

Looking Forward and
Looking Backward:
Cognitive and
Experiential Search

Giovanni Gavetti
Daniel Levinthal

University of Pennsylvania

We used computer simulations to examine the role and interrelationship between search processes that are forward-looking, based on actors' cognitive map of action-outcome linkages, and those that are backward-looking, or experience based. Cognition was modeled as a simple, low-dimensional representation of a more complex, higher dimensional fitness landscape. Results show that, although crude, these representations still act as a powerful guide to initial search efforts and usefully constrain the direction of subsequent experiential search. Changing a cognitive representation itself can act as an important mode of adaptation, effectively resulting in the sequential allocation of attention to different facets of the environment. This virtue of shifting cognitive representation, however, may be offset by the loss of tacit knowledge associated with the prior cognition. ●

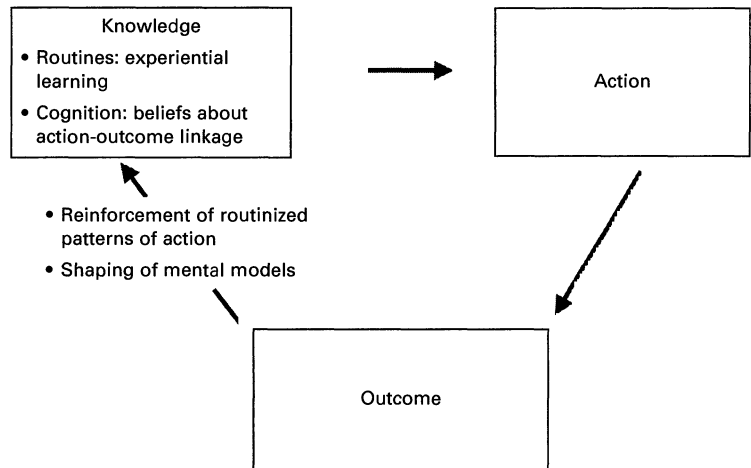
The notion of bounded rationality (Simon, 1955) has been a cornerstone of organizational research (March and Simon, 1958; Cyert and March, 1963) and a basis for two distinct intellectual lineages. One is a perspective focusing on organizational learning (Levitt and March, 1988), especially ideas of local search (Cyert and March, 1963) and the evolution of relatively stable organizational routines (Nelson and Winter, 1982). Such routines reflect experiential wisdom in that they are the outcome of trial and error learning and the selection and retention of prior behaviors. Although bounded rationality highlights the importance of information-processing constraints, as reflected in the role of organizational routines and standard operating procedures (March and Simon, 1958; Cyert and March, 1963), it does not negate the possibility of action based on a logic of consequences (March, 1994). Indeed, the notion of bounded rationality has helped spawn a second research tradition that focuses on individuals as explicitly considering the possible consequences of the choices they make (March and Simon, 1958; Simon, 1991). In this tradition, bounded rationality is manifest primarily in the limited or imperfect cognitive representations that actors use to form mental models of their environment (Thagard, 1996). Such representations both simplify the complexity of spatial relationships (Porac, Thomas, and Baden-Fuller, 1989), the interaction among choices and actors at a point in time, and temporal or causal relationships (Weick, 1979). Cognitive representations have been shown to be a critical determinant of managerial choice and action (Tversky and Kahneman, 1986; Huff, 1990; Fiol and Huff, 1992; Walsh, 1995); in particular, a firm's choice of strategy is often a by-product of actors' representation of their problem space (Simon, 1991).

In terms of figure 1, cognitive and experiential based logics of choice can be distinguished as follows. Cognition is a forward-looking form of intelligence that is premised on an actor's beliefs about the linkage between the choice of actions and the subsequent impact of those actions on outcomes. Such beliefs derive from the actor's mental model of the world (Holland et al., 1986). Greater fidelity between the mental model of action-outcome linkages presumably leads to more efficacious choices of action.

©2000 by Cornell University.
0001-8392/00/4501-0113/\$3.00.

●
Research support from the Sol Snider Center for Entrepreneurship at the Wharton School, University of Pennsylvania and the Division of Research, Harvard Business School is gratefully acknowledged. We have benefited from the comments of seminar participants at the Dallas ORSA/TIMS Meeting; Graduate School of Business, University of Chicago; Sloan School of Management, MIT; Anderson School of Management, UCLA; the Wharton School, University of Pennsylvania; University of Wisconsin, and Columbia University. We thank John Lafkas and Sid Winter for their comments on the prior draft and the Associate Editor Rod Kramer and two anonymous reviewers who played an important role in the development of the paper by constructively prodding us to clarify our arguments. Finally, we thank Linda Johanson for her thoughtful editorial help.

Figure 1. Intelligence of action.



In contrast, experiential wisdom accumulates as a result of positive and negative reinforcement of prior choices (Levitt and March, 1988). Choices that have led to what are encoded as positive outcomes are reinforced, while the propensity to engage in actions that have led to negative outcomes is diminished. In this sense, experiential learning offers a form of backward-looking wisdom. In addition, the cognition—the belief about action-outcome linkages—itself may change as a result of prior experiences (Louis and Sutton, 1991). Thus, efforts at sensemaking (Weick, 1995) can be interpreted as a higher-order form of experiential learning.

Although prior work has addressed how experience may lead to changes in cognitive representations (Louis and Sutton, 1991; Weick, 1995), few scholars have addressed the opposite, how cognition influences subsequent processes of experiential learning. Understanding this linkage is important not only for addressing the general question of how cognitive representations affect choice processes but also for exploring the consequences of cognitive change on the cumulated tacit knowledge that has built up through experiential learning. We provide a formal basis for addressing such issues by developing a simulation model that jointly examines a forward-looking logic premised on simplified representations of the actors' world and a backward-looking logic premised on experiential wisdom. In doing so, we address some important substantive questions, as well as link two formerly disparate literatures that have stemmed from the notion of bounded rationality.

COGNITIVE AND EXPERIENCE-BASED CHOICE

Three basic properties distinguish cognitive from experiential-based choice: the mode of evaluation of alternatives, the extensiveness of alternatives considered, and the location of these alternatives relative to current behavior. Perhaps the most central element of these three dimensions is the process by which possible alternatives are evaluated. Cognition permits the assessment of alternatives "off-line" (Lippman and McCall, 1976), that is, actors need not engage in an activity in order to evaluate it. Actors evaluate alternatives

Cognitive and Experiential Search

based on their understanding of the world and the probable consequences of engaging in the proposed behaviors. The standard image of business planning corresponds to notions of off-line search—analysts with yellow pads and spreadsheets comparing a set of alternatives for their financial implications. Conversely, experiential processes inherently require at least partial implementation of an alternative in order to evaluate its efficacy, making them “on-line” evaluation mechanisms (Lippman and McCall, 1976).¹ Actions are tried, their outcomes experienced, and subsequent revisions to the prior actions may occur (Levitt and March, 1988).²

Forward-looking cognitive-based choice processes also tend to differ from feedback-based, experiential processes in the extensiveness of the set of alternatives considered. At the extreme, an experiential process implies that only one alternative at a time may be explored, such that alternatives are explored sequentially. In contrast, an off-line, cognitive process may invoke a broad set of alternative actions. Finally, independent of the extensiveness of the set of alternatives that is considered is the question of the degree to which these alternatives differ both from the organization’s current behavior and from each other. The risk of experimentation with alternatives that differ substantially from current behavior may be intimately tied to the way possible alternatives are evaluated.

