allow for the coordinated reassignment of agents to task environments. Specifically, an organization sends more agents to more attractive task environments, and vice versa. The overall attractiveness of a task environment is the mean performance of all agents within the task environment. The agents are reassigned to task environments based on a comparison of the mean performance X_i of a task environment and the mean performance X_d of all task environments d :

$$(3) F(i) = \frac{X_i}{\sum_1^d X_d}.$$

The frequency $F(i)$ determines the percentage of all agents that get randomly assigned to the task environment i . They then engage in local search within the assigned task environment.

Note that the organization-wide reassignment complements the individual choice of a task environment by the agent. Agents with a high proclivity to explore (high temperature) may abandon the task environment they got assigned to and sample a different task environment in next time step. This could reveal useful information for reassignment by updating the mean performances of task environments. Agents with a very low tendency to engage in exploration (very low temperature), on the other hand, may take advantage of organization-wide coordination by being sent to more attractive task environments. Without organization-wide coordination, these agents may get stuck in low-performing task environments forever, since they do not explore on their own.

However, inert organizations may be reluctant or even unable to reallocate agents to task environments with uncertain potential. We reflect this in our model by allowing for different degrees of coordination in the reassignment of agents. Agents get allocated to a new landscape with a probability between 0 and 1, representing increasing levels or effectiveness of organizational coordination.

3. Results

In our model, organizations face the dual challenge of exploration and exploitation in a dynamic setting. They have to identify an attractive task environment and then proceed to refine a configuration of subtasks. The task environments change over time, so organizations need to balance exploration and exploitation. Organizations therefore have to master organizational search on two levels. Exploration refers to search among task environments and exploitation to the refinement within a particular environment.

We have restricted our analysis to four task environments with $N=50$ attributes and $K=0$ and $K=49$ interactions between them. The task environments change their relative performance over time (see table 1). We report results for organizations consisting of 100 agents. Our results reported here show the average of 30 simulation runs with 250 time steps.

We begin by comparing different types of structural ambidexterity and their performance properties in huge, but simple landscapes ($K=0$). We then contrast structural ambidexterity with the aspiration-level adaptation represented by contextual ambidexterity. Next, we compare the performance in simple and complex task environments ($K=49$). Lastly, we analyze the impact of coordination on the performance of ambidextrous organizations.

Levels of structural ambidexterity

Ambidextrous and non-ambidextrous organizations differ in the composition and recruitment of agents. An ambidextrous organization tends to recruit more extreme types that are either good at exploration (high proclivity to explore) or exploitation (low proclivity to explore). Non-ambidextrous organizations, on the other hand, hire agents with average properties. This is represented in the model by the Beta distribution that an organization draws and replaces its agent from.

Figure 1 compares the average performance of the various levels of ambidexterity over the run of the simulation (250 time steps) in large, simple landscapes (N=50, K=0).⁴

