Hiding the Evidence of Valid Theories: How Coupled Search Processes Obscure Performance Differences among Organizations

Nicolaj Siggelkow
University of Pennsylvania
Jan W. Rivkin
Harvard University

Theorists argue that an organization's high-level choices, such as its organizational design or the attributes of its top management team, should influence its performance, yet empirical researchers have struggled to detect such influence. The impact of high-level choices may appear weak, we theorize, because the choices are embedded in coupled search processes. A coupled search process exists in an organization when managers search for high-level choices that shape the search for low-level, operational choices, which in turn determine performance. Using a simulation model, we show that coupled search processes obscure the performance impact of high-level choices through two mechanisms: (1) a survivor effect, arising because firms that persist with poor high-level choices are those that luckily happened on good low-level choices despite their poor high-level choices, and (2) a wanderer effect, arising when firms use good high-level choices to find good low-level choices and achieve strong performance but then wander toward poor high-level choices. We show that these effects are particularly strong in stable environments, and we identify empirical strategies that can tease out the true performance impact of high-level choices.●

The quest to detect the performance impact of high-level organizational choices, such as organizational designs or the attributes of top management teams, has often frustrated empirical researchers. In the field of organizational design, for example, scholars have made strong theoretical arguments that certain designs should yield higher performance than others (Lawrence and Lorsch, 1967; Galbraith, 1973; Mintzberg, 1979). Yet a critical review of scores of empirical studies published when researchers' interest in this arena peaked reported a vexing array of "mixed, ambiguous, and near-zero associations" between organizational structure and performance (Dalton et al., 1980: 61). Likewise, evidence on the performance impact of attributes of top management teams (TMTs) has been very split, leading some researchers to conclude that "demographics-based TMT studies . . . [are] characterized by weak or uninterpretable findings, unexplained phenomena, and unusable prescriptions" (Priem, Lyon, and Dess, 1999: 938) and that "pursuing this line of inquiry further will yield results inconsistent at best and fruitless at worst" (West and Schwenk, 1996: 571). Common to these settings is that these high-level choices are embedded in what we term coupled search processes.

A coupled search process arises in an organization when managers search for a set of high-level choices that shape the search for a set of low-level, operational choices, which in turn determine performance. The high- and low-level search processes typically occur at different frequencies. For instance, at a relatively low frequency, managers search for an effective set of organizational design choices: an allocation of decision rights, an incentive system, a system of information collection and flow, and so forth. Such high-level choices tend to persist for years. These design choices have a profound effect because they shape a second, higher-frequency search process at a lower level: the process of making day-to-day operational choices concerning pricing, sales calls,

0001-8392/09/5404-0602/$3.00.

●

Special thanks to Howard Brenner for programming help. We thank Philip Anderson, three anonymous reviewers, participants at the Organization Science Winter Conference 2007, and seminar participants at the European School of Management and Technology, the Freie Universität Berlin, Harvard Business School, the Ludwig-Maximilians-Universität München, Ohio State University, the MIT Sloan School of Management, the Universität St. Gallen, and the Wharton School for helpful comments. We are grateful to the Mack Center for Technological Innovation at the University of Pennsylvania and the Division of Research of Harvard Business School for generous funding.

production levels, shipping, procurement, etc. Such operational choices have large and immediate impacts on a firm's economic performance. The resulting performance can then trigger a fresh search for a new design. In that way, the two search processes are coupled.

Coupled search processes are common in organizations. At a low frequency, a firm's board of directors searches for a top management team with winning attributes. At a higher frequency, the team seeks to find and implement an effective array of strategic choices (Hambrick and Mason, 1984). Intermittently, a firm searches for partners and develops its position in a network of allied firms; more frequently, the position influences the firm's ability to tap knowledge and search for innovations (Ahuja, 2000). At a slow pace, managers develop cognitive frames—images of their competitive environment and their place in it; more rapidly, these frames shape their search for a strategy (Gavetti and Levinthal, 2000). A firm's resource allocation process, evolving slowly, molds the firm's deployment of its resources, which has a near-term impact on its behavior and success in the marketplace (Bower, 1970). A firm's long-run investment in dynamic capabilities may determine its ability to reconfigure resources, which then have a short-run influence on performance (Teece, Pisano, and Shuen, 1997).

An enduring goal of organizational research is to isolate the performance impact of high-level choices that are embedded in such coupled search processes. This goal has often proven elusive. We use an agent-based simulation to show how coupled search processes may obscure performance relationships and, for instance, leave the empirical literatures on organizational design and the attributes of top management teams with weak, ambiguous results. Two mechanisms create particular problems: (1) a survivor effect, arising because the firms that persist with poor high-level choices are those that luckily happened on good low-level choices despite their poor high-level choices, and (2) a wanderer effect, arising when firms use good high-level choices to find good low-level choices and achieve strong performance but then wander toward poor high-level choices. These effects dampen the discernable impact of high-level choices on performance and prevent populations of organizations from drifting over time toward good high-level choices. We identify firm and environmental conditions that strengthen or weaken the survivor and wanderer effects. These conditions point us to a set of empirical strategies that can help researchers tease out the true performance effects of high-level choices.

THE ELUSIVE RELATIONSHIP BETWEEN HIGH-LEVEL CHOICES AND PERFORMANCE

Prior Explanations

It is natural to expect high-level choices to leave two types of empirical fingerprints. First, firms that adopt good high-level choices should perform better than those that make poor choices, creating a relationship between high-level choices and performance. Second, in a population in which firms can change their high-level choices over time, the balance of the population should drift toward good choices, a process that

March and Olsen (1989) termed "historical efficiency." When these fingerprints have failed to appear in literatures on high-level choices such as organizational designs and top management teams' attributes, researchers have offered a number of explanations.