Figure 2 provides a pictorial representation of different choice processes. The classic model of rational decision making can be positioned in this space as corresponding to the off-line evaluation of the whole set of alternatives, both local and distant. A simplified cognitive representation also permits the off-line evaluation of alternatives both local and distant; however, due to the simplification inherent in such a representation, the set of alternatives considered is likely to be less extensive than in the rational model and to be less precisely characterized. In contrast, processes of local experimentation (Cyert and March, 1963; Nelson and Winter, 1982) are examples of on-line evaluation. The alternative is explored through actual experience, rather than putatively in the form of a mental model (Johnson-Laird, 1983; Holland et al., 1986). Typically, experiential search (i.e., a process of on-line experimentation with a modest set of alternatives) is viewed as being focused on the neighborhood of current activity (March and Simon, 1958; Cyert and March, 1963). The process of logical incrementalism (Quinn, 1980), while also focused on a small set of alternatives in the near-neighborhood of current action, involves the off-line evaluation of these alternatives.

The impact of alternative sampling strategies varies dramatically depending on whether the evaluation mechanism is one of on-line experimentation or off-line cognition. If the evaluation of alternatives is off-line, then variation in the sample is generally an attractive property. If low-outcome draws can be costlessly discarded, then greater variance in the sample, holding the mean constant, increases the expected value of those draws that are adopted. This is the basic intuition behind the recent interest in the idea of “real options” in the business strategy literature (Bowman and Hurry, 1993). In a process of on-line experimentation, however, such variation

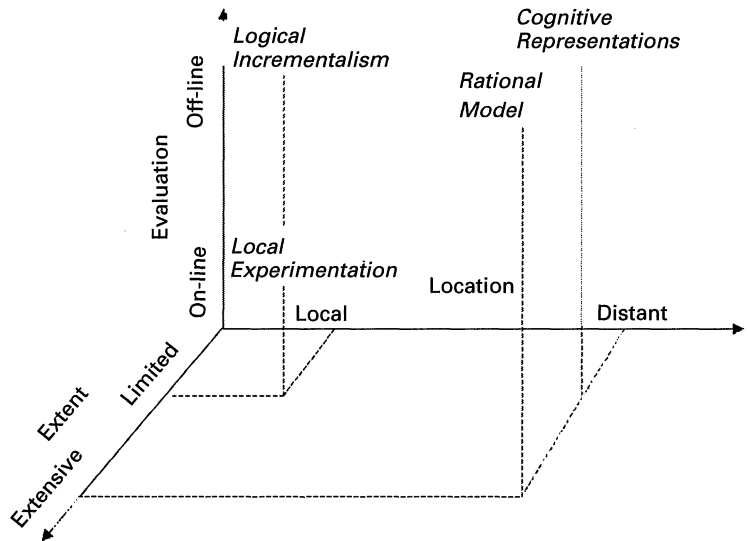
1

For the moment, we are ignoring issues of population-level learning (Miner and Haunschild, 1995) and vicarious learning (March, Sproull, and Tamuz, 1991), though we return to these issues later.

2

Thus, for our purposes, the terms forward-looking, cognitive-based choice, and off-line evaluation are equivalent. They all indicate that choice is based on some mental model or cognitive representation that suggests the outcomes associated with a proposed action. The terms experiential, backward-looking, and on-line evaluation are similarly equivalent. They all refer to choice processes in which knowledge of the linkage between actions and outcomes is derived on the basis of actions taken.

Figure 2. Choice processes.



may prove fatal because the actor experiences the consequences of each experimental draw. Consequently, neighborhood search, searching among alternative actions in close proximity to current behavior (March and Simon, 1958), is an important mechanism of on-line search. Alternative actions in the immediate neighborhood of the current behaviors are not likely to vary greatly in their efficacy from the current performance. As a result, there is typically not tremendous risk associated with experimenting with such alternatives, even though these alternatives represent some variation in action and, consequently, in performance. Therefore, neighborhood search represents a balance of the need to exploit the current wisdom associated with existing actions while, at the same time, engaging in some degree of search (i.e., exploration) for superior alternatives.

Of course, the clear distinction made in the theoretical literature between on-line and off-line search (Lippman and McCall, 1976) is often blurred in actual practice. Manufacturers of new airframes not only engage in ex-ante, off-line evaluation of possible new alternatives in the form of computer simulations of the aerodynamic properties of the proposed forms but also test prototypes in wind tunnels. Wind tunnels, test marketing, and experimental plants represent partial, on-line experimentation. Real economic activity is at stake in these trials. A full commitment of resources, however, is not at stake—no airline passengers are at risk in the wind tunnel, nor are overall sales in danger as an outcome of a marketing experiment with couponing in a particular region. Nevertheless, this distinction between on-line and off-line experimentation is a powerful one, and for simplicity in modeling, we treat it as being dichotomous. We model an off-line, or cognitive, choice process based on a simplified representation of the problem space. We also model experiential search as a process in which alternatives are evaluated as the result of on-line experimentation; furthermore, we postulate a sam-

Cognitive and Experiential Search

pling strategy of neighborhood search consistent with such an evaluation mechanism.

Cognitive Representations

As boundedly rational actors, we cannot envision the full set of alternatives available to us, nor can we completely specify the causal linkages between possible alternative actions and possible outcomes. Our attempts to do so are limited by both the vast number of potentially relevant policy variables and the complex set of interrelationships among these variables. Halford et al. (1994) found that the most complex statistical relation that individuals can process in working memory is a three-way interaction (i.e., three independent variables and one dependent variable, for a total of four dimensions). As a result, behavior is often driven by simplified representations based on implicit theories of the world (Kelley, 1971; Argyris and Schon, 1978). A critical element of expertise is the divergence in experts' representation of a problem (or solution) space from that of novices (Chi, Glaser, and Josselyn, 1981).

The intelligence of choice processes is driven not only by the intelligence of the representation that actors use but also by the computational procedures or algorithms used to identify an action, given a representation (Thagard, 1996). Models of rational choice focus on a particular computational algorithm of choosing the action that maximizes one's payoff. The behavioral inaccuracy of rational choice models, however, may have less to do with the inappropriateness of the algorithm—choose the best alternative—than with the assumption that actors apply that algorithm to the actual problem representation (Camerer, 1997). Therefore, we assume that, given a representation of their environment, actors are able to identify the most attractive action in their simplified cognitive space.

Translating cognition into behavior. The choice of action based on a given representation, however, does not fully characterize behavior. As a result, it is necessary to specify how an incomplete cognition is expressed in terms of actual behaviors. For instance, a management construct such as lean production does not fully specify the particular actions that should be taken on the shop floor or in dealing with suppliers, although it does offer some guiding principles for these issues. As a result, as such concepts diffuse, they may play out rather differently in different contexts (Zbaracki, 1998). As characterized by the cognitive representation, an alternative may be considered to be a template, or outline, of a possible action. Conceptually, there seem to be two basic mechanisms that flesh out such a template. One is the existing set of routines and behaviors of the organization. These actions may serve as defaults for choices that are not specified by the template. Alternatively, the template, and possibly past practices, may serve as a starting point for a process of experiential learning.

Those two processes, the imposition of a set of historical, default routines and experiential search, may be combined in the following manner. The organization chooses a policy within the context of its explicit cognitive representation. An inde-

terminacy results from the fact that the mapping from the cognitive representation to an actual set of behaviors is a one-to-many mapping. This indeterminacy is resolved by the use of existing practices as defaults for the elements not specified by the cognitive representation. With these defaults as a starting point, the organization may subsequently experiment over the set of behaviors that is consistent with the template. As a result, actors with the same cognition may find themselves engaged in distinct sets of behaviors. Cognition provides a guide to choice but does not strictly determine the actual set of behaviors that emerge.

MODEL

Fitness Landscape

A critical factor influencing the intelligence of alternative search processes is the degree to which neighboring alternatives are related to each other in their values. Because it is the correlated nature of the space of alternatives that gives local search processes their power, if adjacent locations in the solution space do not have similar payoffs, then a process of neighborhood search would have no more intelligence than the random selection of alternatives. Kauffman's (1993) NK landscape is well suited to model the degree to which alternative actions are correlated with one another. This framework builds on Wright's (1931, 1932) notion of a fitness landscape. Wright, a biologist interested in the evolution of organisms, considered the mapping of the attributes of organisms (genes) to the fitness level of the overall organism (i.e., the phenotype). Kauffman's representation of a fitness landscape is a simple but powerful framework for considering questions of adaptive learning (Levinthal, 1997). The variable N refers to the number of distinct attributes in an overall policy choice. For instance, in a choice of a firm's business strategy, a number of decisions must be made, including decisions about how the product or service is to be marketed, such as issues of brand name and distribution channels, and how it is to be produced, such as the degree to which activities will be done within the firm or outsourced. The variable K refers to the extent to which the payoff associated with one policy choice depends on other policy choices.