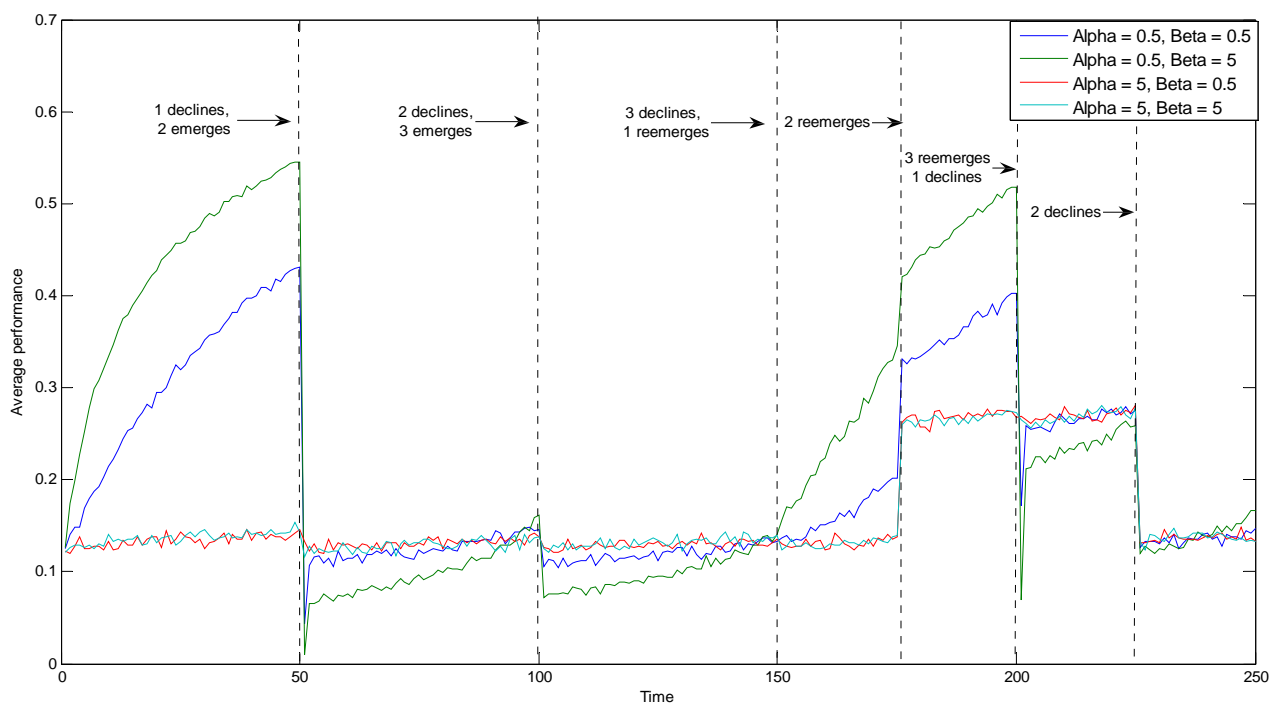


Figure 1: Average performance of different types of ambidexterity (N=50, K=0)

⁴ To test the reliability of our results, we also ran an organization with $\alpha = 10, \beta = 10$. The performance properties are essentially the same as with $\alpha = 5, \beta = 5$. The results are available from the authors upon request.

Managing radical and incremental change: Organizational search in changing environments

The figure shows the evolution of the task environments over time and its impact on the average performance of organizations. Remarkable is the average performance of the ambidextrous organization skewed toward exploitation ($\alpha = 0.5$, $\beta = 5$). The organization comprises a majority of agents with a very low temperature, with the average organizational temperature being 1. The agents find the most attractive task environment quickly and then proceed to refine the configuration within the task environment. After 50 time steps, the environment sharply decreases in performance. Adaptation to the new more attractive landscape is much slower, however, and agents are still engaged in exploring the task environments when the environment changes again. Thus, the drop in agents' performance in time step 100, when the landscapes change again, is not as pronounced as before. After 150 time steps, task environment 1 rebounds, initiating a new round of exploration that quickly zooms in on the new superior landscape. Since most agents have located in task environment 1, average performance declines sharply in time step 200.

The ambidextrous organization ($\alpha = 0.5$, $\beta = 0.5$) is also doing quite well. While this type adapts more quickly to changes of the task environments, it is outperformed when it comes to refining the configurations within a chosen task environment. The reason for this is that the ambidextrous organization consists of more explorers with a very high temperature, with the average organizational temperature being 5. Hence, the organization spends too much time and resources searching among task environments.

Both types of organizations clearly outperform the ambidextrous organization skewed towards exploration ($\alpha = 5$, $\beta = 0.5$) as well as the non-ambidextrous organization with "average" employees. The latter two engage in excessive exploration, effectively falling prey

Managing radical and incremental change: Organizational search in changing environments to the failure trap (Levinthal & March, 1993). Thus, they devote not enough resources to the exploitation of configurations within a task environment.

