Myopia of Selection: Does Organizational Adaptation Limit the Efficacy of Population Selection?

Daniel Levinthal
University of Pennsylvania
Hart E. Posen
University of Michigan

This paper develops and tests a model of the effectiveness of selection processes in eliminating less fit organizations from a population when organizations are undergoing adaptive change. Stable organizational traits, such as a search strategy or routine, do not imply that an organization's performance will remain stable over time or that cross-sectional differences in performance will persist. These properties create the possibility that population-level selection processes will be inefficient in that organizations with potentially superior long-run performance will be selected out. We theorize that organizational-level adaptation often results in fluctuations in current performance across time. These fluctuations may attenuate the degree to which current performance differences among organizations are indicative of future performance. As a consequence, search strategies that generate systematically different performance trajectories, even if they share a common long-run outcome, will generate differing survival rates. These ideas are explored using a formal simulation model employing the framework of NK performance landscapes. Our central finding is that selection may be systematically prone to errors and that these selection errors are endogenous to, and differ markedly across, firms' search strategies.●

The past 40 years has seen the emergence of significant interest in examining the population dynamics of organizations from an evolutionary perspective (Campbell, 1965; Hannan and Freeman, 1977; Nelson and Winter, 1982). A defining feature of the evolutionary perspective is that selection disproportionately removes organizations that are less fit. Yet there is reason to believe that this selection process may not be effective in eliminating organizations undergoing change. First, selection does not operate directly on relatively stable organizational traits, such as strategies or routines that are fundamental to firm-level heterogeneity. Instead, selection operates on the manifestations of those traits in terms of observable performance outcomes such as productivity, patents, new products, product quality and cost, market share, sales, or profits. Moreover, even stable organizational traits, including an organization's search strategy, may lead to performance manifestations that are highly variable. To the extent that organizations' stable traits lead to differences in the nature of learning, they may give rise to significant differences in the reliability of performance over time. Consequently, organizational populations may experience different survival patterns in the selection process not merely based on some inherent superiority of one form or another but also as influenced by the dynamics of change that different forms generate over time.

Second, though change processes and their effects on performance play out over time, selection processes are fundamentally myopic in that they cannot "see" those future effects. Selection is the outcome of a complex web of atomistic and, to some degree, also inherently myopic decisions about the distribution of resources across firms. The quintessential example of myopia is that of individual purchase decisions in product markets, where products' characteristics, such as quality and price, drive consumers' choice and hence

●

selection outcomes for the product. Even forward-looking actors in capital markets are myopic decision makers in that they must construct expectations of future performance solely based on such indicators of past and current performance as sales, profits, or patent filings. One important consequence of the myopia of selection is that even if selection is effective in removing inferior organizations at one point in time, it may be ineffective over time in that it may remove organizations that, had they survived, would have gone on to do well.

The question is whether there are conditions under which myopic selection can identify adapting organizations at intermediate points in time, while they are still in the midst of change, that will be superior in the long run. The existing literature provides some important rationales as to why selection processes may fail in this respect, indicating why inferior organizations, populations, or technologies may dominate the set of survivors (Barron, West, and Hannan, 1994; Carroll and Harrison, 1994; Barnett and Burgelman, 1996; Barnett, 1997). Three broad classes of answers have been proposed. First, selection does not act on a single dimension of performance but, in fact, on a multitude of potential performance criteria. Thus the outcome of the selection process can be as much a function of political and social competition as it is technical competition on the merits of alternatives (Anderson and Tushman, 1990). Second, organizational slack acts as a buffer to stave off failure even when organizations are weak from a technical efficiency standpoint (Levinthal, 1991b; Barnett, Greve, and Park, 1994; Barron, West, and Hannan, 1994). Third, path dependence can arise from network externalities (David, 1985; Arthur, 1989) whereby relative market shares of alternative forms or technology influence selection forces, and thus outcomes at a given point in time are in part a function of prior choices and outcomes (Barron, West, and Hannan, 1994; Carroll and Harrison, 1994; Barnett and Burgelman, 1996; Barnett, 1997).

A different reason, however, is that selection may be systematically prone to errors, which is suggested by the results of Carroll and Harrison's (1994) examination of the extent to which selection processes are historically efficient. Building on March and Olsen (1989), Carroll and Harrison (1994: 720) framed the issue as the extent to which "organizational arrangements observed at a particular point in time represent the unique outcome of some systematic process such as competition." Focusing on differential selection among types of organizations, they employed a computational model of the evolution of two populations of fully inert organizations under an ecological selection regime based on legitimation and competition. They concluded that deviations from historical efficiency may arise from differences in organizations' order of entry into a population due to the positive feedback effects of legitimation. As a result, it is possible for a population of inferior organizations to come to dominate the set of surviving firms if they enter first.

Carroll and Harrison (1994) focused on entry order in their model. Alternatively, as developed here, one can start with the assumption that organizations are engaged in a learning

process and thus exhibit performance manifestations that vary over time. The unfolding of an organization's particular history is then a function of both its idiosyncratic founding conditions and the strategies that guide its learning. Strategies that lead to higher variance in performance over time are likely to increase the extent to which selection is prone to err in selecting out organizations that would have gone on to achieve a high level of performance. In turn, selection errors, in which inferior firms survive selection, would decrease the average performance of the set of surviving firms. When organizations within a population are distinguished by strategies that differ in reliability and performance, an important element of selection consists of sorting among organizations in a single population. Under these conditions, differences across populations in reliability over time alone are sufficient to generate outcomes in which the population with inferior mean long-run performance, but higher reliability, dominates the set of surviving firms.

In the model developed in this paper, we highlight the central role of reliability in performance. In contrast to ecological models in which selection favors reliability (Hannan and Freeman, 1977, 1984), however, our model imposes no a priori relationship between reliability and mortality. Rather, the role of reliability emerges endogenously from a model in which firms that are subject to selection employ strategies that are inert, but these fixed search strategies guide organizational learning.

Important stable attributes of organizations such as search strategies or routines may differ in their contribution not only to organizational performance at a given point in time but also to the variance across time. In turn, these stable traits differentially affect the extent to which selection is prone to errors. We denote this effect on the efficacy of selection as "selectability," which is an organization-level construct that measures the information content embodied in current performance as a critical determinant of the efficacy of selection in identifying organizations that will be superior in the long run. For organizations that exhibit high levels of selectability, intermediate performance manifestations will provide superior information about their long-run performance, and consequently, selection in populations of such organizations will be more effective over time. Thus, even when a set of search strategies across organizations shares common long-run properties, the dynamics by which organizations reach this long-run performance level have important implications for the selectability among organizational populations and the efficacy of selection as a mechanism of population-level adaptation.

Questions about the nature of organizational change and its consequences for the efficacy of selection, though relatively unexplored, are not new. Campbell (1960, 1965), in laying the foundation of evolutionary perspectives in the social sciences, suggested one important factor driving the efficacy of selection. He argued that "if discovery or expansions of knowledge are [to be] achieved, blind variation is requisite"; blind implies that "specific correct trials are no more likely to occur . . . than specific incorrect trials" (Campbell, 1982: 86).

Nevertheless, the organizational analogue differs from its biological foundations in a critical way: organizations engage in adaptive trial-and-error search strategies that are intended to enhance their survival prospects (Simon, 1955; Cyert and March, 1963; Nelson and Winter, 1982). Simon (1969: 95), in contrast to Campbell, argued that such "trial and error is not completely random or blind." This classic debate highlights a central but unresolved question related to the efficacy of selection as a mechanism of population-level adaptation: Is the efficacy of selection invariant to the characteristics of the processes (e.g., blind vs. non-blind variation) that guide organizational change?

To gain insight into the effect of differential reliability on the efficacy of selection, we implemented a formal agent-based model using an NK methodology (Kauffman, 1993; Levinthal, 1997; Rivkin, 2000). This formal structure offers a number of critical advantages over other forms of theorizing. First, our theory involves the intersection of three central features of learning models: partially stochastic outcomes, path dependence, and interdependence among the adapting entities. Under these conditions, verbal theorizing is inadequate for the complexity of theorizing, and closed-form analytical models are nearly impossible to implement (Lane, 1993). Second, our theory is multilevel in that it involves adaptation at both the individual and population levels. The agent-based model allows us both to examine individual histories and to understand how these histories aggregate to population-level outcomes. Finally, agent-based simulation methods allow us to integrate existing theory into the model. For example, rather than treat variance across time as exogenous, our formal modeling approach allows differences in variance to emerge as a consequence of organizational choices in a manner consistent with extant theory.

We started with the premise that organizations are engaged in trial-and-error adaptation intended to enhance their performance in the face of selection pressure. We make three central assumptions in constructing our model: (1) the characteristics of the adaptation process are determined by organizations' search strategies, which we treat as fixed over time; (2) selection is myopic, occurring on the basis of performance at one point in the midst of this adaptation process; and (3) the efficacy of selection is assessed based on realized future performance. We began by examining a baseline model in which we assessed the efficacy of selection across organizations employing a very simple, fully blind adaptation strategy. Over the subsequent four experiments, we compared the baseline model to alternative implementations of trial-and-error search strategies that varied in reliability.

SELECTION ACROSS ADAPTIVE ORGANIZATIONS

Our concern is with the efficacy of selection as a mechanism of population-level adaptation. Selection acts by removing inferior organizations from the population, but evolutionary theory does not imply that selection processes lead to optimal outcomes (Winter, 1964). As Simon (1991: 166) noted, ". . . evolutionists sometimes talk about survival of the fittest. But in fact, natural selection only predicts that survivors will

be fit enough, that is, fitter than their losing competitors." While the term "fitter" implies non-random elimination of less-fit organizations, it leaves open a wide range of possible outcomes. Selection may lead to less than optimal outcomes for two reasons. First, there may be errors in selection at any point in time. For example, consider ten organizations, only three of which will survive selection. Ranking the organizations in terms of performance, if selection is perfectly effective, then only the three top organizations should survive. If, in contrast, selection is error prone, such that the top organization fails and those ranked second through fourth survive, then selection will have increased average fitness but would not have been perfectly effective. Second, less salient in this discourse, but perhaps more important in an organizational context, is that the entities over which selection is operating may themselves be changing over time. In part, this lack of attention in the literature may stem from the importance put on stable organizational traits in evolutionary arguments (Cohen et al., 1996). Yet, as the biological analogy suggests, a fixed genetic structure need not imply a fixed phenotype. Consider the obvious case of human development from fragile newborn to robust toddler, sprouting adolescent, and mature and eventually elderly adult. Clearly, selection at intermediate points in this developmental process, even if it is optimal in terms of removing the weakest individuals at a given point in time, may not be optimal over time because it removes slower developing individuals who may eventually exhibit superior attributes. In the context of selection across adaptive organizations, even optimal selection in the short term need not generate optimal long-run outcomes.