Prior literature suggests two potential reasons why it has been hard to find relationships between high-level choices and performance. First, the proposed relationships might be invalid despite the theoretical foundations. For example, if there is widespread "equifinality" (Doty, Glick, and Huber, 1993; Gresov and Drazin, 1997), such that many different organizational designs or attributes of top management teams produce the same levels of performance, then it becomes hard to specify what design or team composition maximizes performance. Similarly, a multiplicity of environmental contingencies may create conflicting requirements and undermine the notion of a single optimal design or ideal team composition (Galunic and Eisenhardt, 1994). Alternatively, the proposed relationships might be invalid because they miss some additional factor that must be combined with good high-level choices to produce well-implemented low-level choices. For instance, if top management searches for and chooses an excellent strategy but does not communicate it to subordinates well enough, the strategy might not be implemented successfully with appropriate low-level choices despite good high-level choices.

Second, the predicted performance relationships, though fully valid, might be inherently hard to detect empirically (Donaldson, 2001). Detecting them requires that we measure constructs such as the autonomy provided by a design or the diversity of a management team, each of which is difficult to assess and quantify. Moreover, constructs can interact with one another, and results may depend sensitively on fairly arbitrary choices of functional forms (Schoonhoven, 1981). If one does not control for interacting elements, the underlying interdependencies among elements may make it hard to find clear relationships; as Galunic and Eisenhardt (1994: 229) put it, "empirical research typically consists of bivariate analysis, whereas reality is multifaceted."

Prior literature also suggests two reasons why populations of firms might not drift over time toward effective high-level choices. First, managers might adopt organizational designs or corporate boards might select top management teams with attributes that they view as legitimate (Meyer and Scott, 1983) rather than ones that are effective. Consequently, herd behavior may lead a population of firms to lock in on a design or to hire a particular type of management team that does not optimize performance (Staw and Epstein, 2000). Second, if the benefits of changing to another organizational structure or to a new leadership team are outweighed by the direct costs of this adaptation, a firm may stick with its less appropriate high-level choices. Similarly, other sources of inertia might exist, including prior investments or internal political pressures (Hannan and Freeman, 1977) or lack of awareness of alternatives (Miller and Chen, 1994), that would stall drift toward more appropriate high-level choices.

Although these explanations for the mixed empirical evidence are valid, none of them acknowledges the presence of coupled search processes within organizations. As we will show, these processes alone, without any of the explanations offered so far, can make it difficult to detect the performance impact of high-level choices such as organizational designs or the attributes of top management teams. Coupled search processes can obscure the performance impact of high-level choices via two mechanisms: a survivor effect and a wanderer effect.

The Survivor and Wanderer Effects

Survivor effect. The survivor effect is an obscuring mechanism that arises because the firms that persist with poor high-level choices are those that luckily happened on good low-level choices despite their poor high-level choices. The logic of the survivor effect is related to the traditional survivor bias one encounters in much empirical work. Assume there is some high-level characteristic Z that has a negative performance effect. In a simple world without a coupled search process—with a deterministic relationship between Z and performance—firms might quickly realize that Z is bad, and they would stop doing Z. But suppose that after choosing Z or not Z, firms have to engage in some search process for high performance. If this search process has a stochastic element, some firms with Z might be lucky and perform quite well. Consequently, some of these firms would persist in doing Z. Other firms would do poorly with Z and would switch to not-Z or die. As a result, at the end of the search process, firms that are still doing Z are not a random sample of firms with Z but the subset of firms with Z that got lucky in their stochastic draws. Obviously, the same is true of those firms that are not doing Z; the not-Z firms that survive also tend to be the lucky ones. It turns out, however, that the performance bias created by the survival process is greater for those with Z than for those without. This difference mutes the performance gap between the two types of firms.

A simple numerical example helps to clarify this effect. Suppose the performance of a firm is given by Performance = $-Z + \varepsilon$. Z is 1 for firms that engage in Z and 0 for those that do not, so firms that engage in Z take a –1 hit to performance. ε is a normally distributed random variable with a mean of zero and variance of 1. Assume that researchers can observe each firm's performance and whether it engages in Z. To test the effect of Z on performance, the researchers compare the average performance for firms with Z to the average for firms without Z. If all firms are observed, the performance of the firms that engage in Z averages –1, and the performance of the firms that do not averages 0. What happens if firms that have low performance are not observable, for example, because they failed? Figure 1 illustrates the effects. Here the frequency distributions of performance are plotted for firms with Z (gray curve) and without Z (black curve). Assume the bottom 10 percent of all firms are not observable (those to the left of the thick, dashed vertical line). While some firms of both types are excluded, firms with Z are far more likely to be in the bottom 10 percent and be dropped from the sample. The averages of both observed samples are biased upwards by this selection process, but the mean for those with Z rises more than the

Figure 1. Illustration of survivor effect.