Interestingly, all types of organizations take advantage of the reemergence of the first task environment in time step 150, albeit to different degrees. This reveals a long shadow of history. In the model, all agents form beliefs about the attractiveness of task environment. This is captured by the probabilities in the Softmax algorithm that effectively stores prior experience. During the first 50 time steps, task environment 1 is increasingly favored by all agents. This impression is slowly modified rather than completely erased when the task environment declines. Thus, when the task environment reemerges, agents start with a more favorable impression. This also explains why the two successful types of organizations have no trouble finding the attractive task environment in the beginning (where they have no prior beliefs whatsoever). This feature of the model neatly captures the organizational inertia created by prior experience that an ambidextrous organization has to overcome. Moreover, agents adapt more quickly to opportunities when they have some prior, positive experience with it, while they tend to be more reluctant if they had experienced negative feedback in the past.

What are the reasons for the remarkable differences between the different types of organizations? Figure 2 shows the evolution of the average proclivity to explore (the “temperature”) for the four types of organizations over time.




Figure 2: Average proclivity to explore (“temperature”) over time (N=50, K=0)

Only the pure ambidextrous organization ($\alpha = 0.5, \beta = 0.5$) and the ambidextrous organization skewed towards exploitation ($\alpha = 0.5, \beta = 5$) consistently adapt to changing circumstances and shift the balance between exploration and exploitation. The reason is the different hiring or training policies of organizations. The pure ambidextrous organization has an equal probability of recruiting pure explorers or pure exploiters, while intermediate temperatures have a very low probability. Therefore, initially, unlucky exploiters that got trapped in unattractive landscapes get fired and are replaced. After a short burst of exploration, an increasing fraction of underperforming explorers gets fired, since they still wander around between task environments. They get replaced by exploiters that quickly locate in the attractive landscape, because they adhere to prior knowledge and go to the most attractive landscape. They thereby slowly drive out the more adventurous explorers. Hence, the average organizational temperature declines over time. Then, with the change of landscapes in time step 50, the formerly successful exploiters suddenly find themselves disadvantaged and they are replaced by new agents that once again sharply increase the temperature, leading to higher adaptability of the organization as a whole. The same

Managing radical and incremental change: Organizational search in changing environments

mechanism is at work in the ambidextrous organization skewed towards exploitation. Since this type has a higher probability of drawing pure exploiters, the organization is more adept at switching to exploitation, but it takes longer to identify new attractive landscapes, because the new exploiters get assigned to old positions within unattractive landscapes. Thus, the adaptation to a new attractive landscape takes longer than in the pure ambidextrous organization.

The other two types fail to adapt, since they do not systematically replace agents. The ambidextrous organization skewed toward exploration ($\alpha = 5$, $\beta = 0.5$) does not attract enough exploiters within the lifespan of a task environment. Unsuccessful explorers get replaced by explorers that eventually also fail, and so on. The average temperature essentially remains the same during the lifespan of a task environment and the organization fails to adapt its tendency to explore to lower levels. This would be necessary to refine the configuration within a chosen task environment. The same logic applies to the non-ambidextrous organization ($\alpha = 5$, $\beta = 5$). It does not hire or train enough pure exploiters in times of stability and attracts too few pure explorers in times of change.

Structural and contextual ambidexterity compared

In the next step, we compare the structurally ambidextrous organizations with the dynamic adjustment model that captures contextual ambidexterity.⁵ In that model, agents adapt the individual temperature to changes in the current fitness. Figure 3 compares the cumulated performance at the end of the simulation run of the ambidextrous organizations with the dynamic adjustment model. We do not report the cumulated performance of the non-

⁵ Throughout the following we refer to contextual ambidexterity as the dynamic adjustment model to clearly set it apart from the structurally ambidexteritous organizations.

ambidextrous organization, since it corresponds to the wealth profile of the ambidextrous organization skewed toward exploration ($\alpha = 5, \beta = 0.5$).