The fact that selection forces operate on entities that themselves are undergoing a dynamic process of adaptation raises an important issue about the efficacy of selection in an organizational context. A key factor underlying firms' heterogeneity and, in turn, forming the basis for differential selection, is the fact that organizations engage in different strategies as to how they go about establishing an economically viable entity. It is well recognized that distinct strategies may imply differential long-run performance and, as a consequence, may have a significant impact on organizations' survival prospects. It is less well recognized that distinct strategies, because they imply the possibility that they will unfold differently over time, may also have an impact on organizations' survival prospects independent of long-run performance.

In our modeling, we focus on the implications of organizational adaptation for the efficacy with which selection processes identify organizations that will be superior in the long run and remove those that will be inferior. The existing literature in population ecology and organization theory more broadly has focused significant attention on a related but different issue, the implications of organizational change for the mortality risk of individual firms. In doing so, this literature has identified two broad classes of mortality risks: the first related to risk associated with the content of change and the second related to risk in the process of change. In both cases, organizational change, even if it has long-run positive performance

implications, may have short-run negative performance implications that increase mortality.

Risks in the content of change reflect the problems faced by boundedly rational firms in identifying a suitable destination state, or desired set of outcomes. Change from the existing state to a new state entails a significant mortality risk, as the information required to evaluate the performance consequences of such a change may not be available. Moreover, change to a very distant destination state (exploration) entails more risk than change to an adjacent state (exploitation or local search), as adjacent states are more likely to result in similar performance. Thus change leads to instability because organizations implement changes that may have detrimental performance. The instability associated with risk in the content of change arises because of the difficulty of evaluating alternatives without actually implementing them. At the stable extreme, "off-line" evaluation of alternative destination states (Gavetti and Levinthal, 2000), using cognitive representations to evaluate alternatives without implementing them, reduces content risk by eliminating the worst alternatives prior to exposing them to selection forces. At the unstable extreme, content risk is greater when search is blind, because "specific correct trials are no more likely to occur . . . than specific incorrect trials" (Campbell, 1982: 86). When search is blind, alternatives must be implemented and subjected to selection in order to assess their value. Content risk thus reflects the intelligence embodied by change efforts—the ability (or inability) to evaluate alternatives and select only those that have positive and immediate effects on performance.

Process risks reflect the problems that arise because of the embedded nature of organizational structures. Theory distinguishes between changes in core and peripheral organizational elements (Thompson, 1967; Hannan and Freeman, 1984; Levinthal, 1997). Core elements of structure are, by definition, those that are highly interdependent with other organizational elements. In the microcomputer industry, Romanelli and Tushman (1994: 1147) argued, organizational culture, strategy, structure, power distributions, and control systems constitute core elements. In Silicon Valley start-ups, Baron, Burton, and Hannan (1999) studied founders' "organizational blueprints" as a core organizational feature that dictates the extent of formal organizational structures that serve as control and coordination mechanisms. Although changes in peripheral elements are considered to be feasible, research has called into question the survival benefit of changes in core elements (Singh, House, and Tucker, 1986; Haveman, 1992; Amburgey, Kelly, and Barnett, 1993; Barnett and Carroll, 1995; Dobrev, Kim, and Carroll, 2003). The risk of change in core organizational features arises because change requires significant adjustments across a wide variety of interdependent organizational elements, and this in turn engenders instability. Even small changes in core elements may set off a cascade of related changes that disrupt existing organizational linkages, disturb communication channels across organizational boundaries, and disable routines (Hannan and Freeman, 1984). The risk of mortality is further

increased because content and process issues interact—interdependence across core elements of the organization makes it more difficult to evaluate the performance consequences of a change in content.

One strategy organizations use to manage the problem of the complexities of change is to engage in the change processes sequentially rather than in parallel, making individual changes to organizational content in one domain, while holding all else constant. That said, at higher levels of interdependence, firms may need to make concurrent changes across multiple organizational elements (Rivkin, 2000; Sorenson, 2003), and attempts to do otherwise introduce significant organizational instability. For example, Sorenson (2003), in a study of the effects of the interdependence associated with the co-production of components among computer workstation manufacturers, found that interdependence retards performance and increases mortality in all but high-velocity environments. Ethiraj and Levinthal (2004), in a simulation study of organizational adaptation, demonstrated that reciprocal interdependence (Thompson, 1967) leads to significant instability in the face of local search, which can only be overcome by imposing a strict hierarchy.

To the extent that organizational age (and to a lesser extent, size) is a suitable proxy for routines and elaborated organizational structure, empirical work has documented a correlation between age and mortality in the face of change (Delacroix and Swaminathan, 1991; Kelly and Amburgey, 1991; Amburgey, Kelly, and Barnett, 1993; Miller and Chen, 1994). Highlighting the role of organizational structures and routines, Dobrev, Kim, and Carroll (2003) demonstrated, in a study of the U.S. auto industry, that change requiring a reorientation of structures and the introduction of novel routines increases the risk of failure, but change building on existing structures and routines does not. That said, results in many studies have been mixed (Baum and Amburgey, 2001), in part because age (and size) is a proxy not only for structure but also for accumulated resources. For example, research has highlighted the strategic choices that older and larger firms make to buffer themselves from failure (Miner, Amburgey, and Stearns, 1990; Barron, West, and Hannan, 1994).

The organizations literature in general and the population ecology literature in particular thus suggest that change itself exposes a firm to risk (Hannan and Freeman, 1984; Miner, Amburgey, and Stearns, 1990; Haveman, 1992; Amburgey, Kelly, and Barnett, 1993) arising from the difficulty of identifying a suitable destination state (content risk) and interdependence among elements of the elaborated structure (process risk). The causal chain reflects the argument that content and process issues give rise to reductions in the reliability of performance over time that, in turn, increase the risk of mortality. Still unaddressed, however, is the related question of what the consequences of reduced reliability would be for the efficacy of selection as a driver of population-level adaptation.

If selection is to be effective, it must be able to identify organizations at intermediate points in time that will be superior

in the long run. This challenge reflects the difficulties in estimating future performance. Although selection processes may operate through a variety of mechanisms, inevitably, selection must be based on current and past performance. Even if selection processes such as financial markets are forward looking, information about future performance expectations can be based only on current and past data and outcomes. For instance, the early performance of a new venture-capital-backed technology firm—such as the extent to which it has achieved its early sales milestone—can be used to assess its potential. Similarly, in educational contexts, standardized testing is used to assess students' future educational performance. In both cases, the intermediate manifestations of performance before the adaptation process is complete may provide only a noisy signal about long-term performance. More generally, it is unclear to what extent the intermediate performance manifestations of adapting organizations provide a good signal about future performance.

In our modeling, we focus on how heterogeneity across organizations in reliability during the adaptation process affects the quality of the information on which selection acts and, in turn, the efficacy of selection. We illustrate how the information content of current performance as it applies to predicting future performance decreases as the instability with which adaptation proceeds increases. The structure of the model closely reflects its theoretical foundation in the population ecology and organizations literature. Organizations engage in a process of local search for alternatives (Cyert and March, 1963; Nelson and Winter, 1982; Stuart and Podolny, 1996). In such a process, the set of potential destination states is limited to those near to the origin state, and the correlation in performance over time is high relative to exploratory search. But even in local search processes, organizational search choices affect the adaptive dynamics. Building on the prior literature discussed above, the reliability of adaptation is, in part, jointly the result of differences across firms in how search is pursued with regard to (1) the mode by which alternatives are evaluated, which affects the degree of content risk, and (2) the imposition of organizational structure (integrated vs. functional), which affects the degree of process risk. The imposition of structure may be further refined in terms of how effort is allocated over time across elements of structure (parallel vs. sequential allocation of effort). Taken together, we denote these choices collectively as a search strategy, a coherent set of rules governing learning. We have chosen to model the mode of evaluation and the nature of organizational structures because they closely reflect the theoretical foundations of the population ecology literature's focus on reliability, although other elements of search strategies, not modeled here, may also contribute to differential reliability.

First, evaluation to assess the performance impact of an alternative may be conducted on-line or off-line. For on-line evaluation, an organization must implement that alternative in order to assess its implications for performance. As such, on-line evaluation reflects the content risk associated with the difficulties embodied in identifying a destination state for

change. Off-line evaluation allows an organization to implement only those alternatives that are believed to lead to performance improvements. Off-line evaluation decreases the content risk of mortality but does not necessarily eliminate it, because the a priori screening of alternatives may well be imperfect. Although, on average, both methods of evaluation allow organizations to proceed in the direction that leads to marginal performance gains, off-line evaluation does so with greater stability.
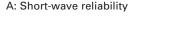
Second, search strategies can differ in how organizations structure the problem-solving search effort. We considered an underlying task structure in which the problem facing organizations can be divided into two sub-problems, such as research and development (R&D) and marketing, or domestic and international business. Given the important role of organizational structure in structural inertia theory, we examined how different structures affect the degree of reliability over time during adaptation. The simplest organizational structure is fully integrated. In this structure, the organization does not treat the problem as decomposable. For example, there are no separate marketing and R&D departments. Instead, search is structured in an integrated manner across both sub-problems. In this sense, firms search across one large problem landscape. Alternatively, organizations may employ a functional structure consisting of two functional units, such as separate R&D and marketing departments. Each functional unit focuses on a particular sub-problem and as such, firms engage in local search on two distinct sub-problem landscapes.
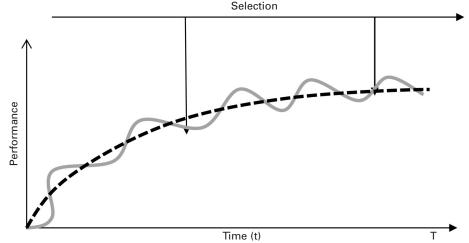
If an organization has a functional structure, a set of search rules governing the allocation of resources across the functional units must be considered. Two fundamentally different search strategies are possible: parallel and sequential. In parallel search strategies, the organization allocates its search effort evenly across both units. In sequential search strategies, it first devotes all search effort to one unit, and only starts searching across the other after it has made a certain degree of progress on the first. As an example of sequential search, entrepreneurial firms in the technology sectors often address the problem of marketing only after technological search has yielded a working prototype. Parallel search enables integrated trial-and-error learning across the entire system, which tends to reveal problems more rapidly and, in doing so, stimulates learning. At the same time, this tight-coupling approach tends to decrease the rate of early performance improvement. This problem arises as adaptation in one domain, which improves performance locally, negatively affects performance in the other domains and thus leads to remedial action. In contrast, sequential search purposefully suppresses learning in one domain in order to enhance learning in another (Baldwin and Clark, 2000), producing superior early performance gains in one unit at the expense of reduced stability over time (i.e., a non-monotonic learning curve).