The issue of interdependence among policy variables has been highlighted in recent empirical work in the human resource literature (Ichniowski, Shaw, and Prennushi, 1997) and analyses of lean production (MacDuffie, 1995). Researchers now speak of systems of human resource practices, premised on the belief that the returns to a particular practice, such as selection policies, depends on other policies, such as training and compensation. Porter (1996) used the tool of activity maps to consider the broad set of interrelationships among a firm's policy choices.


The degree of interrelationship among policy choices has a somewhat counterintuitive implication for the topography of a fitness landscape. The fitness landscape is the mapping from the N policy choices to a payoff value. When the value of K is low and there is little interaction among policy choices, then the fitness landscape is smooth or highly correlated. With a

Cognitive and Experiential Search

low value of K , a change in one policy has little impact on the fitness contribution of other choices. As a result, incremental changes in the vector of N policy variables have a relatively modest impact on overall performance. In contrast, with a high K value, a change in one policy, such as distribution strategy, has implications for the payoff contribution of a large number (K) of other policy choices. In such a setting, even an incremental change in the vector of N policy variables may substantially change the overall payoff level. As a result, the fitness landscape becomes less correlated, or equivalently, more rugged, with a higher K value.

Figure 3 depicts a smooth fitness landscape. Adjacent locations within the landscape tend to have similar fitness values. This landscape has another, related property in that there is only one peak in the landscape. The term peak is defined as a point within the landscape such that any incremental movement from that location will diminish performance.

Figure 3. Smooth fitness landscape.




When there are significant interaction effects among policy variables, there may be a number of local peaks. Figure 4 depicts a landscape with considerable interaction effects in the fitness contribution of policy variables and, as a result, the landscape has a number of local peaks, forming a rugged fitness landscape. Each peak can be viewed as comprising a consistent set of practices—a configuration, in the language of Miller and Friesen (1984). With no interaction effects, the landscape has a single peak and one can identify superior policies or best practices whose superiority is independent of the other characteristics of the firm. With substantial interaction effects, the notion of universal best practice is not valid.

More formally, the fitness landscape is modeled as follows. A policy is characterized as consisting of N attributes where each attribute can take on two possible values.³ Thus, the fitness landscape consists of 2^N possible policy choices, with the overall behavior of the organization characterized by a vector $N\{x_1, x_2, \dots, x_N\}$, where each x_i takes on the value of 0 or 1. The contribution of a given attribute, x_i , of the policy vector to the overall payoff is influenced by K other attributes.

3

The model can be extended to an arbitrary finite number of possible values of an attribute, and the qualitative properties of the model are robust to such a generalization (Kauffman, 1989).

Figure 4. Rugged Landscape.



As a result, the payoff to a particular choice, x_i , can be represented by the following expression: $f(x_i | x_{i1}, x_{i2}, \dots, x_{iK})$. The K variables with which a given element x_i interacts is specified as being the K adjacent elements $(x_{i+1}, x_{i+2}, \dots, x_{i+K})$.⁴ Therefore, each attribute can take on 2^{K+1} different values, depending on the value of the attribute itself (either 1 or 0) and the value of the K other attributes with which it interacts (each of these K values also taking on a value of 1 or 0).

A random number drawn from the uniform distribution from zero to one is assigned to each of the possible $f(x_i | x_{i1}, x_{i2}, \dots, x_{iK})$ combination. Thus, the framework specifies the intensity of interaction effects via the parameter K but provides no restrictions on the particular functional form of the interaction effect. The overall fitness value associated with the full vector of N values, $F(x_1, x_2, \dots, x_N)$, is simply the sum of these individual contributions divided by N : $F(x_1, x_2, \dots, x_N) = \sum_{i=1 \text{ to } N} f(x_i | x_{i1}, x_{i2}, \dots, x_{iK}) / N$.

To illustrate how the fitness landscape is formed, we can consider how payoffs are determined for a policy space where N , the number of dimensions, equals 10 and K equals 3. Suppose that a policy is specified by the array (1,0,0,1,1,1,0,1,0,0). The value of the first element of this array depends on the K successive elements in the array. Thus, the value of 1 in the first element of the array depends on the value of the second through fourth elements of the array. A random number, generated from a uniform distribution ranging from 0 to 1, is assigned to constitute the fitness contribution of a 1 in the first element of the array when there is a 0, 0, and 1 in the second, third, and fourth elements of the array, respectively. A distinct random number is assigned for the case in which there is a 1 as the second element of the array rather than a 0, or any change in the third or fourth elements. This assignment is repeated for each of the N attributes of the organization, and the overall fitness for a particular organization is simply the average for the N attributes.

⁴ The vector is treated as circular for the purposes of determining fitness values. For instance, if $N = 10$ and $K = 3$, then the fitness contribution of the ninth element in the vector will depend on the value of the tenth, first, and second elements in the vector. An alternative formulation of the model is to postulate that the interaction effects are with K randomly chosen other elements of the vector. This alternative formulation results in a similar fitness landscape (Kauffman, 1989).

Cognitive and Experiential Search


Fitness landscape and cognition. We are interested in using the structure of a fitness landscape to model not only processes of experiential learning but also actors' simplified cognitive representation of their decision contexts. In the context of intelligent adaptation on a fitness landscape, the issue of cognition becomes the actors' representation of the landscape. We assume that cognitions are grounded on the actual landscape but that they constitute a simplified caricature of the decision context. A simple way to capture these two properties is to assume that cognitive representations are of lower dimensionality than the actual landscape. This assumption of low-dimensional cognitive representations is consistent with arguments in the cognitive psychology literature (Johnson-Laird, 1983; Halford et al., 1994) and work on managerial cognition (Weick, 1990). It is also consistent with the normative traditions in the management literature that offer low-dimensional typologies such as the Boston Consulting Group matrix (Hax and Majluf, 1984) and generic strategies (Porter, 1980) to help structure the choice of firm strategy.

To capture the notion that cognitive representations of the fitness landscape are a simplification of the actual landscape, actors are assumed to have a representation that consists of N_1 dimensions, where $N_1 < N$. The question remains as to what constitutes the mapping between this simplified representation and the actual landscape. We assume that the cognitive representation, while a simplification, is nonetheless grounded in the actual landscape. More precisely, each point in the cognitive representation is assigned a fitness value equal to the average fitness value of the set of points in the actual fitness landscape that are consistent with this point. For a point in the N_1 dimensional space, there are 2^{N-N_1} points in the actual fitness landscape that are consistent with it. As a result, as N_1 decreases, the cognitive representation becomes increasingly crude. While this representation is an unbiased estimate of the payoff associated with the actual landscape, it is not a particularly good predictor of any given point in the landscape. One could readily imagine a more sophisticated structure, such as a factor analysis of the actual payoff surface that provides a low-dimension representation that maximizes the informativeness of the representation. Thus, we are postulating a representation that poses relatively modest assumptions about actors' cognitive sophistication, while at the same time yielding a representation consistent with the underlying fitness surface.

One summary indicator of the complexity of a fitness landscape is the number of local peaks within the landscape. A point within the landscape is a local peak if any incremental movement in the policy space from that position degrades performance. Local peaks limit the effectiveness of processes of incremental experiential search and result in competency traps (Levitt and March, 1988) in which local evaluation of the landscape suggests that the organization cannot improve its performance when a more global evaluation of possibilities might suggest superior alternatives.