Figure 3: Wealth distributions of structural and contextual ambidexterity (N=50, K=0)

The dynamic adjustment model outperforms the pure ambidextrous organization ($\alpha = 0.5, \beta = 0.5$) as well as the exploration-skewed organization ($\alpha = 5, \beta = 0.5$). It is only outperformed by the ambidextrous organization skewed towards exploitation ($\alpha = 0.5, \beta = 5$). The dynamic adjustment model essentially envisions agents that can easily and seamlessly switch between extreme exploration and exploitation, a feature stressed by the contextual ambidexterity literature. Agents therefore radically increase the temperature when performance declines sharply and they take it down when performance again improves. Figure 4 again looks at the different profiles of organizational adaptation over the simulation run. The aspiration-level model is more skillful in finding an attractive landscape in the beginning, as agents stumble upon the most attractive landscape and then sharply decrease

their temperature. After the decline of the task environment, the agents reignite exploration, effectively mimicking the explorers in the ambidextrous organization.

The dynamic adjustment model, however, gets outperformed by the ambidextrous organization when the first task environment reemerges. The reason for this is again the long shadow of history, since agents still wander around landscapes. The pure exploiters in the ambidextrous organization skewed towards exploitation rely more heavily on prior knowledge and quickly relocate to the reemerging landscape.




Figure 4: Comparison of structural and contextual ambidexterity (N=50, K=0)

Our analysis so far suggests that an ambidextrous organization needs a mechanism to dynamically balance exploration and exploitation over time. A fixed balance will not do. The pure ambidextrous organization achieves adaptation of exploration and exploitation in a dynamic setting by discriminating against one of the two extreme types when it comes to

replacing agents. The same holds true for the ambidextrous organization skewed towards exploitation. It outperforms all other organizations by placing a premium on exploitation, while still allowing for some valuable exploration among landscapes to update the information exploiters then act upon. The dynamic adjustment model, on the other hand, achieves organizational adaptation by radically altering the agents' individual temperature based on performance feedback. The other types of organizations do not systematically adapt to changing circumstances, since they replace agents with just more of the same. Thus, they fail to dynamically balance exploration and exploitation. Diversity of organizational members matters if agents are unable to switch between the requirements of exploration and exploitation .

Simple versus complex landscapes

So far, we have only considered the case of huge, but simple performance landscapes ($K=0$). We now turn to the analysis of very complex landscapes and their effect on performance. The setting is the same as before, with the only difference that the task environments now have pervasive interactions between attributes. Generally, the landscapes become more 'rugged', with many local peaks. Figure 5 shows the cumulated performance for the ambidextrous organization and for the dynamic adjustment model in simple and in complex task environments. We have omitted the remaining three organization types, since they show the same result: The somewhat surprising observation is that these organizations in large, complex landscapes are more successful in terms of average performance than in large and simple landscapes. That is, dynamically balancing exploration and exploitation is actually easier to accomplish in large and very complex environments. How can that be?




Figure 5: Cumulated performance distribution in simple and in complex task environments (N=50, K=0/49)

More complex task environments ($K > 0$) are more ‘rugged’, with many local peaks and performance differences among neighboring configurations differing only in a single attribute becoming relatively more pronounced. This has two important ramifications. First, due to steeper gradients, initial performance increases tend to be greater in more rugged landscape. Second, the local peaks act as attractors (Levinthal, 1997), trapping agents on them. In simple landscapes, on the other hand, agents progress slowly towards the global optimum, which is well past their time horizon. In our model, agents have to negotiate an immensely huge space of 2^{50} possible configurations in each landscape, with an overall time horizon of only 250 time steps. Effectively, they actually have only 100 time steps until a task environment declines in average performance. Therefore, in large and complex landscapes, the danger of getting trapped on a local peak is by far outweighed by the early advantages of sizable performance increases. In other words, agents do better by

Managing radical and incremental change: Organizational search in changing environments

implementing rapid improvements early on, leading to large performance gains, and then by sticking to what works, instead of forever chasing the elusive holy grail of a global optimum. This is essentially what is happening here. If one compares the agents' marginal gains during the first 50 time steps, performance increases much faster in complex landscapes. While the performance gains tend to level out after a while, it takes agents more than 50 time steps to achieve the same average performance in the simple landscape. Eventually, the simple landscapes would surely outperform the agents in the complex task environments, but, with changing landscapes, agents never come around to realize the whole potential.