The challenge to the efficacy of selection across organizations in the midst of adaptation is identifying the organizations that will have superior long-run performance. The selection
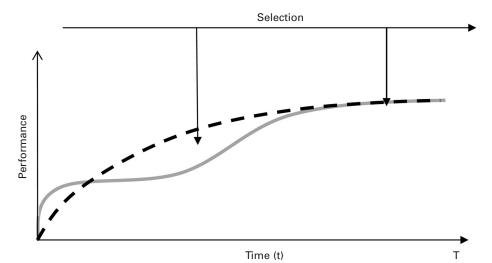
process is significantly confounded by organizational heterogeneity in search strategies that affect both long-run performance and the reliability of performance over time as the organization moves toward the long-run performance outcome. In particular, the path of adaptation over time can be characterized by two basic types of reliability exhibited by organizational learning curves: short-wave reliability and long-wave reliability, as depicted by the grey learning curves in figure 1. While short-wave reliability (figure 1, A) reflects near-term performance cycling, long-wave reliability (figure 1, B) reflects non-monotonicity in performance improvement

Figure 1. Learning curves and intertemporal correlation in performance.

A: Short-wave reliability

B: Long-wave reliability

such that the rate of performance improvement varies considerably over time. To the extent that adaptation is unreliable, selection is less effective at identifying organizations that will be superior in the long run. In this sense, the efficacy of selection is endogenous to the organizational search

strategies that drive differential adaptation trajectories. In our computational model of adaptation, we examine systematic differences in reliability across organizations and the effect of those differences on the information available in the selection process.

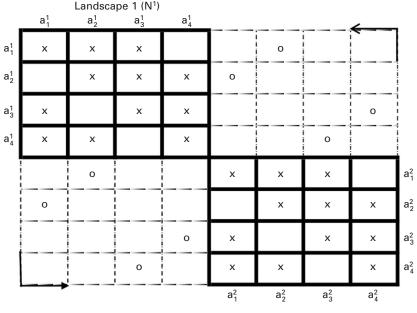## COMPUTATIONAL MODEL OF ADAPTATION IN THE PRESENCE OF INTERMEDIATE SELECTION

We developed an agent-based simulation model to examine the efficacy of selection across heterogeneous organizations undergoing adaptation. We built our model on two central assumptions, discussed above: (1) that adaptation is the result of local search among policy choices guided by search strategies, and (2) selection occurs on the basis of current performance outcomes. We employed as an analytical tool the NK model developed by Kauffman (1993), which has found growing acceptance in the management literature in studies of search (Levinthal, 1997), imitation (Rivkin, 2001), organizational structure (Rivkin and Siggelkow, 2003), and organizational architecture (Ethiraj and Levinthal, 2004). The basic model consists of a structure for defining a fitness landscape (the problem space to be searched), search rules that guide adaptation, and a selection mechanism. Each of these elements is briefly characterized below and a detailed mathematical treatment is included in the Appendix.

Landscape Structure and Fitness Functions

Our primary objective was to explore how the process of adaptation affects the efficacy of selection. We modeled a task environment consisting of two interconnected sub-problem landscapes (e.g., R&D and marketing; domestic and international activity, etc.). The payoff to a set of choices in one problem domain, or landscape, depends on the choices made in the other. We assumed that the boundaries of these two landscapes, along with the level of interdependencies between them, were defined exogenously and thus were not subject to managerial discretion in the short term. Our modeling of this task environment is related to recent efforts to study issues of organizational design (Rivkin and Siggelkow, 2003; Siggelkow and Levinthal, 2003; Ethiraj and Levinthal, 2004), in which interactions across sub-problems (landscapes) play a central role.

A set of choices about how an organization wishes to compete and operate is represented as a location in policy space characterized by a vector of policy choices, $N^1 = (a_1^1, \ldots, a_{n^1}^1)$, and $N^2 = (a_1^2, \ldots, a_{n^2}^2)$, of lengths n^1 and n^2, respectively, where the superscript represents the landscape (sub-problem), and the subscript represents the policy choice. We let each policy choice correspond to a particular binary decision variable, where, for example, the choice of 1 might imply the choice of using flash memory rather than a disk drive in the design of a new portable music player. Consider the example in figure 2. An organization's location on each landscape is represented by vectors of length $n^1 = n^2 = 4$, such that $N^1 = a_1^1 a_2^1 a_3^1 a_4^1$ and $N^2 = a_1^2 a_2^2 a_3^2 a_4^2$. Thus there are $2^{n^1} = 2^{n^2} = 2^4$ possible combinations of policy choices in each space. The firm's position in the overall landscape can be represented by the

Figure 2. Example of a set of firm choices across two landscapes (N^1 and N^2) and the interactions between them.*

*The x's represent interactions across choices within a landscape (sub-problem) while the o's represent interactions across landscapes. For example, the payoff to choice a_1^1 is a function of the choices for a_3^1, a_4^1, but also a_2^2.

vector N of length $n = n^1 + n^2$, such that $N = \{N^1, N^2\} = (a_1^1, \ldots, a_{n^1}^1, a_1^2, \ldots, a_{n^2}^2)$.

Interactions between policy choices can occur both within a landscape and across landscapes (Simon, 1962). As an example, a within-sub-problem interaction in R&D would be the interaction between a laptop's weight and screen size, while an across-sub-problem interaction might occur between R&D and marketing in the form of an interaction between a laptop's weight and choice of consumer- versus business-sales channels. Interactions within a landscape are characterized by the vector K_j^W (length k^W) of policy choices with which policy a_j^1 (or a_j^2) interacts on its own landscape (where j indexes the policy choice). The example in figure 2 demonstrates the interaction space in which $k^W = 2$ (landscape 1 in the upper-left quadrant and landscape 2 in the lower right, which represent two sub-problems). In the example, the payoff to policy a_1^1 (choice of 0 or 1) in landscape 1 is a function of the choices made for the policies in the vector $K_1^W = (a_2^1, a_3^1)$, this relationship between policies is characterized by an "x" in the appropriate cell. In the example in figure 2, if we let $N^1 = a_1^1 a_2^1 a_3^1 a_4^1 = 1110$, then $K_1^W = (1,1)$, and the fitness contribution of policy choice a_1^1 is a function of selecting a 1 for bit a_1^1, as well as selecting $a_2^1 = 1$ and $a_3^1 = 1$. Changing to $a_3^1 = 0$ would alter the fitness contribution made by the choice of $a_1^1 = 1$, while the policy choice setting for $a_4^1 = 1$ or 0 has no effect on the fitness contribution of a_1^1.

Interactions across landscapes are characterized by the vector K_j^B (length k^B) of policy choices in the other landscape that

affect the performance contribution of a policy choice in the focal landscape (the "o"s in the upper right and lower left). That is, interdependencies across activities on the two landscapes complicate the learning process as search activity in one domain both informs and affects progress in the other. In the example, with $k^B = 1$, the fitness contribution of a_1^1 is also a function of the choices made for the policies in the vector $K_1^B = (a_2^2)$. In total, the payoff associated with an individual policy choice is conditional on all policy choices (within and across landscapes) with which it interacts. In our example, the payoff associated with policy choice 1 in landscape 1, a_1^1, is $\pi_1^1(a_1^1 | K_1^W, K_1^B) = \pi_1^1 (a_1^1 | a_2^1, a_3^1, a_2^2)$. The payoffs to each of the $2^{1+k^W+k^B}$ distinct vectors associated with a policy choice are generated by a draw from a uniform distribution $u[0,1]$.

The fitness associated with a position on landscape 1, N^1, is simply the average of the fitness contributions of each of the individual policy choices, such that

$$\Pi_j^1 = \frac{1}{n^1} \sum_{j=1}^{n^1} \pi_j^1(a_j^1 | K_j^W, K_j^B).$$

There is a analogous expression for the fitness associated with N^2. The total fitness is the average fitness across landscapes, such that $\Pi_j = \frac{1}{2}(\Pi_j^1 + \Pi_j^2)$.

Rules Governing Search

Search takes the form of local trial-and-error learning, which corresponds to "hill climbing" in a multiple-peaked performance landscape (Levinthal, 1997). Search is performed by evaluating the consequences for fitness of changing a randomly selected policy choice a_j^L where L is landscape 1 or 2. Search strategies can differ both in how alternatives are evaluated and in the imposition of organizational structure.

Evaluation reflects the mechanism by which the alternatives identified in the trial-and-error search process are assessed and, in particular, the difference between blind and non-blind search. On-line evaluation reflects a context in which it is not possible to determine the value of a change in a single policy element without implementing the change (Gavetti and Levinthal, 2000). Thus "specific correct trials are no more likely to occur . . . than specific incorrect trials" (Campbell, 1982: 86). As a consequence, fitness at the end of the period may be either higher or lower than at the beginning of the period. If fitness is lower, then in the next period, the organization returns to its earlier position and attempts to test another policy choice. Once all single policy choice changes have been examined and found to be inferior, search ceases. In contrast, off-line evaluation is such that, if performance is improved by making the change, it is implemented (Gavetti and Levinthal, 2000), otherwise the potential policy change is discarded. As such, off-line evaluation entails search in which incorrect trials are not subjected to population selection, and as such, off-line search is non-blind. Off-line evaluation is possible when the efficacy of a choice can be determined using such things as cognitive evaluation, computer models, structured experiments, animal models, or prototypes. For

instance, a particular combination of chemicals may be evaluated in vivo in a laboratory, without having to test it directly by selling the chemical in the product market.

Organizational structure reflects differences in the set of constraints imposed on the search process. We implemented two alternative structures. The first structure is integrated. In this structure, an organization chooses a policy choice to evaluate from the entire set of $N = \{N^1, N^2\} = (a_1^1, \ldots, a_{n^1}^1, a_1^2, \ldots, a_{n^2}^2)$ policies, with evaluation based on the change in fitness across the global landscape. The second structure is functional, in which the organization chooses a multidivisional structure that matches that of the task environment. Thus the organization is divided into two functional subunits, with search constrained to within-unit policy changes. When search is conducted in unit 1, only policy choices $N^1 = a_1^1 a_2^1 a_3^1 a_4^1$ may be examined, and only the consequences for local unit 1's performance are evaluated. As noted earlier, the functional structure entails two further sets of choices. Effort can be allocated in parallel or sequentially. In parallel allocation, in each period, the organization attempts one search effort in each functional unit. In sequential allocation, in the early periods (prior to period 30), the organization allocates all resources to search in the first landscape, and only in later periods does search begin in the second landscape.

Figure 3 highlights a typology of search strategies that arise at the intersection of modes of evaluation and the imposition of organizational structure. We implemented four types of search strategies covering four of the quadrants (and sub-quadrants) of figure 3: (Q1) *Baseline:* an integrated structure with on-line evaluation—which represents the simplest search strategy; (Q2) *Integrated:* an integrated structure with off-line evaluation; (Q3a) *Parallel:* a functional structure with parallel allocation of effort and off-line evaluation; and finally (Q3b) *Sequential:* a functional structure with sequential allocation of effort and off-line evaluation.