Figure 5, which is based on the average number of local peaks over a sample of 100 independent landscapes, indi-

Figure 5. Ruggedness of real and cognitive landscape (N = 10).



icates the number of local peaks for a set of representations that vary in their dimensionality (i.e., N_1) for varying levels of K . With $K = 0$, the actual landscape contains a single peak and, as a result, so do all simplifications of the true landscape. As K increases, however, the number of local peaks on the actual fitness landscape rapidly proliferates. In contrast, cognitive representations, even those with fairly high dimensionality (i.e., N_1 values close to N), are rather insensitive to increases in K .

Cognitive representations offer a small set of focal alternatives corresponding to the relatively few local peaks on the cognitive landscape and thereby provide a powerful direction for individual choice processes. The cognitive representation tends to be less rugged, or multi-peaked, than the actual underlying fitness landscape as the result of two by-products of the simplification process. First, the sensitivity in payoffs to movement within the N_1 -dimensional cognitive landscape is reduced by averaging payoffs across the actual set of policy choices in the N -dimensional landscape. Second, while the cognitive simplification is specified as a reduction in the perceived dimensionality of the landscape, it implicitly results in a reduction in the apparent K value, or perceived degree of interaction among the policy variables. Some interaction effects lie outside the cognitive representation—interactions among the $N - N_1$ variables of which the actor is not cognizant are not reflected in the landscape.

Organizations are assumed to choose a policy that their cognitive representation suggests maximizes their payoff. As noted earlier, however, the mapping from the cognitive repre-

Cognitive and Experiential Search

resentation to an actual set of behaviors is a one-to-many mapping. More precisely, there are 2^{N-N_1} actual policy arrays corresponding to each point within the cognitive representation. Thus, the organization in choosing a point within its representation is really choosing a region of the landscape—the region of 2^{N-N_1} possible behaviors that have the highest average payoff. How broad this region is corresponds to the crudeness or precision of the cognition. This indeterminacy may be resolved by the use of existing practices as defaults for the $N-N_1$ elements not specified by the cognitive representation. With these defaults as a starting point, the organization may subsequently experiment over the set of 2^{N-N_1} behaviors that are consistent with the N_1 policy choices determined by the cognitive representation. As a result, actors with the same cognition may find themselves engaged in distinct sets of behaviors. Cognition provides a guide to choice but does not determine the actual set of behaviors that emerge.

Experiential search is characterized as a process of local search. Search is local in that only one element of the N dimensional array is varied at a time. The particular policy variable that is experimented with is chosen at random. If the new array of policy choices increases performances, it becomes the basis for subsequent efforts at local search. Alternatively, if the organization's performance declines, then the organization returns to its prior starting point for its subsequent efforts at local search. The organization is assumed to remember which of the local experiments were unsuccessful. As a result, the organization either identifies a new superior alternative or, after N trials, stops engaging in experiential search and persists in what is a local peak.

ANALYSIS

We use this analytical structure of fitness landscapes and cognitive representations of these landscapes to engage in two sets of analyses. The first examines the interrelationship between cognition and experiential learning processes, while the second set extends this comparison to a setting in which actors' cognitive representations may themselves shift over time. In exploring the basic interrelationship between the role of cognition and experiential search, we carry out three simulations. In the first, we examine how a joint process of experiential and cognitive choice compares with pure experiential search. This analysis illustrates how a cognitive representation may usefully "seed" a subsequent process of experiential learning. In the second simulation, we show that holding to this cognitive frame may have the further virtue of preventing the process of experiential search from wandering off in potentially dysfunctional directions. In the third simulation, we consider how the process of joint cognitive and experiential search competes with a purely experiential process when they are examined in the context of a selection process. Selection pressures exacerbate the advantages of cognition by making more critical the speed at which superior alternatives are identified.

A higher-level form of experiential learning occurs if one shifts one's cognitive representation, rather than shifting the choice of individual policy variables. We examine the effect

of shifting cognitive representation first in a purely cognitive choice process. We then consider the impact of changing cognitive representations in conjunction with experiential learning. The final simulation of these choice processes, in a setting in which the landscape itself changes, illustrates how the trade-off between cognitive adaptation and local search is influenced by the rate of change in the environment.

To ensure that the results reflect the underlying structure of the model and not merely particular realizations of a highly stochastic process, the results are based on the average behavior of organizations over 100 independent runs of the simulation model. For each of these runs, a distinct landscape is specified. Each of these landscapes has the same structure in terms of K and N but is seeded independently. The analysis examines the search behavior and performance of a population of organizations with distinct search strategies operating on the same fitness landscape.⁵

When two search modes are being contrasted, the total number of organizations in the population is set at 90, with 45 organizations in each subgroup. The 45 organizations within one subgroup are assigned a location at random within the actual fitness landscape. If these organizations act, in part, on the basis of a cognitive representation of the landscape, they are seeded with a randomly assigned set of N_1 dimensions that form the basis of their representation. To make the analysis as controlled as possible, the other set of 45 organizations within the population differ in their search mechanism but are clones of the initial set of 45 in their initial position in the actual landscape. In addition, if applicable, the other 45 organizations share the same initial cognitive representation (i.e., the set of N_1 dimensions are the same). An analogous setup is implemented for three contrasting subpopulations, with the only difference being that the number of organizations in each subpopulation is set at 30 organizations.

Cognitive and Experiential Search

We compared processes of forward-looking, cognitive search with the backward-looking intelligence of experiential search and examined how the joint processes of cognitive and experiential search interact. Figure 6 indicates the average performance of a subpopulation of organizations that engage in both cognitive and experiential search. In addition, to isolate the effect of joint cognitive and experiential search, a subpopulation is included that engages exclusively in experiential search.⁶

The addition of cognition, or forward-looking intelligence, dramatically enhances the initial adaptive behavior. Organizations immediately identify the peak with respect to their cognitive representation. While this cognitive peak yields, on average, an actual payoff that is superior to the randomly specified initial location on the landscape, this payoff is still relatively modest when contrasted with the potential payoffs that can be realized on the actual payoff landscape. Having identified the optimal choice of N_1 attributes as suggested by their cognitive representation of the fitness landscape, organizations then explore the remaining $N-N_1$ policy variables experientially.

5

This analysis does not consider issues of the deformation of the landscape as a result of the density of organizations in a particular location, although the analysis of selection pressure does consider diffuse competitive interaction in the form of the effect of relative fitness values on survival rates.

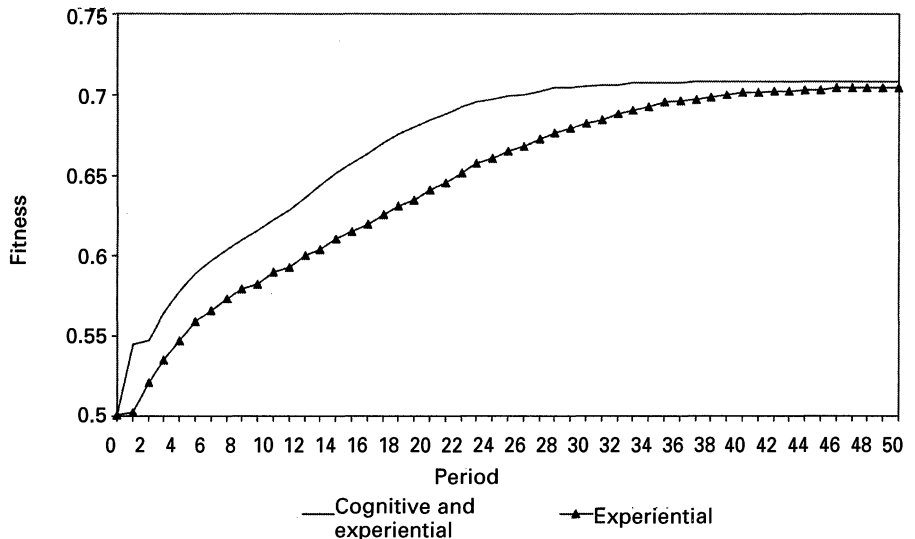
6

A third possibility, a population whose choice is based purely on cognitive search is embedded in the results for the subpopulation engaged in joint cognitive and experiential search. The performance reached in period 1 is the result of cognitive search. In the absence of subsequent experiential search, this same performance value would hold across the subsequent time periods.