Coordination and organizational reallocation of agents

Coordination has a marked effect on the average performance of all the types of organizations considered here. Coordination captures higher-level intervention that sends agents to task environments based on their relative average performance. Figure 6 shows the average cumulated performance for different levels of coordination. With increasing levels of coordination, the performance of all the organizational types sharply increases and converges, blurring the differences between types. We find the same results for large and complex landscapes (N=50, K=49).




Figure 6: Levels of coordination and average cumulated performance (N = 50, K = 0)

With coordination, all organizations become more competent in identifying an attractive landscape, since only the current average performance matters. This contrasts with adaptation and selection of task environments by the individual agents. They base their choices of task environments also on the cumulated mean performance. Past experiences with task environments still influence today’s choices. This type of path-dependency is absent in organization-level coordination. So, even if just one agent stumbles upon an attractive landscape with a high average performance, it sends a strong signal to the organization as a whole, regardless of the agents’ individual beliefs about the task environments. This also implies that at least a modicum of explorative activity is needed to sense changes in the task environment. However, coordination does not have much bite if it is not complemented by an effective reallocation device that reassigns agents to task

Managing radical and incremental change: Organizational search in changing environments environments. Hence, with an increasing effectiveness of reallocation, organizations adapt more quickly to changes in the task environments.

This also highlights that a sensing technology – the ongoing exploration of new task environments – needs to be matched with an effective reallocation device that channels exploitative search efforts into new promising ventures. Our results therefore reinforce and differentiate a critical point in research on ambidexterity, the role of senior management (Tushman, O'Reilly, 2004; Smith, Tushman, 2005). Senior management's role in making ambidexterity effective goes well beyond allowing for exploration in an organization. Without some level of exploration, an organization fails to identify new promising task environments. At least as critical is senior management's ability to reallocate resources to new promising ventures. This does not only reinforce success, but even more so supports the eventual transition from mere discovery (exploration) to refinement (exploitation) in the new venture. For this reason, coordination and the reallocation of resources play a critical part in organizational adaptation by strengthening the effectiveness of both exploration and exploitation.

4. Discussion and conclusion

Our results point to a number of interesting propositions and avenues for further research on the ambidextrous organization and beyond. To cope with the twin challenge of incremental and radical environmental change, organizations must be able to dynamically adjust the balance between exploration and exploitation. Exploration relates to the discovery of new attractive environments; exploitation to the refinement of configurations within a task environment. Organizations need to search among as well as within task environments.

Managing radical and incremental change: Organizational search in changing environments

A fixed balance between exploration and exploitation, as represented by the non-ambidextrous organization, will not do. It fails to secure the gains from incremental change by exploitation, while not providing for enough experimentation when organizations have to adapt to entirely new organizations. Likewise, a structurally ambidextrous organization geared towards exploration, foregoes the performance benefits of refinement search by exploitation. The other two types of structural ambidexterity, on the other hand, achieve organizational adaptation by systematically creating and adapting the diversity of specialists within its boundaries. In times of rapid change pure exploiters get replaced with employees specializing in exploration, allowing for the discovery of new attractive ventures. In structural terms, the positions get redefined toward a higher disposition for exploration. Likewise, after the discovery of an attractive opportunity, exploiters drive out explorers, responding to the need for exploitation and initiating a transition towards a prolonged period of refinement. However, the remaining explorers still allow the organization to sense changes in the task environments. This not only points to the critical role of organizational structures in bringing about ambidexterity, but even more so stresses the role of an 'ambidextrous' personnel policy in hiring, firing, and training of employees. Structural ambidexterity is enabled by diversity among specialists and its systematic adaptation of it over time. Success depends on the commitment to a fixed policy that nurtures and supports diversity among employees and the policy needs to be matched to the ambidextrous organizational structure that provides the positions for specialists.