Figure 3. Evaluation mode versus organizational structure: A matrix of alternative strategies.

Selection

Because we were interested in the efficiency of selection, we modeled adaptation and selection (organizational exit) without considering the entry of replacement organizations.[1] Selection regimes vary in their intensity—weak versus strong. Selection was implemented as the outcome of a stochastic process such that an organization's probability of survival is a function of its fitness relative to the fitness of the best firm at a given point in time (Wilson and Bossert, 1971). We implemented a proportional selection model, such that the probability of survival at a given point in time is $s_{i,t} = (\Pi_{i,t}/\mathrm{Max}(\Pi_t))^{Z_t}$, where $\Pi_{i,t}$ is the fitness of organization i at time t, $\mathrm{Max}(\Pi_t)$ is the maximum fitness achieved by any firm at time t; and Z_t is the strength of selection. The proportional selection model imposes competitive consequences on organizations because the probability of failure for a given organization increases with fitness of the leading organization in the population at any point in time.[2]

SIMULATIONS

We used the fitness landscape described above as an analytical tool to analyze the efficacy of selection across alternative search strategies. All simulations are for an n = 20 landscape consisting of two sub-problems where each sub-problem corresponds to a single landscape with $n^1 = n^2 = 10$ interdependent policy choices. This in turn generates approximately one million (2^{20}) policy choice combinations represented by unique locations on the combined landscape. The interaction structure within each sub-landscape is specified as $k^W = 5$, with $k^B = 2$ interactions across landscapes. A simulation run lasts for 100 periods and includes 200 organizations seeded randomly on the landscape. In each period, an organization is allocated two search attempts. All results are averaged over 100 runs, and thus we evaluated the performance of a total of 20,000 organizations in each simulation. Our analysis proceeded in four experiments. In the first experiment, we examined the dynamics of performance improvement and, in particular, the reliability of adaptation across alternative search strategies. In the second experiment, we examined the efficacy of continuous selection (on the basis of current performance) within populations of organizations that were homogeneous in terms of search strategies. In doing so, we assessed how different strategies affected the ability of selection to act as a mechanism of population-level adaptation, to determine if there are differences across strategies in selectability—the information content of current performance that allows selection to screen out organizations in a manner corresponding to their long-run performance. In the third experiment, we examined alternative selection regimes, first by enabling selection on the basis of past as well as current performance and, second, by implementing discrete selection at a single point in time. In the final experiment, we allowed organizations pursuing different search strategies to compete and examined the implications of differences across strategies in selectability for the demographics of surviving firms.
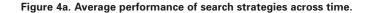
Experiment 1: Reliability of Alternative Search Strategies

We began by assessing the performance trajectories of alternative search strategies in the absence of selection. We did so by running independent simulations for each of the four distinct search strategies in quadrants 1 through 3 in figure 3 above. Results are presented in figures 4a–4c: (a) average performance across time; (b) cross-sectional variation (standard deviation of fitness across firms) within the population of organizations following a given strategy; and (c) intertemporal short wave reliability measured as $Corr(F_{t-1}, F_t)$. The central performance measure of fitness in figure 4a shows that although the rate of performance increase differs markedly across the strategies, all four strategies appear to generate a common asymptote with an average fitness of .69. In fact, the long-run results are not quite identical, although they are indistinguishable at the three-digit level. In particular, functional parallel strategies, by generating perturbations in one landscape that have an impact on search in the other landscape, periodically enable firms to escape a poor basin of attraction for a basin with a superior local peak. But these differences in long-run performance across strategies are only noticeable at higher levels of cross-landscape interdependence than was implemented in this paper and have no bearing on our results.

All strategies use the same amount of effort in the search process. Thus differences in the rate of progress can only come from differences in search-strategy performance characteristics driven by the mode of evaluation and the imposition of organizational structure. The central qualitative observation is that the baseline and sequential strategies not only exhibit slower rates of progress but also entail far greater cross-sectional variability, as shown in figure 4b. For the baseline strategy, these observed results can be attributed to the on-line evaluation, which results in some fraction of alternatives generating performance decrements. These errors lead to a significant reduction in short-wave reliability that both decreases the rate of performance improvement and increases cross-sectional variation, particularly between periods 15 and 25, as shown in figure 4c.

Less obvious is the sharp contrast in behavior among the three search strategies that entail off-line evaluation of alternatives (Quadrants 2 and 3 in figure 3). Integrated search, which differs from the baseline strategy because it implements only performance-enhancing alternatives, proceeds with greater reliability and, as a consequence, more rapidly than baseline strategy search. The two functional search strategies, parallel and sequential, match the underlying task structure by imposing organizational structure and decomposing the problem across two functional units. For example, while marketing and R&D functions may be interconnected in terms of solutions, firms typically decompose the problem across two functional units. Functional search strategies ignore the between-unit interdependence in activities and outcomes because doing so reduces coordination costs.

Despite this commonality, the dynamics of performance of parallel and sequential search strategies are markedly differ-
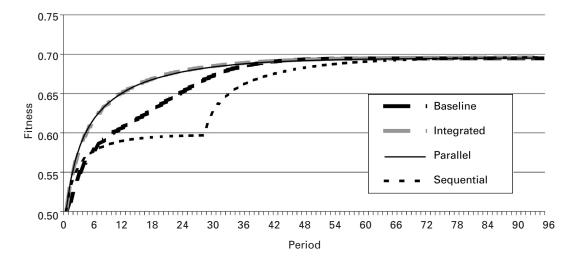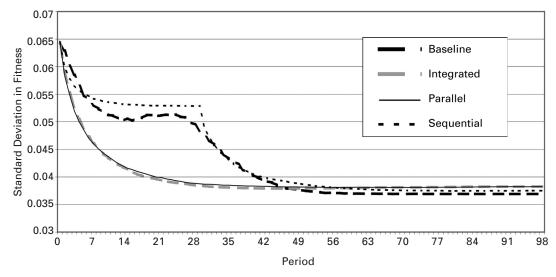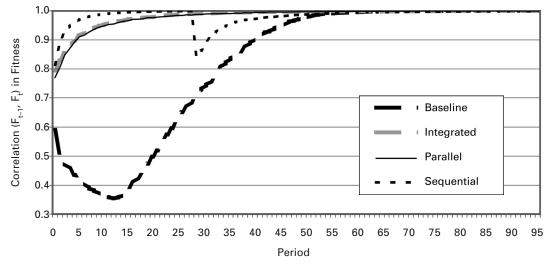
Figure 4a. Average performance of search strategies across time.

Figure 4b. Cross-sectional variation (S.D. across organizations).

Figure 4c. Short-wave reliability: Correlation in performance across time.

ent. The parallel-strategy firms exhibit performance that closely matches that of the integrated-strategy firms. In contrast, sequential-strategy firms, by focusing on only one landscape in the early periods, limit the potential for performance improvement. In particular, the potential improvement is on average half that of the potential improvement that a parallel-strategy firm can expect. Sequential-strategy firms, however, exhibit even higher levels of short-wave reliability in early periods by allocating all their search effort to one sub-problem landscape. This higher early reliability leads to rapid early performance gains that are quickly offset by the limited room for improvement in the landscape in which search efforts are initially focused. In addition, this decoupling of search introduces a second issue: the adaptation trajectory of sequential-strategy firms is highly non-monotonic. After period 30, when sequential organizations commence search in the second landscape, fitness improves dramatically, cross-sectional variation decreases, and short-wave reliability returns to pre-period 30 levels.

All search strategies examined in the first experiment led to adaptation that converged to a common long-run performance. Nevertheless, differences across search strategies in the mode of evaluation of alternatives and the implementation of organizational structure gave rise to significant differences in the level of short-wave reliability. This property, in turn, drove both the rate of performance improvement and the level of cross-sectional variation.
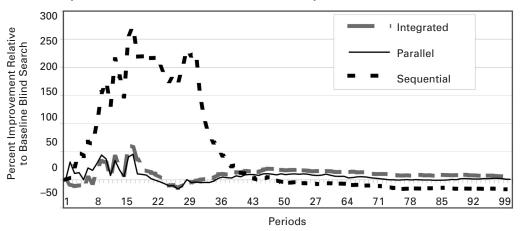
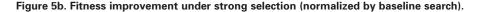### Experiment 2: Implications of Reliability for the Efficacy of Selection

In the second experiment, we examined the consequences of differences in reliability for the efficacy of selection as a mechanism of population-level adaptation, to assess the extent to which selection increases average population-level performance across the different search strategies by eliminating inferior firms. Our proposition is that selection will be more effective at identifying superior firms when those firms pursue search strategies that exhibit higher levels of reliability in the adaptation process.
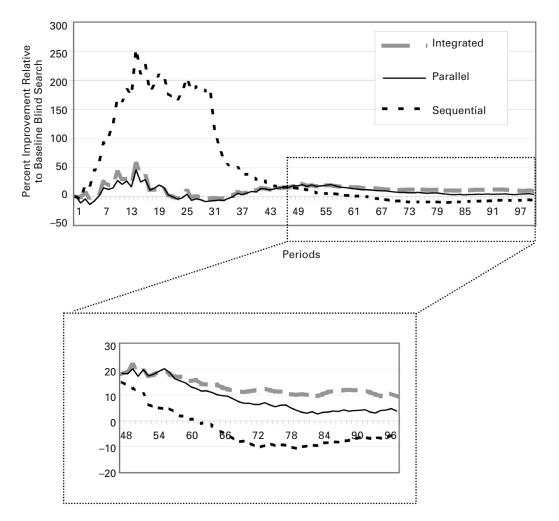
We ran this experiment under two different selection regimes: (1) weak selection, which removes approximately 50 percent of the population over the 100-period experiment, and (2) strong selection, which eliminates approximately 95 percent of firms over 100 periods. The performance variable of interest was the average improvement in population-level fitness generated by selection, as it removes organizations deemed to be inferior, normalized by the baseline adaptation results. The results are presented in figures 5a (weak selection) and 5b (strong selection). The figures show that strong selection over a population of integrated-search firms increases fitness 9.34 percent more than does such selection across a population of firms pursuing the baseline adaptation model. Because the results are consistent across both regimes (although less intense in the weak selection regime), we focus our discussion on the strong selection regime.

Under the strong selection regime, both integrated and parallel search strategies outperform the baseline search strategy

Figure 5a. Fitness improvement under weak selection (normalized by baseline search).