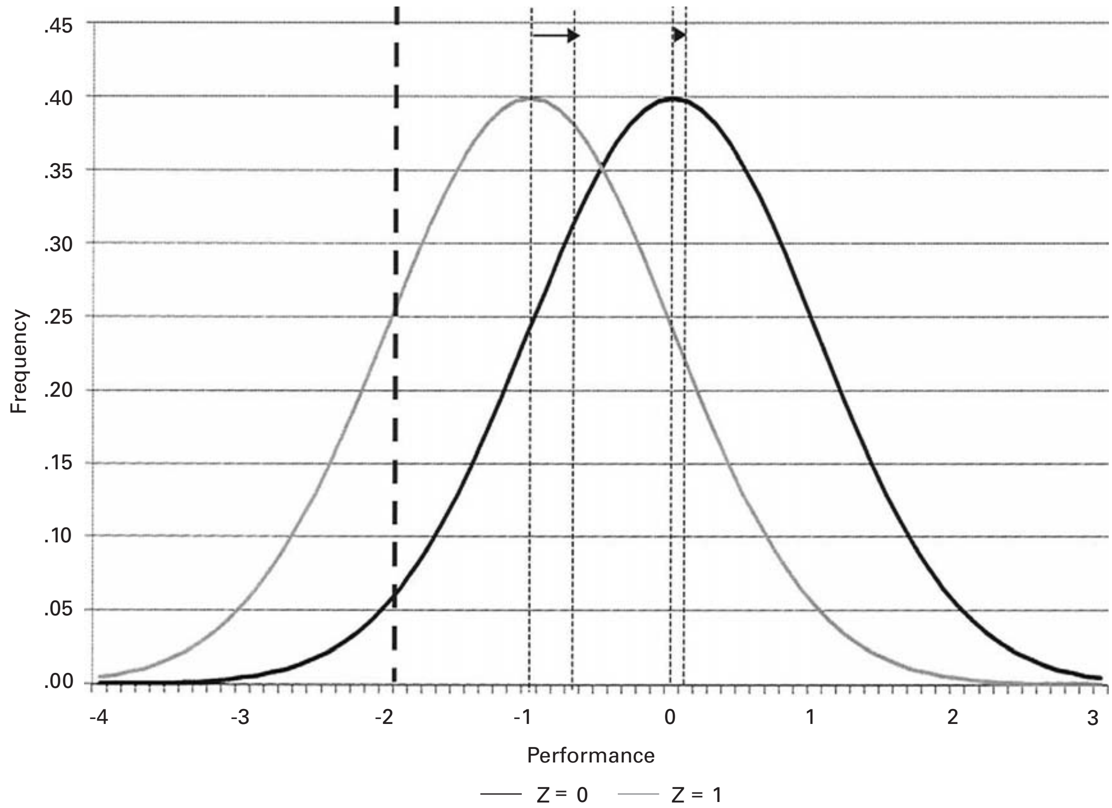

mean for those without Z. (The thin, dotted lines in the figure denote the mean levels before and after dropping the bottom 10 percent.) As a result, the performance difference between firms with Z and those without Z shrinks. The stronger the selection, i.e., the higher the proportion of firms that are not observed, the more the performance difference will be dampened. For instance, in our example, the difference of 1.00 between firms with and without Z becomes only .69 when the bottom 10 percent are not observed and continues to shrink to .56 (20 percent not observed), .37 (50 percent not observed), and .17 (80 percent not observed).

Wanderer effect. The wanderer effect, the second mechanism through which coupled search processes dampen performance differences, arises when firms use good high-level choices to find high-performing low-level choices but then wander toward poor high-level choices while retaining their good low-level choices and strong performance. The survivor effect can arise in a number of settings, but the wanderer effect is particular to coupled search processes. With a coupled search process, high-level choices shape but do not uniquely determine low-level choices, and low-level choices truly determine performance. Consequently, the following may occur. A firm may start with a good high-level choice X, which helps it to discover high-performing low-level choices. At some point, it might experiment with a different high-level choice Y, wandering in the space of high-level choices. Assume that this new high-level choice is inferior to the old one, in the sense that a firm that chose Y throughout its history would, on average, perform worse than a firm that chose X permanently. Even though the firm has changed its high-level choice from X to Y, it may stick with its good low-level choices, because it is trying to maintain performance and the new, less suitable high-level choice Y does not guide it to better low-level choices. Late in this coupled search process, an outside observer might see firms with apparently poor high-level choices performing well because they have stuck with good low-level choices to which earlier, good high-level choices guided them. The observer might misattribute the current high performance to the currently observable high-level choices.

For a concrete example of the wanderer effect, imagine a firm that once had a management team with a great set of attributes. That team uncovered an effective set of strategic choices. Later, the board of directors started to experiment with the composition of the management team. It hired a management team with dysfunctional attributes, but these managers were at least smart enough not to mess up the successful, prior strategic choices. Consequently, they display high performance despite their attributes. If this happens enough in a population of firms, a researcher who looks at the relationship between a management team's attributes and performance could easily conclude that the composition of the management team has little or no effect on performance.

Impact of Adaptation Trigger, Adaptation Mode, and Environmental Turbulence

Our aim in this paper is not only to identify the mechanisms through which coupled search processes obscure the relationship between high-level choices and performance but also to understand the conditions under which these mechanisms, and consequently performance dampening, are likely to be strong. Prior literature has identified three such contingencies: (1) the trigger for adaptation (time or performance level), (2) the mode of adaptation (incremental or mimetic), and (3) the degree of environmental turbulence.

Adaptation trigger. The first contingency concerns what triggers a firm to adapt its high-level choices. Some observations suggest that the passage of time spurs firms to search for new high-level choices, such as a new organizational design (Brown and Eisenhardt, 1997). At regular intervals, a firm tries out something new. A second trigger for high-level search, more in line with what March and Simon (1958) or Cyert and March (1963) might suggest, is poor performance. In this case, a firm is more likely to search for new high-level choices when its performance is poor relative to its aspiration level, which in turn is set by the firm's observation of others' performance levels. The two triggers for adaptation are likely to have quite different influences on the two dampening mechanisms that we described before. When adaptation is time-driven, a lot of wandering takes place, regardless of firms' performance levels. Even great performers are likely to wander away from the high-level choices that guided them to great performance. When high-level change is performance-driven, in contrast, we would expect the survivor effect to be amplified. With performance-driven adaptation, poor performers rapidly alter their high-level choices. Thus, for a firm to stick with a poor set of high-level choices, it must have gotten particularly lucky. Thus we hypothesize:

Hypothesis 1a (H1a): Time-driven adaptation will lead to a larger wanderer effect than will performance-driven adaptation.

Hypothesis 1b (H1b): Performance-driven adaptation will lead to a larger survivor effect than will time-driven adaptation.