Contextual ambidexterity relies on highly skilled generalists to achieve a dynamic balance between exploration and exploitation. A generalist in a contextually ambidextrous organization is an individual that may seamlessly move between the contrasting

requirements of exploration and exploitation. The balance is dynamically adjusted by performance feedback. Experimentation therefore drives the dynamic balance between exploration and exploitation. This view of ambidexterity closely corresponds to the evolutionary perspective of backward looking experimentation and adaptation. However, it also requires employees that are capable of skillfully making a radical transition from the strict adherence to prior knowledge to breaking with the past. Structural ambidexterity requires a team of specialists to make this transition instead of highly skilled individuals (Greve, 2006). Thus, there may be different paths to the goal of ambidexterity.

Our results also stress the importance of sensing technologies and reallocation devices for achieving ambidexterity. An organization must be able to sense changes in the environment and respond by reallocating employees to new attractive opportunities. The sensing technology is encapsulated by having an on-going modicum of experimentation and exploration, even in times of stability. This reinforces the importance of organizational slack in the discovery of new opportunities. Organizational adaptation to new ventures may be reinforced by the organization-wide coordination and reallocation of resources. Instead of exclusively relying on individual level adaptation, the performance of both structural and contextual ambidexterity may be increased by firm-level coordination. The sensing technology creates a signal that is made useful for exploiters by their reallocation to more attractive ventures. Since coordination by senior management is considered to be a property of structural ambidexterity, this finding supports the complementarity of structural and contextual ambidexterity. Coordination may speed up and reinforce individual adaptation. Once again, for coordination, performance feedback is critical, since coordination solely depends on current relative performance. As depicted in our model, coordination is,

essentially, 'memory-free'. Our results show that coordination could be a way to overcome the long shadow of prior history, since it takes time until agents have updated their expectations or beliefs about the relative advantages of task environments. Coordination helps them in breaking free from the past. The results of our model also stress the importance of unlearning, since adaptation would be more smooth if agents would not only adhere less to what they have learned, but reset their memory when the environment changes radically.

However, we paint a very stylized and positive picture of organizational coordination in our model. The performance differences between environments are large and non-ambiguous. In addition, senior management perfectly evaluates the performances in different environments and may swiftly act upon it (e.g. Knudsen and Levinthal, 2007). This shows that the critical role of organizational coordination and the reallocation of effort by senior management need more attention in future research.

5. References

- Adler, P., B. Borys. 1996. Two types of bureaucracy: Enabling and coercive. *Adm. Sc. Quart.* **41**(1) 61-89.
- Adler, P., B. Goldoftas, D.I. Levine. 1999. Flexibility Versus Efficiency? A Case Study of Model Changovers in the Toyota Production System. *Org. Sc.* **10**(1) 43-68.
- Benner, M. J., M. L. Tushman. 2003. Exploitation, exploration, and process management: The productivity dilemma revisited. *Acad. Man. Rev.* **28**(2) 238-256.
- Birkinshaw J., C. Gibson. 2004a. The antecedents, consequences, and mediating role of organizational ambidexterity. *Acad. Man. J.* **47**(2) 209-226.
- Birkinshaw J., C. Gibson. 2004b. Building ambidexterity into an organization. *MIT Sloan Man. Rev.* **45**(4) 47-55.
- Burns, T., G.M. Stalker. 1961. *The Management of Innovation*. Tavistock, London.
- Cyert, R.M., J.G. March. 1963. *A Behavioral Theory of the Firm*. Blackwell, Oxford.
- Duncan, R. B. 1976. The ambidextrous organization: designing dual structures for innovation. Kilmann, R.H., L.R. Pondy, D. Slevin, eds. *The management of organization*. North-Holland, New York, 167-188.
- Ethiraj, S.K., D. Levinthal. 2004. Modularity and Innovation in Complex Systems. *Man. Sc.* **50**(2) 159-173.
- Galbraith, J. 1982. Designing the innovating organization. *Org. Dyn.* Winter 5-25.
- Gavetti, G. 2005. Cognition and Hierarchy: Rethinking the Microfoundations of Capabilities' Development. *Org. Sc.* **16**(6) 599-617.
- Gavetti, G., D. Levinthal. 2000. Looking Forward and Looking Backward: Cognitive and Experiential Search. *Adm. Sc. Quart.* **45**(1) 113-137.
- Gibson C., J. Birkinshaw. 2004. The antecedents, consequences, and mediating role of organizational ambidexterity. *Acad. Man. J.* **47**(2) 209-226.
- Greve, H. 2003. *Organizational Learning from Performance Feedback – A Behavioral Perspective on Innovation and Change*. Cambridge University Press, Cambridge.
- Greve, H.R. (2007), Exploration and exploitation in product innovation, *Industrial and Corporate Change* (forthcoming).