Figure 5b. Fitness improvement under strong selection (normalized by baseline search).

in terms of long-run fitness improvement, by 9.34 and 3.85 percent, respectively, suggesting that selection is more effective at increasing population-level performance across either integrated or parallel search than the baseline search strategy. In addition, the difference in improvement between inte-

grated and parallel search is large, given that the results from experiment 1 showed only relatively small differences in short-wave reliability. This highlights an as yet unexplored cost of decomposing problems into constituent parts. Even small differences in reliability that arise through the imposition of organizational structure can alter the ability of selection to facilitate effective population-level adaptation.

The results for the sequential search strategy are more complex. Selection-driven performance improvements are an order of magnitude larger for sequential search in the early periods, reaching a peak of 250 percent of the rate exhibited by the baseline search strategy, but this dissipates rapidly with the activation of search efforts related to the second sub-problem (landscape 2). In the long run, selection across firms pursuing a sequential search strategy is less effective than equivalent selection across the baseline blind-search model by 6.7 percent.

To assess the significance of these differences in long-run performance contingent on surviving selection, we conducted unpaired t-tests of the mean performance across each strategy at period 100. The results of this test for the strong selection regime are presented in table 1. Each data point in the table contains the difference in fitness at period 100 as the column less the row. Thus after selection, the average fitness of a population of firms pursuing the integrated search strategy is .0047 greater than that of a population of firms pursuing the baseline strategy, a difference that is highly significant at the .001 level. There are two central results from this table. First, selection is significantly less effective across firms pursuing the baseline blind search strategy than either the integrated or the parallel strategies. Second, selection is significantly less effective across the sequential strategy than all the other search strategies examined.

To examine why selection is more or less effective across some search strategies, we conducted two further sets of analyses. First, we calculated the selection-error rate observed at period 100 to determine how many of the firms (out of 200) that survived selection were actually in the top 100 firms in period 100. The results of this analysis, which are presented in figure 6, are read as follows: in the weak selection regime with firms pursuing the baseline search

Table 1

Difference in Mean Fitness at Period 100 under a Strong Selection Regime*

Strategy	Baseline	Integrated	Parallel
Integrated	.0047••••	—	—
	(.0013)	—	—
Parallel	.0027•	–.0020	—
	(.0015)	(.0015)	—
Sequential	–.0035•••	–.0081••••	–.0061••••
	(.0015)	(.0015)	(.0017)

• $p < .10$; •• $p < .05$; ••• $p < .01$; •••• $p \leq .001$.
* Standard errors are in parentheses. The difference in fitness is the average row fitness less the average column fitness, with a positive difference indicating greater row fitness.

strategy, of the 102 firms that survive selection, 41 of the surviving firms were not in the top 102 of the 200 firms, thus generating a 39-percent error rate. In general, the results of the error-rate analysis are consistent with the performance-improvement examination. A shift from the baseline search strategy to integrated search leads to a decrease in the error rate from 71 percent to 65 percent, with the error rate increasing to 66 percent for parallel search, consistent with the long-term performance differences exhibited in figure 5b. The error rate across sequential search was the highest observed, at 77 percent.

Thus the efficacy of selection as a mechanism of population-level adaptation, highlighted by differences in population fitness under selection, is driven by the extent to which the selection process is error prone, removing organizations in earlier periods that would have gone on to do well if they had survived. In particular, for all search strategies except sequential search, there appears to be a strong relationship between the extent of short-wave reliability and selection-error rates. As short-wave reliability increases, error rates appear to decrease, and selection becomes more effective.

Second, we examined the nature of information available for selection at a given point in time, focusing on how heterogeneity across organizations in the reliability of adaptation affects the quality of the information on which selection acts. The challenge over time for selection is in estimating future performance. To what extent do the intermediate performance manifestations of adapting organizations provide a good signal about future performance? If differences across search strategies affect the degree of reliability of adaptation, then this may distort the information provided by contemporaneous performance, rendering it less informative about long-run performance and in turn limiting the effectiveness of intermediate selection.

To examine this issue, we measured long-wave reliability by calculating the extent to which fitness at a given point in time is correlated with long-run fitness, $Corr(F_t, F_{100})$, which provides a direct measure of the informativeness of intermediate performance levels. For example, if the correlation between fitness at period 30 and 100 is 100 percent, then fitness at period 30 is perfectly informative about long-run fitness. If, in contrast, the correlation is only 5 percent, then intermediate fitness provides little useful information about long-run fitness. The results of this analysis are shown in figure 7, which plots the correlation in fitness between period t and period 100 across each of the search strategies. For example, in period 15, the correlation between contemporaneous and long-term fitness for the integrated strategy is approximately 52 percent—that is, fitness at period 15 explains 52 percent of the variance in long-term performance. In contrast, for the baseline blind-search strategy model, the correlation is only 36 percent. For the sequential strategy, the correlation drops further to 27 percent.

For the baseline, parallel, and sequential strategies, the short-wave and long-wave reliability results are broadly similar in terms of rank ordering. In contrast, for the sequential strate-

Figure 6. Selection error rates across search strategies at period 100.
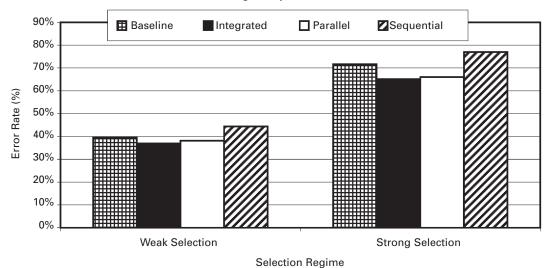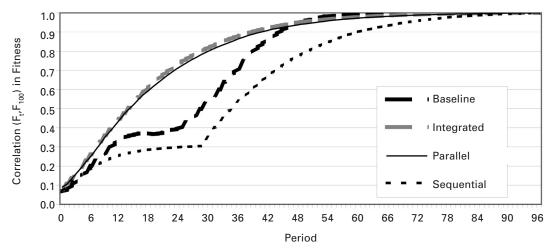

Figure 7. Long-wave reliability: The information content in contemporaneous performance.

gy, while the short-wave reliability is superior to that of the baseline strategy, the long-wave reliability is greatly inferior. This reversal stems from the temporal decomposition of the search effort in the sequential strategy. Because the two sub-problems (landscapes) are interdependent, the onset of search in the second landscape in period 30 makes obsolete much of the firm-level learning in the first landscape. As a consequence, the selective removal of inferior sequential-strategy firms prior to period 30 is much more likely to be in error. The early performance improvement generated by selection across sequential-strategy firms, shown in figure 5b above, leads to rapid improvement in the population's performance and is consistent with the very high level of short-wave reliability that this strategy demonstrates at the outset of the simulation. But this advantage is quickly overwhelmed by the cost associated with a significant decrease in long-wave reliability. In the long run, this reliability deficit leads to selection errors that decrease the ability of selection to act as a mechanism of population-level adaptation.

Finally, we examined the extent to which our results were sensitive to differences across strategies in cross-sectional variation. We implemented a rank-order selection regime that was insensitive to cross-sectional variation. Our results (available from the authors) were robust to this alternative specification. Differences in cross-sectional variation do affect the efficacy of selection, but when there are also substantial differences in reliability across strategies, differences in cross-sectional variation will tend to be swamped by the contribution of these differences in reliability over time.

Experiment 3: Efficacy of Alternative Selection Regimes

The previous experiment demonstrated that the efficacy of selection varied significantly across search strategies as a function of the reliability of adaptive efforts. In that experiment, selection was continuous, occurring in each period, and was implemented on the basis of contemporaneous performance alone. In the third experiment, we explored the efficacy of two alternative selection regimes that, although arbitrary in their construction, reflect important dimensions of selection regimes observed in practice: (1) continuous selection with the use of historical information and (2) discrete selection at period 15.

We implemented selection on the basis of historical, rather than simply current performance. We used twenty periods of historical data and ran the simulation models from experiment 2 to ascertain if more information enhances the efficacy of selection. As an indication of the efficacy of selection, we considered the error-rate results, which are reported in figure 8. The height of the bars represents the percentage of organizations that survive selection erroneously. The numbers above the bars represent the comparison with the results from experiment 2 graphed in figures 5a and 5b. Thus, for the baseline blind-search strategy, the use of historical data under a strong selection regime led to an error
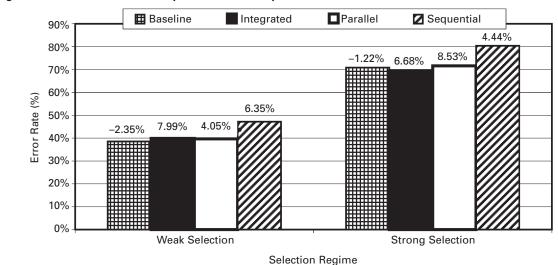
Figure 8. Selection error rate with performance history.*

* Numbers above the columns represent an increase or decrease in the error rate from the continuous selection mode results graphed in figure 6.

rate of 71 percent, which is a reduction of 1.22 percent relative to the error rate observed in experiment 2 (e.g., current performance only). In contrast, for the integrated, parallel, and sequential strategies, the error rate increased by 6.68, 8.53, and 4.44 percent, respectively.

The results are informative in two respects. First, for the more reliable strategies (integrated, parallel), historical data are less useful predictors of future success than is current performance, because performance at any point in time for these monotonic strategies provides information on the lower bound of the local peak, and therefore historical performance data from points lower down the peak is less informative about future performance than is current performance. Thus the counterintuitive result is that including historical data diminishes the efficacy of selection by increasing error rates. Second, for the less reliable strategies, which exhibit non-monotonic performance gains, the results diverge. The baseline strategy suffers from a deficit of short-wave reliability. Because adaptation is blind, in any given period, the organization may implement a trial that reduces performance. Here, the use of historical data, which in effect averages backward, captures prior-period results that were at times superior (higher up the local peak) to the current period, thus providing better information on which selection could act and resulting in a lower error rate. In contrast, sequential search suffers from a deficit in long-wave reliability. In this case, the use of historical data implies that after period 30 (the onset of landscape 2 search), pre-period 30 data are still considered, and as a consequence, the error rate increases.

Discrete selection represents a regime with intermittent rather than continuous selection and is in this sense similar to that of a venture capital regime. We examined the outcome of a discrete selection event at period 15. Error-rate results on the efficacy of selection are reported in figure 9.
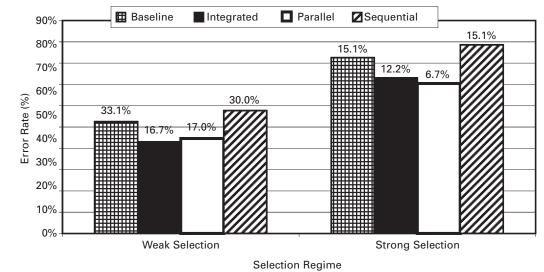
Figure 9. Discrete selection error rate with performance history.*

* Numbers above the columns represent an increase or decrease in the error rate from the continuous selection mode results graphed in figure 6.