124/ASQ, March 2000

Cognitive and Experiential Search

Figure 6. Joint cognitive and experiential versus experiential search (N = 10, K = 3, N1 = 3).*



* For cognitive and experiential search, the average ratio of standard deviation to mean equals 0.099; for experiential search it equals 0.129.

Most of the performance enhancement that arises with joint cognitive and experiential search comes from the subsequent slow but steady improvement associated with the experiential search effort. This result stems from the complexity of the actual payoff landscape and the relative crudeness of the cognitive representation. As the number of dimensions of the cognitive representation (N1) increases towards N, the role of experiential search is reduced. Even though much of the performance improvement is associated with experiential search, however, the initial period of cognitive search has a persistent effect on performance. Experiential search draws an organization toward a local peak in the payoff landscape. The particular local peak that is reached via experiential search is a function of the organization's starting position in the fitness landscape.

Those points within the landscape that lead to a common peak via a process of local search are said to belong to the same "basin of attraction" (Kauffman, 1993).⁷ The extent of a basin of attraction is positively correlated with the height of the local peak with which it is associated (Kauffman, 1993). Thus, local experiential search is intelligent not only in that it leads an organization to a local peak in the landscape but, if organizations are dispersed randomly over the landscape, local experiential search tends to lead organizations to relatively higher local peaks. Despite the intelligence of local experiential search, cognition helps an organization to identify, on average, superior basins of attraction. The set of policy choices suggested by the cognitive representation need not correspond to a local peak in the actual fitness landscape, let alone a global peak, but the global peak of the organization's cognitive landscape generally corresponds to an attractive region of the actual landscape. While the cognitive representation is crude, cognitive peaks nonetheless tend to lie in superior basins of attraction with respect to the actual fitness landscape.


7

One could think of this as a valley that circumscribes a mountain. Starting from any point in this valley, an upward climb leads to a common point, the peak of the mountain.

The cognitive representation provides not only a powerful suggestion for an initial choice of organizational form but also a useful discipline on subsequent efforts at experiential search. Figure 7 indicates the performance over time of two subpopulations that engage in both cognitive and experiential search. Both subpopulations in the initial period choose the N1 policies suggested by their cognitive representations and subsequently engage in experiential search. The difference in their behavior lies in the fact that one of the subpopulations is restricted to sustain this choice of N1 policies, while the other subpopulation engages in experiential search over the whole set of N variables. Those organizations not constrained by their cognitive template wander much more broadly over the landscape than is desirable, some of them sufficiently far from the initial point in the landscape suggested by their representation that they end up in an inferior basin of attraction.

Effects of selection pressure. Although pure experiential learning approaches the performance of the joint cognitive and experiential search process, there is a considerable period of time during which performance under the two processes substantially differs. As a result, the efficacy of cognition is seen more strongly if we consider a competitive ecology within the population of organizations, with organizations exiting and entering the population, and relatively poorly performing organizations (i.e., those with low fitness values) tending to be selected out of the population. The standard representation of such a selection mechanism in the mathematical biology literature is to assume that the probability of mortality is $1 - F/F_{Max}$ where F is the focal organism's fitness level,

Figure 7. Cognitions and basins of attraction (N = 10, K = 3, N1 = 3).*



* For cognition and constrained experiential search, the average ratio of standard deviation to mean fitness value equals 0.097; for cognition and unconstrained experiential search, it equals 0.181.

Cognitive and Experiential Search

and F_{Max} is the fitness value of the maximally fit organism in the population (Wilson and Bossert, 1971). This F_{Max} value need not be the highest possible fitness value in the landscape but, rather, corresponds to the maximal fitness value obtained by an organization within the population. Organizations that exit are replaced by replicating a randomly chosen organization from the population of surviving organizations in that period.

Figure 6 indicated that the average performance of the subpopulation of organizations engaged in experiential search and the subpopulation engaged in both cognitive and experiential search becomes similar in the absence of any selection pressure. When selection pressure is introduced, however, organizations engaged in cognitive search tend to dominate the population of organizations. Figure 8 indicates the changing mix of cognitive and non-cognitive organizations when the initial population is seeded with equal numbers of the two types of organizations.

The superiority of joint cognitive and experiential search reflected in figure 8 is compounded when this analysis is carried out in a landscape with a greater degree of interdependence among the policy variables (K) and, in turn, the complexity of the landscape.⁸ Although a higher value of K degrades the fidelity of the cognitive representation for a given value of N_1 , a higher K increases the dispersion of performance realized by a population of purely experiential organizations. Increases in K increase the ruggedness of the actual landscape and cause purely experiential organizations to be trapped by poor local peaks. As a result, the superiority of the joint cognitive and experiential search over purely experiential search indicated in figures 6 and 8, generated with a landscape with a K value of 3, is enhanced for landscapes with a higher K value.


Shifting Cognitive Representations

An important form of adaptation is the shift in cognitive representations themselves (Louis and Sutton, 1991; Weick, 1995). Different representations may have more or less fidelity to the actual environment in which an organization operates. In addition, by shifting representations, organizations may also enhance their performance by sequentially allocating attention to elements of the choice problem. Yet such shifts also pose risks. Most obviously, there is the risk that the new representation is an inferior characterization of the environment. Even with a shift to a superior cognitive representation, however, there is the risk of reorganization itself (Hannan and Freeman, 1989; Amburgey, Kelly, and Barnett, 1993). The shift in policies prompted by the new representation may result in the loss of the experiential wisdom accumulated in the context of the prior representation. To understand the possibly adaptive role of shifting a representation and how its adaptiveness may depend on the role of experiential knowledge, we undertook two sets of analyses. We first examined the effect of shifting cognitions in the absence of experiential search and then considered the implications of shifting cognitive representations in the context of joint cognitive and experiential search. Lastly, we examine the impli-

8

The characterization of complexity by the intensity of interaction effects was introduced by Simon (1969).

Figure 8. Survival rates of cognitive and experiential organizations (N = 10, K = 3, N1 = 3).



cations of shifting cognitive representations in the context of a fitness landscape that itself changes over time.

Under a process of cognitive search, the organization first makes a policy choice based on its understanding of payoffs as characterized by a set of N_1 attributes. The remaining $N - N_1$ policy variables are left fixed. With a shift of representation to a new set of dimensions, a new set of variables becomes subject to conscious choice. Whether such shifts in representation facilitate intelligent adaptation depends on the degree to which the fitness landscape is correlated, or the value of K . In the limit, with an uncorrelated fitness landscape, there is no carryover in the intelligence of choices made on the basis of one representation to another representation. With a correlated fitness landscape, however, shifting representations and the associated sequential attention to the set of policy variables is an effective mechanism for dealing with the inevitable simplification of any given representation.


If representations change over time, what determines the frequency of change and how is the new representation chosen? Clearly, there is a wide range of possibilities. To provide some sense of the range of behaviors, consider an extreme setting in which the frequency and the choice of representation are randomly determined, as well as a setting in which there is some intelligence in both the timing of a shift to a new representation and the choice of the new representation itself. Random shifting of representations implies no instrumental logic as to when a shift occurs, nor an informed choice of representation when the organization does shift. As a result, such a process serves as a useful baseline against which to compare more contingent choice processes.