Adaptation mode. The second contingency focuses on how a firm that is adapting its high-level choices selects new high-level choices. A firm can do so via incremental changes or by mimicking other firms. Incremental change reflects internal, small-scale experimentation. Incremental and local search, in which "members of organizations begin by looking at close alternatives to improve performance" (Barnett and Sorenson, 2002: 291), has often been used in formal models to reflect managerial behavior (March and Simon, 1958). Empirical support for this assumption in the context of searching for new organizational designs is provided by Colombo and Delmastro (1999: 264), who found, in a sample of 438 Italian metalworking plants, that organizational change with respect to decision autonomy at different levels in the organization was "characterized by a process of marginal adaptation instead of radical modification."

In contrast, with mimicry, firms do not experiment locally; they imitate the high-level choices of other firms. Prior conceptual and empirical work (e.g., Haveman, 1993) indicates that firms are more likely to imitate high-performing firms than low-performing firms. Thus, in the terminology of Haunschild and Miner (1997), firms engage in "outcome imitation" by choosing to copy a high-level choice that appears to have good performance consequences. Intuition suggests that mimicry will lead to less wandering because mimicry allows firms to wander only toward the successful, not freely. Intuition also suggests that mimicry will generate a larger survivor effect. When all firms can rapidly leap toward successful high-level choices, the only way to persist with a poor high-level choice is to be very lucky. All but the very lucky will quickly move far away from the poor high-level choices. Thus we hypothesize:

Hypothesis 2a (H2a): Mimicry will lead to a smaller wanderer effect than will incremental experimentation.

Hypothesis 2b (H2b): Mimicry will lead to a larger survivor effect than will incremental experimentation.

Environmental turbulence. The third contingency we consider addresses whether the mapping from low-level choices to performance (the environment) is stable or changes over time. If the environment is stable, then as long as a firm sticks with its low-level choices, its performance level remains the same, because the high-level choices do not matter directly for performance. In that context, wandering is quite likely. The kind of wandering with impunity shown in the example above, the firm that hired a dysfunctional group of managers who at least stuck with the successful set of low-level choices of their predecessors, works only with a stable environment. If the environment changes, then the new group is tested quickly. The team is revealed to be dysfunctional, and the board might start looking for a new management team that can reproduce the earlier success in a new and changed environment. Thus wandering is less likely to occur in turbulent environments:

Hypothesis 3a (H3a): Turbulent environments will lead to a smaller wanderer effect than will stable environments.

Intuition also suggests that the survivor effect will be smaller in a turbulent environment. The survivor effect is driven by firms that persist with poor high-level choices because they were lucky with their low-level choices. In a turbulent environment, however, a firm would have to be lucky with its low-level choices repeatedly to be a survivor. This is simply unlikely to happen. As a result, we expect survivor effects to be very small in turbulent environments:

Hypothesis 3b (H3b): Turbulent environments will lead to a smaller survivor effect than will stable environments.

Together, hypotheses 3a and 3b suggest that turbulent environments will lead to less dampening of expected performance relationships than will stable environments.

MODEL

Our goal was to study how coupled search processes affect the ability to detect underlying relationships between high-level choices and ensuing performance. The model we present below is fairly general, yet to make the exposition concrete, we situate our description in the context of organizational design. We model firms that make high-level design decisions and low-level activity choices. Guided by its design choices, each firm constantly searches for activity choices that produce good performance. Less frequently, each firm revisits its design choices, searching for designs that may allow it to improve performance even further.

We used simulation modeling, which "is particularly useful when the theoretical focus is longitudinal, nonlinear, or processual, or when empirical data are challenging to obtain" (Davis, Eisenhardt, and Bingham, 2007: 481)—all of which are true in studying the impact of design on performance. Several prior studies have used simulations to assess population-wide changes in the distribution of organizational designs (e.g., Carley and Svoboda, 1996; Carroll and Harrison, 1994; Siggelkow and Levinthal, 2003, 2005; Strang and Macy, 2001), but we know of no prior efforts to simulate a coupled search process in which firms endogenously choose when and how to change their designs.

The agent-based simulation approach we used allowed us to model the high- and low-level search processes directly and let the consequences of these processes emerge from the simulation. At the same time, the model allowed us to eliminate a host of other issues that would be difficult to control for in empirical studies, including the alternative explanations for non-results noted above. In the model, high-level choices truly influence performance; there are no missing factors that can produce implementation errors in low-level choices; relevant constructs can be measured perfectly; managers pursue performance and not legitimacy; changing high-level choices is costless; and all firms are endowed with identical starting conditions and capabilities for change. Moreover, we made the problem of searching for high-level choices as simple as possible to avoid interactions among various elements of design. In particular, we focused on a single, one-dimensional element of organizational design: the degree to which, in a multilayered firm, department managers are allowed to narrow the information that flows to superiors, i.e., their degree of autonomy to narrow down options before they have to turn to superiors. If coupled search processes make it difficult to detect the optimal degree of this one-dimensional variable, even more problems for analysis are likely to arise when designs are more complex. Lastly, the model gave us control over the environments in which firms operate, which allowed us to pinpoint the effect of environmental contingencies on our results. Our model consists of three parts: environments in which firms operate, search for high-performing configurations of low-level activities, and search for better high-level choices.

Environments

We conceptualize a firm's management team as facing a system of interdependent decisions about its activities (Porter, 1996; Siggelkow, 2002). Each firm must decide, for instance, how much to train its sales force, whether to field a broad product line or a narrow one, or whether to pursue basic research or not. Moreover, a number of these low-level choices interact with each other in influencing firm performance. For instance, the value of having a well-trained sales force might increase as a firm broadens its product line. A firm's environment is, in its simplest conception, a mapping from the firm's possible sets of interdependent, low-level activity choices to performance levels.