The main result is that discrete selection exhibits a higher selection-error rate than does continuous selection. Under the weak selection regime, the error rate increase is between 16.7 and 33.1 percent. Under the strong selection regime, the increase in error rates ranges from 6.7 to 15.1 percent. Two observations are of note. First, error rates grow faster under the weak selection regime. Second, the implementation of discrete selection, rather than continuous selection, comes at a cost that differs across search strategies. In particular, although the basic pattern of the efficacy of selection across alternative search strategies remains constant, the error rate increases faster across the less reliable baseline and sequential strategies than it does across the more reliable integrated and parallel strategies. The largest increase in error rate occurs in selection over the sequential strategy. The efficacy of the results from discrete selection, however, will be partly a function of the time at which selection is implemented. Because our interest was in modeling a selection regime similar to that of venture capital, we examined early selection, but as the time of discrete selection moves later, the efficacy of selection will improve. At the extreme, with selection at period 100, selection will simply be proportionate to long-run performance.

The results of experiment 3 suggest that the general ranking of alternative strategies in terms of selectability remains relatively unchanged under alternative selection regimes. As one would expect, a decrease in information on which selection acts (discrete selection) reduces the efficacy of selection, though this reduction in efficacy is less severe in a strong selection regime. Counterintuitively, the results suggest that in most cases, the use of more information (historical selection) also serves to reduce the efficacy of selection.

Experiment 4: Competitive Implications of Differences in Selectability

The previous three experiments demonstrated that the efficacy of selection varies significantly across search strategies as a function of the reliability of adaptation. In experiment 4, we examined the consequences of differences across strategies in selectability when two organizational forms (i.e., sets of organizations pursuing different search strategies) are competing for dominance in a population. Each simulation involved 100 runs of the model with 200 firms split between the two competing search strategy types.

Two central measures that are of interest are the relative proportion of survivors of each type of organization and the probability that a given type of organization will be the winning organization. We ran six separate simulations (three each under the weak and strong selection regimes) in which populations of two types of firms compete: (a) baseline vs. integrated, (b) baseline vs. parallel, and (c) baseline vs. sequential. In the absence of selection, all four strategies (Quadrants 1 through 3 in figure 3) generate identical long-run performance. Contingent on surviving selection, however, this equifinality disappears. Figure 10 illustrates the population demographics results, the proportion of survivors in period 100 that are non-baseline-strategy firms. Figure 11 illustrates the win-
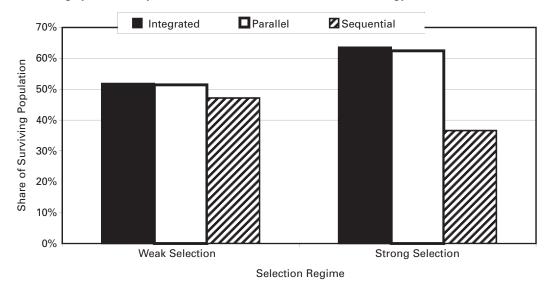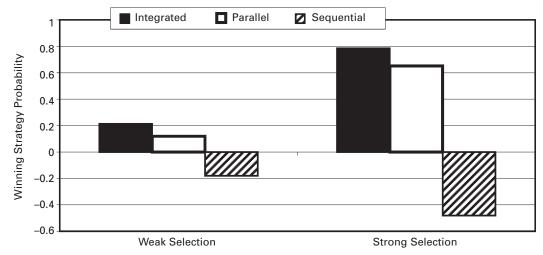
Figure 10. Demographics of competitive outcomes (versus baseline search strategy).

Figure 11. Winner of competitive outcomes (versus baseline search strategy).

ner of the competition, graphing the probability that a non-baseline-strategy firm has (or is tied for) the highest performance at period 100. On the basis of both performance metrics, the integrated and parallel strategies generate superior outcomes in competition with the baseline strategy. For example, under strong selection, the integrated-strategy firms come to account for 63 percent of the population of surviving firms and an integrated-strategy firm exhibits the highest performance amongst the survivors nearly 80 percent of the time. In contrast, the sequential-strategy firms lose to baseline-strategy firms on both performance metrics. The results suggest that differences in the efficacy of selection (driven by differential reliability) across strategies have direct implications for the outcome of competition between alternative organizational forms. Because selection is more error prone within the populations of less reliable search strategies, many of the latently superior members of such sub-pop-

ulations are mistakenly eliminated, significantly reducing the likelihood that one such organization will be the winner. As a consequence, organizations pursuing search strategies that are more reliable, exhibiting superior selectability, come to dominate the population of surviving organizations. In this sense, reliability and survival are correlated even when selection occurs on the basis of performance. This final experiment highlights differential selectability, driven by differences in the reliability of adaptive efforts, as an important factor in determining the outcome of competition in evolutionary systems in which selection occurs across organizations that are themselves undergoing adaptation.

DISCUSSION AND CONCLUSION

Evolutionary perspectives on firms and industries have become quite prominent in the strategy and organizations literatures (Nelson and Winter, 1982; Hannan and Freeman, 1984; Baum and Singh, 1994; Burgelman, 1996). Central to the evolutionary argument is the presence of some process of competitive exclusion and differential selection. Furthermore, an important strand of the strategy and organizations literature has taken a firm-centered view of evolutionary dynamics in analyzing the path-dependent processes of firm-level adaptation and learning (cf. Helfat, 2003). Although there has been some recognition of the need to consider the interrelationship between the processes of adaptation and selection (Levinthal, 1991a), and some important steps have been taken in that direction (March and Shapira, 1992; Barnett, Greve, and Park, 1994; Barnett and Hansen, 1996; Denrell and March, 2001; Greve, 2002), the micro dynamics of the two processes have only begun to be explored. Our objective in this paper was to use a simulation to explore the efficacy of selection as a mechanism of population-level adaptation when individual organizations themselves are adapting. Our central finding from the simulations is that the efficacy with which selection discriminates among adaptive organizations, and as a consequence, the ability of selection to drive population-level adaptation, is a function of the process of organizational change itself. Search strategies that lead to learning outcomes that exhibit high reliability are more selectable in that selection is less prone to error when acting across firms that are pursuing highly reliable strategies.

We adopted the assumption of inertia with respect to organizations' "genetic structure," but this premise does not negate the possibility, or even the likelihood, of profound changes in the observed properties of an organization, what a biologist would refer to as its phenotype, and, as a consequence, the performance outcomes it experiences. In our analysis, the stable fundamental organizational property was the firm's search strategy, both how alterative initiatives are evaluated (on-line vs. off-line) and the structural arrangements of the organization (integrated vs. functional). The search strategy in turn determines the broad characteristics of the dynamics of organizational adaptation.

In examining selection, efficacy can be considered from two very different perspectives. First, from a simple cross-sectional perspective, to what degree are poorly performing

firms selected out of the population at any given point in time? Second, from a longitudinal perspective, are those entities that are on a superior latent adaptive trajectory more likely to survive? The longitudinal problem exists because organizations are adapting, and selection is inherently myopic. Even the self-consciously forward-looking selection efforts of investment analysts and venture capitalists are necessarily constructed on the basis of current and historical data. We held constant cross-sectional efficacy and focused our attention on the longitudinal problem, which only exists when there is adaptation: Under what conditions can myopic selection demonstrate the possibility of prospective intelligence, identifying adapting organizations at intermediate points in time, in the midst of change, that will be superior in the long run?

Results of our model point to the fact that the efficacy of selection is endogenous to the dynamics of organizational adaptation. Important stable attributes of organizations such as strategies or routines may differ not only in their contribution to organizational fitness or performance at a given point in time but may also have differential impacts on the efficacy of selection over time. We denoted this effect on the efficacy of selection as selectability, an organization-level construct that measures the information content embodied in current performance as a determinant of the efficacy of selection in identifying organizations that will be superior in the long run. For organizations that exhibit high levels of selectability, intermediate performance provides superior information about long-run performance, and consequently, selection is more effective longitudinally within populations of such organizations.

An important driver of selectability is the extent to which adaptation proceeds more or less reliably over time. Our finding on reliability is closely related to the ideas developed in the population ecology literature, in particular, the early argument that selection occurs on the basis of reliability (Hannan and Freeman, 1984). Although later work has relaxed this argument (cf. Hannan, 2005: 53), our results suggest that selection on performance alone is sufficient to generate a correlation between reliability and survival: the role of reliability emerges endogenously from a model in which inert strategies guide organizational learning in a competitive ecology. The correlation between reliability and survival occurs because for patterns of adaptation that are highly reliable over time, an organization's performance at any point in time provides better information for selection to act upon in filtering out inferior organizations (i.e., it is more selectable). For example, the integrated search strategy exhibits higher selectability than does the baseline blind strategy. As a consequence, selection across integrated-strategy firms is more effective at removing poor-quality organizations. Conditional on survival, firms using integrated search on average outperform those engaged in blind search, even though unconditional on survival, they exhibit identical long-run performance. That said, though our model highlights the process by which selection favors reliability, it was only tested in a stable environment. Left unexplored in our paper are the implications of

our model for the consequences of reliability in performance under conditions of environmental turbulence (Dobrev, Kim, and Carroll, 2003).

Taking the myopia of selection into account has important implications for our understanding of the interrelationship between the processes of adaptation and selection, in particular, for understanding the relative merits of exploration versus exploitation strategies. March's (1991) well-known result is that exploratory strategies (with higher variance) tend to favor long-run performance, while exploitive strategies favor near-term performance. But strategies that are effective in the long run must survive a series of short-run selection demands (Levinthal and March, 1993). The myopia of selection suggests that prospects for the long-run performance advantage of exploration must be offset by the higher performance variance over time that tends to obscure estimates of an organization's underlying quality. This difficulty of estimating performance in turn increases the risk that latently superior firms will fail and inferior firms will survive. It would even suggest that firms with poor endowments would choose to pursue an exploration strategy, not simply because it may enhance their long-run performance prospects but also because higher variance masks the performance deficit between them and their rivals.

These results provide an alternative explanation for the observation that inferior organizations often come to dominate the set of surviving firms. Scholars have argued that this might occur because performance is not the sole determinant of the outcome of selection (Anderson and Tushman, 1990). Our study suggests an additional explanation: that selection may be systematically prone to error. In contexts in which selection is acting across organizations engaged in adaptation, the reliability with which adaptation occurs is centrally important. Strategies that are inferior may dominate if they exhibit more reliability in their adaptive efforts.

In examining errors in the selection process, we uncovered an important relationship between the intensity of selection and its error rate. Researchers have long argued that higher munificence (a less intense selection regime) gives organizations access to external resources (Finkelstein and Hambrick, 1990), which in turn increase their decision latitude and enhance their performance (Baum and Wally, 2003). We identified an alternative mechanism by which munificence may enhance performance, in this case, population performance. In particular, selection that is less intense (more munificent) is subject to fewer errors in selectively removing inferior firms. If the distinctive properties of organizations, such as tacit and organizationally embedded knowledge, are lost with the dissolution of a firm, then less intense selection, with its lower error rate may enhance social welfare.