- He Z.L., P.K. Wong. 2004. Exploration vs. exploitation: An empirical test of the ambidexterity hypothesis. *Org. Sc.* **15**(4) 481-494.
- Kauffman, S.A. 1993. *The origins of order*. Oxford University Press, New York.
- Knudsen, T., D.A. Levinthal. 2007. Two faces of search: Alternative generation and alternative evaluation, *Org. Sc.* **18**(1), 39-54.
- Lawrence P., J. Lorsch. 1967. *Organization and Environment: Managing differentiation and integration*. Harvard University, Boston, MA.
- Levinthal, D.A. 1997. Adaptation on rugged landscapes, *Man. Sc.* **43**(7) 934-951.
- Levinthal, D.A., J.G. March. 1993. The myopia of learning, *Strat. Man. J.* **14**(8) 95-112.
- Lubatkin, M.H., Z. Simsek, Y. Ling, J.F. Veiga. 2006. Ambidexterity and Performance in Small-to Medium-Sized Firms: The Pivotal Role of Top Management Team Behavior. *J. Man.* **32**(5) 646-672.
- Luce, R.D. 1959. *Individual Choice Behaviour*. Wiley, New York.
- March, J.G. 1991. Exploration and exploitation in organizational learning. *Org. Sc.* **2**(1) 71-87.
- Miles, R.E., C. Snow. 1978. *Organizational strategy, structure and process*. McGraw-Hill, New York.
- Miner, A.S. 1991. Organizational Evolution and the Social Ecology of Jobs. *Am. Soc. Rev.* **56**(6) 772-785.
- O'Reilly, C., M.L. Tushman. 2004. The ambidextrous organisation. *Harvard Bus. Rev.* **82**(4) 74-81.
- Rosenkopf, L., A. Nerkar. 2001. Beyond local search: Boundary-spanning, exploration, and impact in the optical disc industry. *Strat. Man. J.* **22**(4) 287-307.
- Smith W.K., M.L. Tushman. 2005. Managing strategic contradictions: A top management model for managing innovation streams. *Org. Sc.*, **16**(5) 522-536.
- Sorenson, O. 2002. Interorganizational Complexity and Computation. Baum, J.A.C., ed. *Companion to Organizations*. Blackwell, Oxford. 664-685
- Tushman, M.L., C. O'Reilly. 1999. Building Ambidextrous Organizations. *Health Forum J.*, March/April 20-23.
- Tushman, M.L., C. O'Reilly. 1996. Ambidextrous organizations: Managing evolutionary and revolutionary change. *Cal. Man. Rev.* **38**(4) 8-30.

Managing radical and incremental change: Organizational search in changing environments

Tushman, M.L., W.K. Smith. 2002. Organizational Technology. Baum, J.A.C., ed. Companion to Organizations. Blackwell, Oxford. 387-414

Van Looy, B., T. Martens, K. Debackere. 2005. Organizing for Continuous Innovation: On the Sustainability of Ambidextrous Organizations. *Creat. & Innov. Man.* **14**(3) 208-221.

Vermorel, J., M. Mohri. 2004. Multi-Armed Bandit Algorithms and Empirical Evaluation Paper presented at *European Conf. Machine Learning*, 2004.