In our model, each simulated firm has to make decisions concerning N activities, and we designate the realized choices by $a_1, \ldots, a_N$. For simplicity, we assume that a firm can resolve each decision in either of two ways. For instance, a_1 may represent the decision to invest in more training of the sales force ($a_1 = 1$) or not ($a_1 = 0$), while a_2 may represent the decision to increase product breadth ($a_2 = 1$) or not ($a_2 = 0$). A firm thus has a total of 2^N possible configurations of low-level activities.

In computational studies of firms as interdependent systems, it has become common to visualize the payoffs to these choice configurations as a performance landscape (e.g., Levinthal, 1997; Gavetti and Levinthal, 2000; Rivkin, 2000; Lenox, Rockart, and Lewin, 2006; Levinthal and Posen, 2007). A performance landscape consists of N "horizontal" dimensions, each representing one of the N decisions a firm has to make, and one "vertical" dimension, which records the payoff to each of the possible choice configurations. A performance landscape is thus a mapping of any possible vector of firm activities $\mathbf{a} = (a_1, a_2, \ldots, a_N)$ to a performance value $V(\mathbf{a})$. This mapping is created in our simulation with a variant of the NK model (Kauffman, 1993), which has been employed in a number of organizational studies (for a survey, see Sorenson, 2002). It is assumed that the contribution of each individual activity choice a_i to firm payoff V is affected by the state (0 or 1) of the choice itself and by the states of K other choices $\mathbf{a}_{-i}$. For instance, the value that a firm derives from training its sales force is likely to be influenced by whether the firm decides to increase its product breadth.

Denote the contribution of choice a_i by $c_i(a_i, \mathbf{a}_{-i})$. For each landscape, the particular values of all possible c_i's are determined by drawing randomly from a uniform distribution over the unit interval, i.e., $c_i(a_i, \mathbf{a}_{-i}) \sim u[0, 1]$. The performance of a specific set of choices $\mathbf{a}$ is then given by the average of the N contributions: $V(\mathbf{a}) = [c_1(a_1, \mathbf{a}_{-1}) + c_2(a_2, \mathbf{a}_{-2}) + \ldots + c_N(a_N, \mathbf{a}_{-N})]/N$. In all simulations reported in the body of the paper, we set N to 8 and K to 7. Thus we consider performance landscapes that have a high degree of interdependence. Drawing on Simon (1962) and Thompson (1967), who considered interdependence as a key driver of complexity, we could also say that these environments are highly complex.

Environmental turbulence, a contingency variable, is implemented as follows. In stable environments, firms operate on the same landscape for their entire life-histories (1,000 periods). In turbulent environments, correlated shocks reshape the landscape periodically. In particular, once a landscape is created, every 50 periods each contribution value c_i is replaced by $0.2*c_i + 0.8*u$, where u is a new draw from a uniform distribution over the unit interval.

Search for High-performing Configurations of Low-level Activity Choices

In the first period of each simulation, firms are placed on a randomly chosen point of the landscape, that is, endowed with a randomly chosen set of activity choices. In subsequent periods, each firm tries to find higher-performing sets of choices, higher locales on the performance landscape. In modeling how firms search a performance landscape, we focus on a two-layered organizational design, in which department managers search for improvements in their respective departments and send proposals to senior managers, who in turn coordinate departmental choices. As described in the next section, firms may, over time, change the degree to which department managers are allowed to narrow the information flow to superiors. We focus on this particular organizational design decision for three reasons. First, the degree of departmental autonomy to narrow down options before superiors have to be involved is a central issue of organizational design. Second, this choice can be modeled as a one-dimensional object for search—how many proposals have to be sent up—thereby providing a very simple setting for organizational search. Third, as shown below, a clear relationship exists in our model between the richness of information flow and performance. Due to the many interdependencies among the low-level activity choices, a firm benefits from a rich information flow and, hence, relatively little autonomy. We do not claim that this particular organizational structure is in any way optimal. Other structures, including ones that stress more horizontal information exchange, may outperform this one (Siggelkow and Rivkin, 2005), but to simplify the organizational search process as much as possible, we focus on adaptation of design within this structure.

Search for better sets of low-level activity choices proceeds as follows. We assume that the firm's eight decisions are split between two department managers, A and B. The first manager has responsibility for the first four decisions, and the second manager has responsibility for the remaining four decisions. In each period, each department manager evaluates all four local alternatives to the status quo within his or her department and ranks them from most attractive to least attractive for the department. An alternative is local if it differs in just one choice. For instance, if the status-quo choices for a department manager are 0000, then a department manager would evaluate 1000, 0100, 0010, and 0001. Department managers rely wholly on local search to locate better alternatives; they cannot, for instance, observe and copy the low-level choices of other, successful firms. If manager A has control over the first four decisions, and manager B has control over decisions 5–8, then manager A evaluates each alternative by computing $V_A = [c_1(a_1, \mathbf{a}_{-1}) + c_2(a_2, \mathbf{a}_{-2}) + c_3(a_3, \mathbf{a}_{-3}) + c_4(a_4, \mathbf{a}_{-4})]/4$, while manager B computes $V_B = [c_5(a_5, \mathbf{a}_{-5}) + c_6(a_6, \mathbf{a}_{-6}) + c_7(a_7, \mathbf{a}_{-7}) + c_8(a_8, \mathbf{a}_{-8})]/4$.

After evaluating and ranking the alternatives, each manager sends the P most preferred proposals to senior management. P captures the richness of a firm's vertical information flow. A low value of P reflects a firm in which managers are expected to or permitted to narrow down options a great deal before turning to superiors. A high value of P reflects a firm in which senior managers want to review many alternatives themselves. P must be between 1 and 5, because department managers need to send up at least one proposal and cannot send up more than the four local alternatives and the status quo.