An additional counterintuitive result is that using more information in the selection process—using historical data in addition to current data—does not necessarily improve the efficacy of selection. In fact, using more data improved the efficacy of selection only when search efforts suffered from a lack of short-wave reliability because of a blind search strat-

egy. The efficacy of selection across strategies that exhibited relatively high short-wave reliability (e.g., integrated and parallel) was strongly diminished by the use of historical data. For search strategies that result in infrequent large shocks, and a consequent lack of long-wave reliability, that devalue older learning (e.g., sequential strategy), the use of historical data also diminishes the efficacy of selection.

Many important empirical contexts have selection that is discrete in time, rather than continuous. Perhaps the dominant example is that of venture capital. Our results suggest that discrete selection at early points in time is significantly less effective than is continuous selection. This result does not rule out the viability of discrete selection regimes. When carrying out selection is costly, a trade-off between frequency and efficacy is to be expected. For example, in the venture capital context, the frequency of selection is in part a function of the cost of engaging in a round of selection. At the same time, our results demonstrate that in the transition from continuous to discrete selection, error rates grow faster under a weak selection regime than a strong selection regime. Our model thus suggests that one would expect to observe significantly more intense selection in a discrete selection regime than in a continuous selection regime and thus a much higher failure rate. This result is independent of the arguments for high failure rates that are commonly put forward about technological or market uncertainty and opportunity costs for venture capitalists.

Finally, our examination of the efficacy of selection also enhances our understanding of the eternal question, what is a good strategy? Our results suggest an interesting answer. All else held constant, in competitive contexts in which survival is conditional on the performance of other firms, higher-variance strategies (March, 1991) decouple mean performance from survival and in doing so, alter the probability of survival for individual firms. At the population level, however, the survival of firms with inferior mean performance reflects a long-run selection error. This result is particularly important when increasing reliability is costly and thus reliability must be traded off against long-run performance. For example, a modicum of exploration may enhance long-run performance, relative to the local search we modeled, but it does so at the cost of decreased reliability. In such cases, counterintuitively, an individual firm's attempt to enhance its performance may diminish the efficacy of selection as a mechanism of population-level adaptation and lead to a set of surviving firms that exhibit inferior performance. The high variance strategy may be superior at the individual firm level but inferior at the population level.

A logical implication of the claim that selection makes systematic errors in identifying organizations that would go on to do well is that, from the perspective of individual organizations, pursuing a strategy that increases selectability need not be advantageous. For example, organizations with poor initial strategic positions that engender limited long-run prospects may be best served by pursuing an adaptation strategy that purposefully obfuscates their performance outlook, reducing their selectability. Such a strategy may

enhance their own survival prospects, even though it reduces the efficacy with which selection is able to enhance population-level performance. In this sense, systematic errors in the selection process, which are socially deleterious, may well represent strategic opportunities for organizations.

Organizations engage in learning processes to enhance their performance prospects. Implicit in this organizational objective is that such action enhances their survival prospects. Yet learning is a long and drawn out process—end states of learning may take months or years to achieve. Under the influence of the myopia of selection, the dynamics of organizational adaptation are a critical determinant of both the likelihood of survival for any particular organization and, more broadly, the efficacy of selection as a mechanism of population-level adaptation.

REFERENCES

Amburgey, T. L., D. Kelly, and W. P. Barnett
1993 "Resetting the clock: The dynamics of organizational change and failure." Administrative Science Quarterly, 38: 51–73.

Anderson, P., and M. L. Tushman
1990 "Technological discontinuities and dominant designs: A cyclical model of technological change." Administrative Science Quarterly, 35: 604–633.

Arthur, W. B.
1989 "Competing technologies, increasing returns, and lock-in by historical events." Economic Journal, 99: 116–131.

Baldwin, C. Y., and K. B. Clark
2000 Design Rules: The Power of Modularity. Cambridge, MA: MIT Press.

Barnett, W. P.
1997 "The dynamics of competitive intensity." Administrative Science Quarterly, 42: 128–160.

Barnett, W. P., and R. A. Burgelman
1996 "Evolutionary perspectives on strategy." Strategic Management Journal, 17: 5–19.

Barnett, W. P., and G. R. Carroll
1995 "Modeling internal organizational change." Annual Review of Sociology, 21: 217–236.

Barnett, W. P., H. R. Greve, and D. Y. Park
1994 "An evolutionary model of organizational performance." Strategic Management Journal, 15: 11–28.

Barnett, W. P., and M. T. Hansen
1996 "The Red Queen in organizational evolution." Strategic Management Journal, 17: 139–157.

Baron, J. N., M. D. Burton, and M. T. Hannan
1999 "Engineering bureaucracy: The genesis of formal policies, positions, and structures in high-technology firms." Journal of Law, Economics and Organization, 15: 1–41.

Barron, D. N., E. West, and M. T. Hannan
1994 "A time to grow and a time to die: Growth and mortality of credit unions in New York City, 1914–1990." American Journal of Sociology, 100: 381–421.

Baum, J. A. C., and T. Amburgey
2001 "Organizational ecology." In J. A. C. Baum (ed.), Blackwell Companion to Organizations: 304–326. Malden, MA: Blackwell.

Baum, J. A. C., and J. V. Singh
1994 Evolutionary Dynamics of Organizations. New York: Oxford University Press.

Baum, J. R., and S. Wally
2003 "Strategic decision speed and firm performance." Strategic Management Journal, 24: 1107–1129.

Burgelman, R. A.
1996 "A process model of strategic business exit: Implications for an evolutionary perspective on strategy." Strategic Management Journal, 17: 193–214.

Campbell, D. T.
1960 "Blind variation and selective retention in creative thought as in other knowledge processes." Psychological Review, 67: 380–400.
1965 "Variation and selective retention in socio-cultural evolution." In H. R. Barringer, G. I. Blankstein, and R. W. Mack (eds.), Social Change in Developing Areas: A Reinterpretation of Evolutionary Theory: 19–48. Cambridge, MA: Schenkman.
1982 "Evolutionary epistemology." In H. C. Plotkin (ed.), Learning, Development, and Culture: Essays in Evolutionary Epistemology: 73–107. Chichester, UK; New York: Wiley.

Carroll, G. R., and J. R. Harrison
1994 "On the historical efficiency of competition between organizational populations." American Journal of Sociology, 100: 720–749.

Cohen, M. D., R. Burkhart, G. Dosi, M. Egidi, L. Marengo, M. Warglien, and S. Winter
1996 "Routines and other recurring action patterns of organizations: Contemporary research issues." Industrial and Corporate Change, 5: 653–698.

Cyert, R. M., and J. G. March
1963 A Behavioral Theory of the Firm. Englewood Cliffs, NJ: Prentice-Hall.

David, P. A.
1985 "Clio and the economics of QWERTY." American Economic Review, 75: 332–337.

Delacroix, J., and A. Swaminathan
1991 "Cosmetic, speculative, and adaptive organizational change in the wine industry: A longitudinal study." Administrative Science Quarterly, 36: 631–661.

Denrell, J., and J. G. March
2001 "Adaptation as information restriction: The hot stove effect." Organization Science, 12: 523–538.

Dobrev, S. D., T.-Y. Kim, and G. R. Carroll
2003 "Shifting gears, shifting niches: Organizational inertia and change in the evolution of the U.S. automobile industry, 1885–1981." Organization Science, 14: 264–282.

Ethiraj, S. K., and D. Levinthal
2004 "Bounded rationality and the search for organizational architecture: An evolutionary perspective on the design of organizations and their evolvability." Administrative Science Quarterly, 49: 404–437.

Finkelstein, S., and D. C. Hambrick
1990 "Top-management-team tenure and organizational outcomes: The moderating role of managerial discretion." Administrative Science Quarterly, 35: 484–503.

Gavetti, G., and D. A. Levinthal
2000 "Looking forward and looking backward: Cognitive and experiential search." Administrative Science Quarterly, 45: 113–137.

Greve, H. R.
2002 "Sticky aspirations: Organizational time perspective and competitiveness." Organization Science, 13: 1–17.

Hannan, M. T.
2005 "Ecologies of organizations: Diversity and identity." Journal of Economic Perspectives, 19: 51–70.

Hannan, M. T., and J. Freeman
1977 "The population ecology of organizations." American Journal of Sociology, 82: 929–964.
1984 "Structural inertia and organizational change." American Sociological Review, 49: 149–164.

Haveman, H. A.
1992 "Between a rock and a hard place: Organizational change and performance under conditions of fundamental environmental transformation." Administrative Science Quarterly, 37: 48–75.

Helfat, C. E. (ed.)
2003 The SMS Blackwell Handbook of Organizational Capabilities: Emergence, Development, and Change. Malden, MA: Blackwell.

Kauffman, S. A.
1993 The Origins of Order: Self-Organization and Selection in Evolution. New York: Oxford University Press.

Kelly, D., and T. L. Amburgey
1991 "Organizational inertia and momentum: A dynamic model of strategic change." Academy of Management Journal, 34: 591–612.

Lane, D. A.
1993 "Artificial worlds and economics, part I." Journal of Evolutionary Economics, 3: 89–107.

Levinthal, D. A.
1991a "Organizational adaptation and environmental selection: Interrelated processes of change." Organization Science, 2: 140–145.
1991b "Random walks and organizational mortality." Administrative Science Quarterly, 36: 397–420.
1997 "Adaptation on rugged landscapes." Management Science, 43: 934–950.

Levinthal, D. A., and J. G. March
1993 "The myopia of learning." Strategic Management Journal, 14: 95–112.

March, J. G.
1991 "Exploration and exploitation in organizational learning." Organization Science, 2: 71–87.

March, J. G., and J. P. Olsen
1989 Rediscovering Institutions: The Organizational Basis of Politics. New York: Free Press.

March, J. G., and Z. Shapira
1992 "Variable risk preferences and the focus of attention." Psychological Review, 99: 172–183.

Miller, D., and M. J. Chen
1994 "Sources and consequences of competitive inertia: A study of the U.S. airline industry." Administrative Science Quarterly, 39: 1–23.

Miner, A. S., T. L. Amburgey, and T. M. Stearns
1990 "Interorganizational linkages and population dynamics: Buffering and transformational shields." Administrative Science Quarterly, 35: 689–713.

Nelson, R. R., and S. G. Winter
1982 An Evolutionary Theory of Economic Change. Cambridge, MA: Belknap Press of Harvard University Press.

Rivkin, J. W.
2000 "Imitation of complex strategies." Management Science, 46: 824–844.
2001 "Reproducing knowledge: Replication without imitation at moderate complexity." Organization Science, 12: 274–293.