Cognitive and Experiential Search

Intelligence in the shift to a new representation should presumably reflect some insight about the possibility for performance improvement latent in a shift to an alternative representation. Relative performance, or fitness value, is a simple way to reflect this potential for performance improvement. The worse an organization's relative performance, the more likely it is that exploring alternative representations may lead to a higher fitness level. Population-level learning (Miner and Haunschild, 1995), or vicarious learning (March, Sproull, and Tamuz, 1991) from the cognitions of other, successful organizations in the population, is used as a basis for intelligence in the choice of a new representation. Imitating another organization's fully articulated set of behaviors is quite challenging (Szulanski, 1996), if not impossible (Rivkin, 2000), but imitation of what the key factors are by which to dimensionalize an environment is much more readily diffused. There is considerable evidence in the social psychological literature that concepts diffuse through social learning processes (Bandura, 1977). Management consultants and the business press are important facilitators of the diffusion of such concepts among business organizations (Abrahamson, 1996).

Figure 9 explores the effect of these processes of population level learning with respect to cognitive representations. The figure indicates the average fitness level over time for three subpopulations of organizations. The first retains a fixed representation over the course of the simulation. Organizations in this subpopulation may differ in their individual representation (i.e., their set of N1 values), but the representations themselves remain fixed over time. The second curve

Figure 9. Search via shifting representations (N = 10, K = 3, N1 = 3).*



* For fixed cognition, the average ratio of standard deviation to mean equals 0.153; for random change in representation, it equals 0.134; for semi-intelligent change in representation, it equals 0.117.

reflects the behavior of a subpopulation of organizations that engage in random shifts of their cognitive representations. Their behavior is random in two respects. First, the probability of a shift in cognition is unconnected to the organization's performance but is specified as a fixed probability each period. Second, the choice of a new cognition is not premised on any evidence of the power of the alternative cognition. One of the N1 dimensions is simply replaced at random by one of the N-N1 other dimensions.

The third curve reflects some modest degree of intelligent or at least contingent action. The intelligence is driven by comparisons between the organization and the overall population of organizations. The probability of a shift in cognition is a function of relative fitness, a comparison of the organization's fitness level with the maximally fit organization in the overall population. This sensitivity to relative performance is treated in a highly simplified manner as a performance threshold. If the organization's relative performance falls below a fixed percentage of the maximally fit organization, then it engages in a shift to a new cognitive representation.

The choice of a new cognitive representation is assumed to reflect some degree of intelligence as well. The organization imitates the cognition of one of the leading organizations in the overall population. As argued earlier, the replication of the choice of critical dimensions by which to evaluate one's environment seems much more plausible than the replication of a fully articulated set of behaviors. For the purposes of figure 9, the set of leading organizations was defined as those organizations whose fitness value lies in the top third of the population, and the performance threshold that provides an impetus to recategorization was set at 25 percent below the performance of the maximally fit organization.⁹ Shifting cognitions, whether at random or on a somewhat informed basis, enhances performance over time. With a fixed cognition, the organization immediately identifies the global peak with respect to its cognitive representation. Associated with this point is a set of N1 policy variables. Without experiential learning on the remaining N-N1 variables or a shift to a new cognitive representation, there is no basis for moving from the position identified in the initial period. What underlies the systematic increase in performance when cognitions shift, even randomly, is that with a fixed or inert cognition, actors immediately identify the global peak as defined by their cognitive representation. If a new cognitive representation is adopted, then a distinct set of N1 policy parameters will be identified on the basis of this new cognitive representation, but since the fitness landscape is correlated, there is still some intelligence associated with the choice of policies under the prior cognition. These policies are unlikely to be the ideal choices based on the policies suggested by the new cognition; however, they are likely to be superior to a random specification of the N-N1 variables that lie outside a given cognitive representation. In this manner, shifting cognitive frameworks effectively results in a sequential allocation of attention to different facets of the true landscape.

Chasing cognitive rainbows. The analysis of the adaptive implications of changing cognitive representations was car-

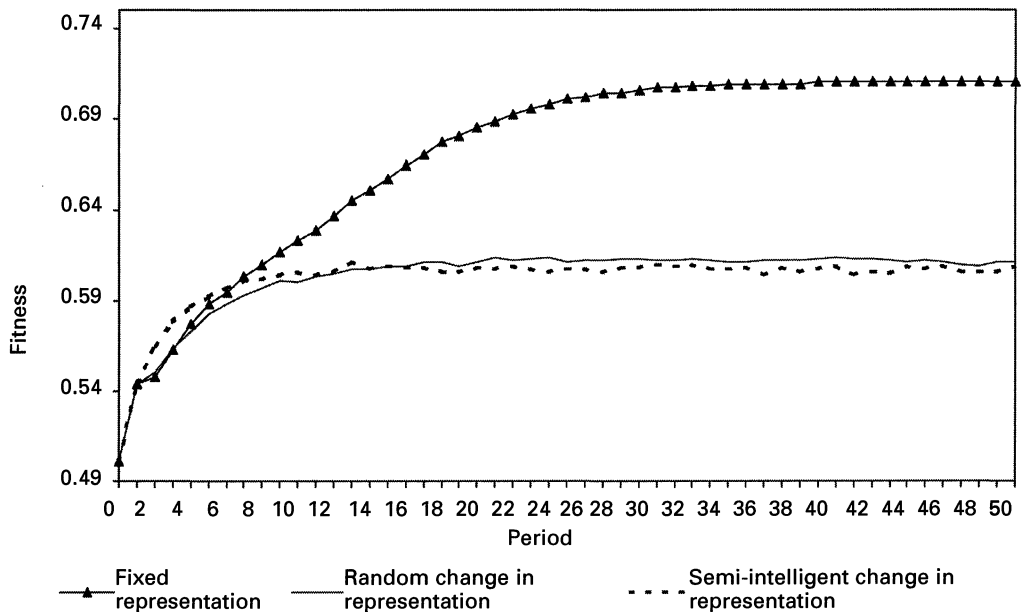
⁹ These results were rather insensitive to the precise parameter values chosen to characterize the processes of population learning. For instance, refining the set of leading organizations to a smaller set, such as restricting the definition of leading organizations to be the top 10 organizations rather than the top 30, does not have a significant impact on performance. Similarly, decreasing the performance threshold that induces search for a new cognition from 75 percent of the most maximally fit organizations to 50 percent, or increasing the stringency of the threshold by setting it at 85 percent, has a negligible impact on performance.

Cognitive and Experiential Search

ried out in a setting in which experiential learning was absent. If search is driven both by cognition and experiential learning, then changing one's cognitive representation poses an additional risk. Not only may one shift to a new representation that is inferior to the prior one, but the shift to a new policy array based on the new cognitive representation may cause the organization to negate the experiential wisdom that it has accumulated. The new representation motivates the organization to shift some of the policy variables that had emerged over time through the process of experiential search. The intelligence associated with this history of experiential learning over the space of the N-N1 parameters not reflected in the cognitive representation is negated in the cognitive evaluation of alternatives. This is akin to the risk of organizational change identified by Hannan and Freeman (1989) and empirically examined by Amburgey, Kelly, and Barnett (1993) in their analysis of organizational adaptation and mortality. Amburgey, Kelly, and Barnett found that organizational change substantially increases the immediate risk of mortality, even though these changes may lower the long-run hazard rate.

Figure 10 indicates the average performance over time for three subpopulations of organizations that vary with respect to whether they shift their cognitive representation. If organizations shift their representation, they are modeled as doing so in the same manner as in the prior analysis of pure cognitive adaptation in figure 9. The critical difference between this and the prior analysis is that the organizations are engaged in both cognitive-based search and experiential search processes. Contrary to the prior analysis of shifting cognitive representations in the absence of experiential learn-

Figure 10. Change in cognitive representation and experiential search (N = 10, K = 3, N1 = 3).*




* For fixed representation, the average ratio of standard deviation to mean fitness value is 0.083; for random change in representation, it is 0.128; for semi-intelligent change in representation it is 0.120.

ing, the subpopulation of organizations with a fixed cognitive representation now has substantially superior performance. While the sequential allocation of attention to different elements of the policy space facilitated adaptation in a setting of purely cognitive search, this process of sequential attention leads to the organization foregoing considerable experiential wisdom that more than offsets such gains.