Senior management focuses on coordinating the actions of the firm's two departments. Using the departments' proposals and status-quo choices, senior management pieces together at random one composite alternative for all eight activity choices; assesses it in light of the interests of the firm as a whole, evaluating the overall $V(\mathbf{a})$; compares it with the status quo; and implements the option if it yields a higher payoff for the firm than the status quo does. The activity configuration implemented by senior management forms the starting point for further search at the departmental level in the next period. Heterogeneity arises between firms with the same value of P because each senior management team constructs a composite alternative by choosing at random among the proposals from its two departments. A firm whose senior management happens to pick a good combination of proposals to evaluate may

Figure 2. Flow chart of simulation model capturing search for high-level, organizational design choices and low-level activity configurations.

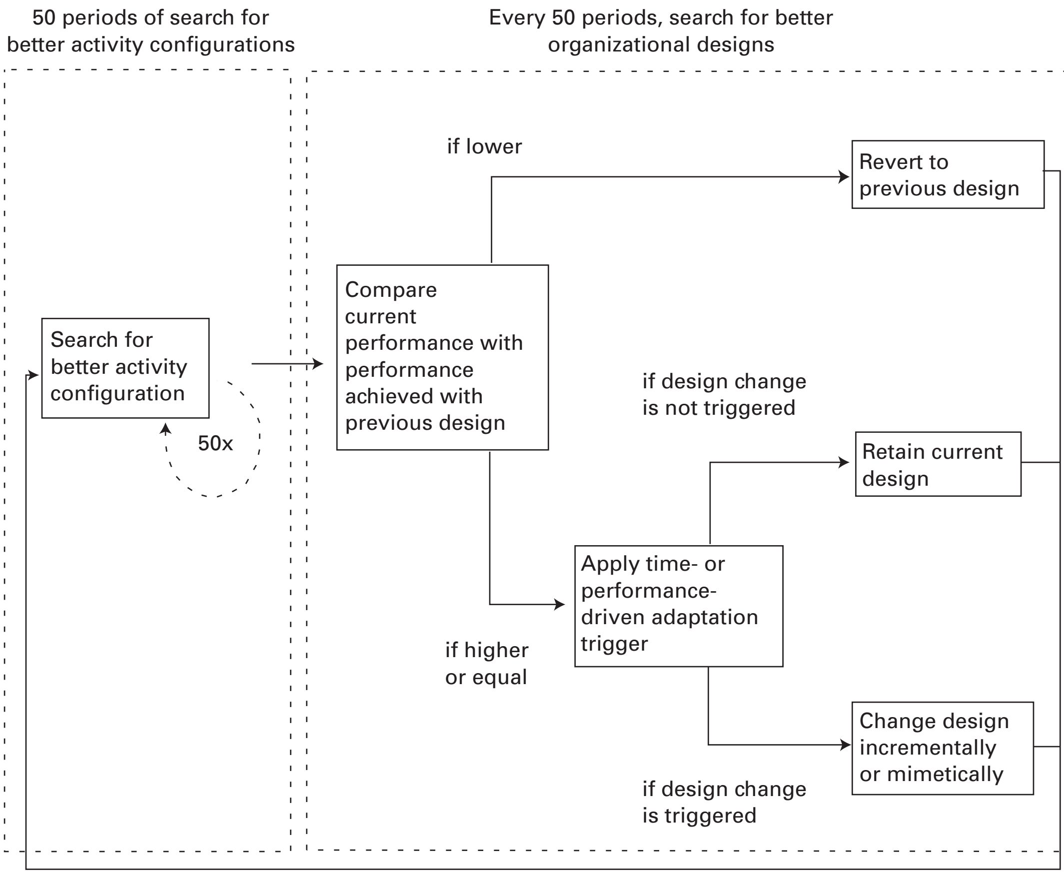

improve its performance faster than a firm with less lucky senior management.

Search for Better Organizational Designs

Figure 2, a flow chart of the simulation, shows how the high-level search for a design is linked to the low-level search for the activity configurations responsible for performance. In contrast to the search for better activity configurations, which occurs every period, the search for better high-level, organizational design choices occurs less frequently. Each firm considers a design change every 50 periods. It first has the opportunity to compare the final performance achieved under the current design with the final performance achieved under its most recent different design. If the performance under the current design is worse than the performance under its most recent design, the firm reverts to the most recent design, which completes the adaptation process for this 50-period cycle. Thus a firm is allowed to recover from a failed experiment, at least with respect to its design; it still retains its current set of choices. If a firm does not revert, it proceeds to search for a better design, as described below.

We distinguish between two different triggers for adaptation. If the firm has a time-driven trigger for adaptation, it always changes its design every 50 periods. If the firm has a performance-driven trigger for adaptation, it searches for a new design with a probability that is inversely correlated with its relative performance. In particular, let Π be the performance of the focal firm, let Π_{max} be the maximum performance of all firms in the landscape, and let Π_{min} be the minimum performance. Then the focal firm reorganizes with probability $p = (\Pi_{max} - \Pi)/(\Pi_{max} - \Pi_{min})$. This implies, for instance, that the firm with the best performance never reorganizes, and the firm with the worst always reorganizes. Should all firms have the same performance, i.e., $\Pi_{max} = \Pi_{min}$, we assume that reorganization occurs for sure. With performance-driven change, a firm's two search processes are coupled in both directions: a firm's organizational design shapes its search for better activity configurations, and the success of its search for better activity configurations affects how frequently it searches for better organizational designs.