Rivkin, J. W., and N. Siggelkow
2003 "Balancing search and stability: Interdependencies among elements of organizational design." Management Science, 49: 290–311.

Romanelli, E., and M. L. Tushman
1994 "Organizational transformation as punctuated equilibrium: An empirical test." Academy of Management Journal, 37: 1141–1166.

Siggelkow, N., and D. A. Levinthal
2003 "Temporarily divide to conquer: Centralized, decentralized, and reintegrated organizational approaches to exploration and adaptation." Organization Science, 14: 650–669.

Simon, H. A.
1955 "A behavioral model of rational choice." Quarterly Journal of Economics, 69: 99–118.
1962 "The architecture of complexity." Proceedings of the American Philosophical Society, 106: 467–482.
1969 The Sciences of the Artificial. Cambridge, MA: MIT Press.
1991 Models of My Life. New York: Basic Books.

Singh, J. V., R. J. House, and D. J. Tucker
1986 "Organizational change and organizational mortality." Administrative Science Quarterly, 31: 587–611.

Sorenson, O.
2003 "Interdependence and adapt-
ability: Organizational learning
and the long-term effect of
integration." Management
Science, 49: 446–463.

Stuart, T. E., and J. M. Podolny
1996 "Local search and the evolu-
tion of technological capabili-
ties." Strategic Management
Journal, 17: 21–38.

Thompson, J. D.
1967 Organizations in Action: Social
Science Bases of Administra-
tive Theory. New York:
McGraw-Hill.

Wilson, E. O., and W. H. Bossert
1971 A Primer of Population Biolo-
gy. Sunderland, MA: Sinauer
Associates.

Winter, S. G.
1964 "Economic natural selection
and the theory of the firm."
Yale Economic Essays, 4:
225–272.

APPENDIX: Construction of the NK Simulation

Task Environment

Firms operate in a task environment consisting of two nearly decomposable
sub-problem spaces (landscapes). A location in landscape 1 is represented
by the vector N^1 (length n^1) of policy choice attributes, and similarly, a loca-
tion in landscape 2 is represented by the vector N^2 (length n^2). Taken togeth-
er, a location in the combined landscape can be described by the row vector
N (length $n = n^1 + n^2$) of firm policy choice attributes such that:

$$N = \{N^1, N^2\} = (a_1^1, a_2^1, \ldots, a_{n^1}^1, a_1^2, a_2^2, \ldots, a_{n^2}^2). \tag{1}$$

For notational convenience, we have omitted the subscript denoting the
firm, i, from the discussion that follows. We assume that each policy choice
is binary (1 or 0) and therefore there are 2^n possible combinations of organi-
zational choices.

The topography of the landscape is a function of the number of interactions
between policy choices, which influence the degree to which the payoff of a
focal policy choice depends on the value of other policy choices. Interactions
can occur both within a landscape and between landscapes (the upper-right
and lower-left quadrants in figure 2).

Consider a single policy choice j on landscape L, a_j^L. For notational conve-
nience, we let $L = 1$ or 2 where $\neg L = 2$ if $L = 1$ and $\neg L = 1$ if $L = 2$. Let K_j^W
represent the vector of policy choices of length k^W ($a_j^L \notin K_j^W$) within land-
scape L that affects the payoff to policy a_j^L. The k^W policies with which a_j^L
interact are specified as the k^W adjacent elements such that:

$$K_j^W = \begin{cases} (a_{j+1}^L, \ldots, a_{j+k^W}^L) & \text{if } k^W > 0 \\ 0 & \text{if } k^W = 0. \end{cases} \tag{2}$$

In addition, let K_j^B represent the vector of policy choices of length k^B on the
other landscape $\neg L$ that affect the payoff to policy choice a_j^L. Specifically, we
randomly fill the vector K_j^B such that on average there are k^B policies with
which a_j^L interact. Thus, this notation implies symmetry in values above and
below diagonal. Though we implemented and ran a model with symmetry,
and the results are identical, the models employed in the paper are not sym-
metrical.

$$K_j^B = \begin{cases} (a_{-1}^{\neg L}, \ldots, a_{-k^B}^{\neg L}) & \text{if } k^B > 0 \\ 0 & \text{if } k^B = 0. \end{cases} \tag{3}$$

We defined the fitness contribution of an individual policy choice j on land-
scape L, a_j^L, as π_j^L, such that:

$$\pi_j^L = \pi_j^L (a_j^L | a^L \in K_j^W, a^{\neg L} \in K_j^B) = \pi_j^L (a_j^L | K_j^W, K_j^B). \tag{4}$$

Myopia of Selection

The fitness contribution of each policy choice, π_j^L, can take $2^{1+k^W+k^B}$ unique values that are assigned as independent and identically distributed draws from a uniform distribution u[0,1].

When $k = k^W = k^B = 0$, the payoff to a policy choice depends on that choice alone. As k increases, the payoff to a policy choice depends not only on that choice but also on the value of k other policies. In general, low levels of interactions result in a smooth landscape (at the extreme, when $k^W + k^B = 0$, the landscape has only a single peak), and high levels of interactions result in a very rugged landscape with many local peaks, reflecting a highly complex task environment. Asymmetry in the levels of interactions within and between the two nearly decomposable problem spaces, where $k^W \gg k^B$, produces a landscape that can be described as nearly decomposable. In the extreme, when $k^B = 0$, the landscape is fully decomposable.

The fitness of a firm with the policy choice (location) vector N^L on landscape L is defined as the average of the fitness contribution of each of its N^L policy choices such that:

$$\Pi^L = \frac{1}{n^L}\sum_{j=1}^{n^L}\pi_j^L(a_j^L \,|\, a_-^L \in K_j^W, \, a_{-}^{\neg L} \in K_j^B) = \frac{1}{n^L}\sum_{j=1}^{n^L}\pi_j^L(a_j^L \,|\, K_j^W, \, K_j^B). \tag{5}$$

A firm's fitness on the combined landscape is the average of its fitness on each individual landscape such that:

$$\Pi = \frac{1}{2}(\Pi^1 + \Pi^2). \tag{6}$$

Search

Adaptation is modeled as a local search process in which a firm selects an innovation from an alternative set consisting of all policy-choice sets that vary from the current set by a single policy attribute. Thus at any location on an individual sub-problem landscape (L = 1 or 2), there are n^L alternatives from which to choose.

Firms are allocated two search efforts per period and as such, each period is subdivided into two sub-periods $t = t_a$ and $t = t_b$. Search varies across organizational structures and methods of evaluation which can be divided into two main forms based on organizational structure: (a) functional and (b) non-functional. The parallel search strategy, which employs a functional organizational structure with off-line evaluation, proceeds as follows. In the first sub-period of a given period t such that $t = t_a$, the firm selects at random one of its N_t^1 landscape 1 policy choices to evaluate for possible change. Consider the policy choice $a_{j,t}^1$ and its inverse $\hat{a}_{j,t}^1$ (recall that policy choices are binary such that a flip entails switching a 1 to a 0 and vice versa). Without moving to the new location in landscape 1, the firm compares its expected fitness after the flip to its fitness prior to the flip and moves to the new location (accepts the flip) only if the move results in an improvement in fitness. The evaluation of fitness is local, such that the firm considers only landscape 1 performance in assessing the outcome of a search attempt. After concluding search in landscape 1 (L = 1), the firm then conducts search in sub-period $t = t_b$ in landscape 2 (L = 2). More formally, if:

$$\frac{1}{n^L}\sum_{m=1}^{n^L}\pi_{m,t}^L(a_{m,t}^L; \, a_{\cdot(m,j)}^{L}{}_{,t}, \, \hat{a}_j^L \,|\, K_{m,t}^W K_{m,t}^B) >$$

$$\frac{1}{n^L}\sum_{m=1}^{n^L}\pi_{m,t}^L(a_{m,t}^L; \, a_{\neg(m,j),t}^L, \, a_{j,t}^L \,|\, K_{m,t}^W, \, K_{m,t}^B)$$

$$\text{then} \begin{cases} N_t^L = (a_{\neg j,t}^L, \, \hat{a}_{j,t}^L) \text{ otherwise,} \\ N_t^L = (a_{\neg j,t}^L, \, a_{j,t}^L). \end{cases} \tag{7}$$

The second functional structure gives rise to the sequential search strategy, which is identical to the parallel strategy with the exception of the allocation of effort across the two sub-problem landscapes. Prior to period 30, sequen-

tial-strategy firms employ both search efforts in landscape 1. Post period 30, they follow the parallel strategy outlined above, allocating one effort to each sub-problem landscape.

Integrated search strategies come in two forms, baseline blind (on-line evaluation) and integrated (off-line evaluation). For these strategies, a policy choice is chosen from the global landscape $N_t = \{N_t^1, N_t^2\} = (a_1^1, a_2^1, \ldots, a_{n^1}^1 a_1^2, a_2^2, \ldots, a_{n^2}^2)$. Two search efforts are made, one each in sub-periods $t = t_a$ and $t = t_b$. The performance consequences of each search effort is evaluated analogously to equation 7, but taking into account the global (rather than sub-problem) landscape. On-line evaluation is done by moving immediately to the new location regardless of the performance consequences. If the flip causes a performance decrease, then at the start of the next period, the firm returns to its earlier location. Off-line evaluation is conducted (as earlier) with a move to the new location only if the move improves performance.

Selection

We examined the outcome of the developmental process under selection regimes that vary in timing, intensity, and duration. The probability of firm i surviving a selection event in period t is defined as:

$$s_{i,t} = \left(\frac{\Pi_{i,t}}{Max(\Pi_t)} \right)^{Z_t}. \qquad (8)$$

where $Max(\Pi_t)$ is the maximum fitness at time t by all firms in the population. Z_t defines the strength of selection. Because time is discrete, we treat Z_t as a vector of real numbers $[0, \infty]$ that define the strength of selection in each time period from $t = 0$ to $t = t_{max}$ such that $Z_t = (z_t, \ldots, z_{t_{Max}})$. A random number is drawn from a uniform distribution $u[0,1]$. If $s_{i,t}$ is greater than or equal to the realization of the draw, then the firm survives. When $Z_t = 0$, the probability of survival for firm i in period t is $s_{i,t} = 1$ (100%). As Z_t increases, the probability of survival conditional on the firm's fitness declines.

Simulation Model

The simulation models 200 firms as they search under risk of selection for 100 periods ($t_{max} = 100$). The simulation proceeds as follows. At time $t = 0$, each firm is randomly allocated an initial location in both landscape 1, $N_{i,0}^1$, and landscape 2, $N_{i,0}^2$. Each period t is divided into two sub-periods (parts a and b), and firms are allocated two search efforts per period. At the end of each full period, after all firms have completed their search effort, firms are subject to selection. Results are averaged over 100 runs of the simulation.