Shifts in cognition, however, may prove valuable in the context of changes in the fitness landscape. A new cognitive representation may more effectively capture the new landscape than a prior representation. For instance, there have been fundamental shifts in technology and the basis of competition in the computer industry. Different cognitive representations are suggested by these shifts, from the vertically integrated business model of mainframes and minicomputers, to a PC-based computing platform, to the current network-oriented views of computing (Moschella, 1997). In addition, with dramatic changes in the fitness landscape, prior experiential wisdom is rendered largely obsolete. As a result, the virtue of rapidly identifying attractive regions of the landscape, via a cognitive process as a result of a shift in cognitive representation, can compensate for the loss associated with foregone experiential wisdom. Such a situation is reflected in figure 11, which depicts the performance of organizations that shift their representations relative to those with a fixed representation in the context of a dynamic fitness landscape. Every five periods, the landscape is perturbed by respecifying a new random draw for the fitness contribution of eight of the ten policy choices. At the point of environmental change, the relative performance of organizations that shift their cognitive representation improves, but only informed shifts in representations are sufficiently useful to

Figure 11. Cognitive dynamics in changing environments (N = 10, K = 3, N1 = 3).



Cognitive and Experiential Search

compensate for the loss of experiential wisdom. Subsequent to such a shift in the fitness landscape, the fitness level of organizations with fixed cognitive representations and those with shifting representations converge.

DISCUSSION AND CONCLUSION

The notion of bounded rationality offers us two different logics of choice and action: a backward-looking logic of stimulus-response learning and a forward-looking logic of consequences. Organizational behavior comprises both forms of intelligence. Our results indicate that some central phenomena lie at the interface of these two logics and can be explained only by considering them jointly. Although prior research has addressed how experience influences the formation of cognitive frameworks (Weick, 1995), little attention has been devoted to how cognition affects the accumulation of experiential wisdom. By focusing on this linkage and analyzing the consequences of cognitive change on accumulated experiential knowledge, we derived three sets of results.

First, cognitive representations play an important role in seeding and constraining the process of experiential learning. Even the crude cognitive representations modeled here provide a powerful starting point for subsequent efforts at experiential learning. This advantage is reflected in both asymptotic performance levels (figure 6) and, more strongly, survival rates in a competitive context (figure 8). This role of cognition is accentuated when the fitness landscape is more complex, or rugged. Cognition is useful not only in seeding the process of experiential search on a particular location in the fitness landscape but also in constraining the process of experiential search from wandering to less attractive regions of the landscape (figure 7). Experiential learning tends to consist of local, hill-climbing processes, the search over the space of alternatives in the neighborhood of current practices. Although such mechanisms are a powerful means to explore a complex problem space, they are inherently constrained by the local topography in which the actor lies. The standard pathology that results from this constraint is a competency trap (Levitt and March, 1988), in which incremental change efforts degrade performance but more substantial change efforts may identify a superior point on a fitness surface.

In contrast, cognitive processes allow a broader examination of the fitness surface, but unless an actor is endowed with the omniscience assumed in economic analyses (Milgrom and Roberts, 1990), there is no reason to presume that the global peak will be identified. We have created a structure that allows us to explore a middle ground between the myopia of local hill climbing and such omniscience. Imperfect cognitive representations of the fitness space help identify promising regions of the landscape. The incomplete template suggested by these representations is then fleshed out through a process of local experiential search.

Our second set of results highlights the effect of changing cognitive representations. Such changes can be an important form of adaptation in two different respects. First, the new representation may consist of a better mental model of the actor's environment, reflecting either weaknesses in the prior

representation or environmental shifts that render a previously adequate representation less effective. Second, even in a stable world, a shift among equally valid, or invalid, representations may enhance adaptation by facilitating the sequential attention to different facets of the actor's environment (figure 9). As a result, even with a limited representation, it becomes possible over time to span the full dimensionality of the problem space. Meyer and Gupta (1994) made a similar argument in their discussion of the need to shift performance metrics occasionally, as did Eccles and Nohria (1992) in their discussion of organizational change. For these authors, the impetus for change is not necessarily a change in the external environment but that the potential for further improvement has become limited given the current structure. The set of performance metrics and, more broadly, an organization's structure importantly influence actors' perceptions of their problem space. Only by offering a fresh perspective with a shift in performance metrics (Meyer and Gupta, 1994) or organizational structure (Eccles and Nohria, 1992) is further improvement possible.

The final set of results points to the dangers of such shifts in cognition and helps reconcile conflicting perspectives on organizational change (Tushman and Romanelli, 1985; Hannan and Freeman, 1989). Organizational change prompted by a shift in cognitive representation is costly in that prior, experiential wisdom may be largely negated (figure 10). When there is a high degree of interdependence among actions, the wholesale shift in behavior driven by a new cognitive representation may result in a tremendous loss of experiential wisdom. One is then faced with a trade-off between the benefits of a shift to a new, potentially more attractive region of the fitness landscape, engendered by the shift in cognitive representation, and the immediate performance decrement that stems from the loss of this experiential wisdom. The results thus provide further insight into the risks of reorganization and strategic change. In a fixed fitness landscape with experiential learning, shifts in cognitive representations result in a sharp decline in fitness level and, in turn, a sharp increase in the risk of organizational mortality. This behavior is consistent with the empirical observations of Amburgey, Kelly, and Barnett (1993) and, more generally, with ideas of structural inertia (Hannan and Freeman, 1984). As suggested in the analysis that underlies figure 11, however, the value of an updated cognition is enhanced and the loss of experiential wisdom is attenuated by the partial obsolescence of that wisdom as a result of environmental change. Consistent with this argument, empirical evidence suggests that strategic change plays a strong adaptive role in turbulent environments (Baum, 1990; Zajac and Kraatz, 1993; Tushman and Rosenkopf, 1996).

Cognitions and experience are also linked in ways that are not well reflected in the current modeling effort. Cognitive representations themselves are clearly an outcome of efforts at sensemaking with respect to prior experiences (Holland et al., 1986; Weick, 1995). Although we allowed for the possibility that the search for alternative cognitive representations may be prompted by poor performance outcomes, we did

Cognitive and Experiential Search

not model an explicit inferential process and the testing of actors' representation of their environment. In addition, the representations modeled here are quite stylized and atheoretic. They are a simple, mechanical mapping from a high dimensional space to a lower dimensional space. Clearly, more needs to be done to characterize more accurately the formation of cognitive representations and their persistence, or adaptation, over time. At the same time, the results derived here suggest that even simple models of the world have a tremendous potential to guide search processes.

A final limitation of the current analysis, and correspondingly an important opportunity for future work, is the lack of modeling of intraorganizational cognitive processes. While we have considered, to some degree, social learning processes (Bandura, 1977) across organizations, we have not explored such processes within organizations. Cognitive representations are challenged not only by experiential learning and feedback from the environment but also by conflict and persuasion within the organization.

Intelligent action is driven both by one's understanding of the world and adaptive responses to prior experiences. The former is limited by one's representation or mental model of the world, while the latter is constrained by the limited number of experiences that one may have relative to the vast set of possible actions that one may take. As a result, cognitive and experiential processes are complementary. Cognitive search is broad in that it considers a wide array of alternatives simultaneously, but misspecified in that these alternatives are evaluated on the basis of an incomplete mental model of the world. In contrast, experiential search is narrow because it enables actors to explore only a small set of alternatives at a given moment but lets them test these alternatives on the basis of the actual environment rather than a mere representation of the environment. Models of bounded rationality should incorporate both forms of intelligence.

REFERENCES

Abrahamson, E.

1996 "Management fads." *Academy of Management Review*, 21: 254-285.

Amburgey, T., D. Kelly, and W. Barnett

1993 "Resetting the clock: The dynamics of organizational change and failure." *Administrative Science Quarterly*, 38: 51-73.

Argyris, C., and D. Schon

1978 *Organizational Learning: A Theory of Action Perspective*. Reading, MA: Addison-Wesley.

Bandura, A.