Once a firm decides to reorganize, it has to determine how to change its design. In the simulation, the only element of design that a firm can change is P, the richness of information exchange. If a firm uses incremental experimentation to search for a new level of P, it either increases or decreases P by one unit with equal probability. Should P reach its upper or lower bound, the firm will move away from the boundary with a probability of 1/2 and remain at the boundary with a probability of 1/2. If a firm uses mimicry, it will imitate the level of P of another firm, most likely a high performer. In particular, if Π_i is the performance of any given firm i, then a focal firm m will copy another firm j, j ≠ m, with probability $p = \Pi_j/\Sigma_{i \neq m}\Pi_i$. For example, suppose that there are five firms with performances as follows: Firm 1: 0.3; Firm 2: 0.7; Firm 3: 0.4; Firm 4: 0.6; and Firm 5: 0.8. If Firm 5 reorganizes using the mimicry mode of adaptation, it has a 0.3/(0.3 + 0.7 + 0.4 + 0.6) = 15 percent chance of copying Firm 1's design, a 0.7/(0.3 + 0.7 + 0.4 + 0.6) = 35 percent chance of copying Firm 2's design, and so forth. Each firm chooses the design it intends to mimic before any firm actually implements its changes so that firms do not accidentally mimic a just-adopted design of another firm. Firms mimic only P, the high-level design choice, not the low-level activity choices of effective firms.

Common Parameters

To focus on the effect of organizational design, we eliminate all other heterogeneity across firms in each simulation: all firms start at the same time, with the same activity configuration, performance level, and managerial search and evaluation capabilities. At the beginning of each simulation, firms differ only in their initial degree of information flow, with P being set to 1, 2, 3, 4, or 5. We use five firms of each type, creating a population of 25 firms on each landscape. All firms on the same landscape use the same search process for new organizational designs. For each landscape, we run the simulation for 1,000 periods, and we repeat each simulation on 1,000 landscapes, which are created by randomly re-drawing all contribution values c_i. The reported results are thus always averages over 1,000 observations. Performance values are reported as fractions of the highest performance possible on each landscape.

Table 1

Results of Simulations for Firms in Stable Environments

Final design	(1) Baseline	(2) Time-driven incremental	(3)	(4) Time-driven incremental*	(5)	(6) Perf.-driven incremental	(7)	(8) Time-driven mimicry	(9)	(10) Perf.-driven mimicry	(11)
P = 1	.840	.909	19.3%	.904	8.2%	.896	15.3%	.915	17.0%	.917	7.6%
P = 2	.866	.912	19.5%	.902	14.6%	.905	17.0%	.907	21.1%	.915	10.0%
P = 3	.881	.914	20.4%	.909	20.9%	.912	19.6%	.910	19.3%	.919	14.9%
P = 4	.898	.918	20.2%	.919	25.7%	.922	22.8%	.917	19.7%	.918	23.5%
P = 5	.916	.921	20.7%	.924	30.6%	.928	25.3%	.920	22.8%	.921	44.0%
Dampening		84%		73%		57%		94%		94%	
Degree of dampening due to P = 1:											
Survivors		8%				27%		7%		44%	
Wanderers		83%				46%		91%		56%	
Number of observations required to find significant differences	21	612		215		77		3,752		4,119	

* Results for designs that led the firms to their sticking points, rather than the firms' final designs.

RESULTS

Effects of Adaptation Triggers and Modes

Table 1 reports simulation results for firms that operate in stable environments. Column 1 reports the baseline: the average performance in period 1,000 of firms that cannot change their designs. These are the performance levels that would arise for each level of P without a coupled search process—that is, if firms were assigned immutable designs and only low-level search for better activities occurred. The results show that with many interdependencies among activities, a rich information flow is very valuable: performance increases monotonically with P. For instance, firms with P = 5 have an average performance of .916, whereas firms with P = 1 have an average performance of only .840, statistically a highly different value. The intuition behind this baseline result is that highly complex landscapes are very "rugged," with many sticking points (Rivkin and Siggelkow, 2003). At a sticking point, there is no alternative configuration of the N choices that the actors within the firm would consider and that meets the approval of enough actors to be adopted. Once a firm reaches a sticking point, its search has come to an end, assuming the firm does not change its design and the environment does not change afterwards. Rugged landscapes create the possibility that firms will get stuck with suboptimal choices. Rich information flow increases the degree of exploration in which a firm engages and thereby prevents the firm from getting stuck too quickly, boosting its long-run performance.

Time-driven incremental search process. Given the relatively large performance difference between firms with P = 1 and P = 5 in column 1 and the absence of confounding circumstances, one would expect that when firms can adjust

Figure 3. Performance by organizational design in stable environments.