1977 *Social Learning Theory*. Englewood Cliffs, NJ: Prentice-Hall.

Baum, J.

1990 "Inertial and adaptive patterns in organizational change." In L. Jauch and J. Wall (eds.), *Academy of Management Best Papers Proceedings*: 165-169.

Bowman, E., and D. Hurry

1993 "Strategy through the options lens: An integrated view of resource investment and the incremental-choice process." *Academy of Management Review*, 18: 760-782.

Camerer, C.

1997 "Progress in behavioral game theory." *Journal of Economic Perspectives*, 11: 167-188.

Chi, M. T. H., P. J. Feltovich, and R. Glaser

1981 "Categorization and representation of physics problems by experts and novices." *Cognitive Sciences*, 5: 121-152.

Cyert, R., and J. G. March

1963 *A Behavioral Theory of the Firm*. Englewood Cliffs, NJ: Prentice-Hall.

Eccles, R., and N. Nohria

1992 *Beyond the Hype*. Boston: Harvard Business School Press.

Fiol, C. M., and A. S. Huff

1992 "Maps for managers: Where are we? Where do we go from here?" *Journal of Management Studies*, 29: 267-285.

- Halford, G., W. Wilson, J. Guo, W. Gayler, J. Wiles, and J. Stewart**
1994 "Connectionist implications for processing capacity limitations in analogies." In K. Holyoak and J. A. Barnden (eds.), *Advances in Connectionist and Neural Computation Theory*, vol. 2: Analogical Connections: 363–415. Norwood, NJ: Ablex.
- Hannan, M., and J. Freeman**
1984 "Structural inertia and organizational change." *American Sociological Review*, 49: 149–164.
1989 *Organizational Ecology*. Cambridge, MA: Harvard University Press.
- Hax, A., and N. Majluf**
1984 *Strategic Management: An Integrative Perspective*. Englewood Cliffs, NJ: Prentice-Hall.
- Holland, J., K. Holyoak, R. Nisbett, and P. Thagard**
1986 *Induction: Processes of Inference, Learning, and Discovery*. Cambridge, MA: MIT Press.
- Huff, A. S.**
1990 *Mapping Strategic Thought*. New York: Wiley.
- Ichniowski, C., K. Shaw, and G. Prennushi**
1997 "The effects of human resource management practices on productivity: A study of steel finishing lines." *American Economic Review*, 87: 291–313.
- Johnson-Laird, P. N.**
1983 *Mental Models: Towards a Cognitive Science of Language, Inference, and Consciousness*. Cambridge, MA: Harvard University Press.
- Kauffman, S.**
1989 "Adaptation on rugged fitness landscapes." In D. Stein (ed.), *Lectures in the Sciences of Complexity*: 517–618. Reading, MA: Addison-Wesley.
1993 *The Origins of Order*. New York: Oxford University Press.
- Kelley, H.**
1971 *Attribution in Social Interaction*. Morristown, NJ: General Learning Press.
- Levinthal, D.**
1997 "Adaptation on rugged landscapes." *Management Science*, 43: 934–950.
- Levitt, B., and J. G. March**
1988 "Organizational learning." In *Annual Review of Sociology*, 14: 319–340. Palo Alto, CA: Annual Reviews.
- Lippman, S., and J. McCall**
1976 "The economics of job search: A survey." *Economic Inquiry*, 14: 155–187.
- Louis, M., and R. Sutton**
1991 "Switching cognitive gears: From habits of mind to active thinking." *Human Relations*, 44: 55–75.
- MacDuffie, J. P.**
1995 "Human resource bundles and manufacturing performance: Organizational logic and flexible production systems in the world auto industry." *Industrial and Labor Relations Review*, 48: 197–221.
- March, J. G.**
1994 *A Primer on Decision Making*. New York: Free Press.
- March, J. G., and H. A. Simon**
1958 *Organizations*. New York: Wiley.
- March, J. G., L. S. Sproull, and M. Tamuz**
1991 "Learning from samples of one and fewer." *Organization Science*, 2: 1–13.
- Meyer, M., and V. Gupta**
1994 "The performance paradox." In L. L. Cummings and B. M. Staw (eds.), *Research in Organizational Behavior*, 16: 309–369. Greenwich, CT: JAI Press.
- Milgrom, P., and J. Roberts**
1990 "The economics of modern manufacturing." *American Economic Review*, 80: 511–528.
- Miller, D., and P. Friesen**
1984 *Organizations: A Quantum View*. Englewood Cliffs, NJ: Prentice-Hall.
- Miner, A., and P. Haunschild**
1995 "Population level learning." In L. L. Cummings and B. M. Staw (eds.), *Research in Organizational Behavior*, 17: 115–166. Greenwich, CT: JAI Press.
- Moschella, David**
1997 *Waves of Power: The Dynamics of Global Technology Leadership, 1964–2010*. New York: American Management Association.
- Nelson, R., and S. Winter**
1982 *An Evolutionary Theory of the Firm*. Cambridge, MA: Harvard University Press.
- Porac, J. F., H. Thomas, and C. Baden-Fuller**
1989 "Competitive groups as cognitive communities: The case of Scottish knitwear manufacturers." *Journal of Management Studies*, 26: 397–414.
- Porter, M.**
1980 *Competitive Strategy*. New York: Free Press.
1996 "What is strategy?" *Harvard Business Review*, 74: 61–78.
- Quinn, J. B.**
1980 *Strategies for Change: Logical Incrementalism*. Homewood, IL: Irwin.
- Rivkin, J.**
2000 "Imitation of complex strategies." *Management Science* (in press).
- Simon, H. A.**
1955 "A behavioral model of rational choice." *Quarterly Journal of Economics*, 69: 99–118.
1969 *The Sciences of the Artificial*. Cambridge, MA: MIT Press.
1991 "Bounded rationality and organizational learning." *Organization Science*, 2: 125–134.
- Szulanski, G.**
1996 "Exploring internal stickiness: Impediments to the transfer of best practice within the firm." *Strategic Management Journal*, 17: 27–43.
- Thagard, Paul**
1996 *Mind: Introduction to Cognitive Science*. Cambridge, MA: MIT Press.
- Tushman, M., and E. Romanelli**
1985 "Organization evolution: A metamorphosis model of convergence and reorientation." In L. L. Cummings and B. M. Staw (eds.), *Research in Organizational Behavior*, 7: 171–222. Greenwich, CT: JAI Press.
- Tushman, M., and L. Rosenkopf**
1996 "Executive succession, strategic reorientation and performance growth: A longitudinal study in the U.S. cement industry." *Management Science*, 42: 939–953.
- Tversky, A., and D. Kahneman**
1986 "Rational choice and framing of decisions." *Journal of Business*, 59: S251–S278.

Cognitive and Experiential Search

Walsh, J.

- 1995 "Managerial and organizational cognition: Notes from a trip down memory lane." *Organization Science*, 6: 280-321.

Weick, K.

- 1979 *The Social Psychology of Organizing*, 2d ed. Reading, MA: Addison-Wesley.
- 1990 "Cartographic myths in organizations." In A. S. Huff (ed.), *Mapping Strategic Thought*: 1-10. New York: Wiley.
- 1995 *Sensemaking in Organizations*. Thousand Oaks, CA: Sage.

Wilson, E., and W. Bossert

- 1971 *A Primer on Population Biology*. Sunderland, MA: Sinauer Associates.

Wright, S.

- 1931 "Evolution in Mendelian populations." *Genetics*, 16: 97-159.
- 1932 "The role of mutation, inbreeding, cross-breeding and selection in evolution." *Proceedings XI International Congress of Genetics*, 1: 356-366.

Zajac, E., and M. Kraatz

- 1993 "A diametric forces model of strategic change: Assessing the antecedents and consequences of restructuring in the higher education industry." *Strategic Management Journal*, 14: 83-102.

Zbaracki, M.

- 1998 "The rhetoric and reality of total quality management." *Administrative Science Quarterly*, 43: 602-636.