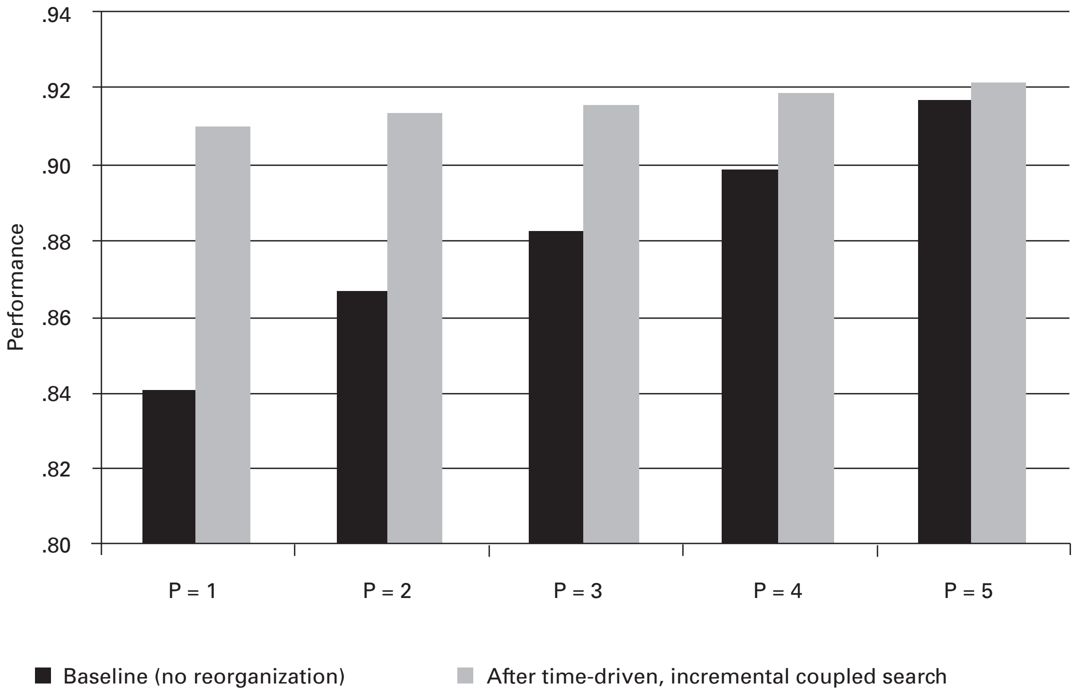

their designs, those with higher levels of P would continue to outperform those with lower levels of P. Moreover, one would expect firms to be able to find the more appropriate designs, so that better designs would appear more frequently in the population of firms than would less appropriate designs. We test these expectations, first, in a coupled search process with time-driven incremental change. Every 50 periods, each firm reorganizes by adjusting its P up or down by 1 or reverts to its prior design. Column 3 in table 1 shows the percentage of firms that end the simulation with each level of P, while column 2 shows the average final performance of these firms. For these firms and for the baseline firms in column 1, figure 3 graphs average performance as a function of P.

Two results in columns 2 and 3 stand out. First, there is hardly any drift toward the more appropriate design (column 3). The simulation started with an equal distribution of firms across the different levels of P (20 percent each), and by the end, we find 19.3 percent of firms with P = 1 and 20.7 percent of firms with P = 5. Second, the performance penalty associated with a poor design appears to be drastically smaller than the baseline would lead one to expect (column 2 vs. column 1; and figure 3). Whereas the performance difference between the best (P = 5) and worst (P = 1) designs was .076 when firms could not change their designs, now the performance difference is only .012, roughly one-sixth as large. The performance differences among designs, while still present, are severely dampened. To gauge the size of this effect, we construct a metric called *dampening*. Let $\Pi_b(P1)$ and $\Pi_b(P5)$ be the baseline performance of firms with P = 1 and P = 5 when they cannot change their design, and let Π(P1) and Π(P5) be the performance of firms that have P = 1 and P = 5 at the end of each simulation. Then *dampening* = 1 – [Π(P5) – Π(P1)]/ [Π_b(P5) – Π_b(P1)]. A dampening of 100 percent would imply that the entire performance effect of design is concealed by coupled search processes.

As reported in the bottom panel of table 1, about 84 percent of the performance difference between firms with best and worst designs is dampened by the coupled search processes. This performance dampening is driven almost exclusively by the higher performance of firms with poor designs rather than the lower performance of firms with good designs. To understand why firms with seemingly poor designs have high performance, we examined the survivor and wanderer effects. To gauge how much dampening is created by P = 1 survivors, we identify survivors by tagging firms whose performance at the end of period 50 equals the performance in period 1,000 and whose starting design equals its final design. Each of these firms found a sticking point with its initial design, never found a better configuration of activity choices, and ended up with its original design. Given that firms using the time-driven, incremental search process change their design every 50 periods, we would expect relatively few survivors to exist, and, in fact, only about 11 percent of all P = 1 firms observed in period 1,000 are survivors. As expected, the performance of these survivors, .892, is significantly higher than the performance of the average baseline P = 1 firm of .840. Overall, the degree of dampening created by survivors is 8 percent, as shown in table 1.[1]

To gauge the size of the wanderer effect, we identify those firms that ended up with P = 1 yet found their sticking point with a different level of P. For instance, about 26 percent of all firms that have P = 1 as their final design actually reached their final performance level with a P = 5 design at some earlier time and then remained at this performance level as they wandered to a design of P = 1. These firms' average performance is a relatively high .922 and is being attributed to the P = 1 design because that is the design observed at the end of the simulation for these firms. If firms follow a time-driven incremental search process, such wanderers—firms that achieved their final performance with a design that is different from the current design—are rampant. About 89 percent of all firms with P = 1 at the end of the simulation reached their sticking points (and final-period performance) with a design different from P = 1. Overall, the degree of dampening created by wanderers is 83 percent, as shown in table 1.[2]

Given that in this case dampening is caused mainly by the wanderer effect, perhaps a researcher should have tracked down for each firm the design that really led to its final performance and attributed final performance to that design rather than to the final design. Columns 4 and 5 of table 1 present the results of doing this. Tracking down the design that led to final performance reveals a larger drift toward the appropriate design. On average, only 8.2 percent of all firms reach their final sticking point with a P = 1 design, while 30.6 percent of all firms reach their final sticking point with a P = 5

1
Let Π_b(P1) and Π_b(P5) be the mean performance of firms with P = 1 and P = 5 when they cannot change their design, i.e., the baseline performance. Let Π_S(P1) be the mean performance of P = 1 survivors and α the fraction of final P = 1 firms that are survivors. Then, the degree of dampening created by P = 1 survivors is given by α*[(Π_S(P1) – Π_b(P1)]/[Π_b(P5) – Π_b(P1)].

2
Let Π_b(P1) and Π_b(P5) be the mean performance of firms with P = 1 and P = 5 when they cannot change their design. Let Π_w(P1) be the mean performance of P = 1 wanderers and β the fraction of final P = 1 firms that are wanderers. Then, the degree of dampening created by wanderers is given by β*[(Π_w(P1) – Π_b(P1)]/[Π_b(P5) – Π_b(P1)]. A third category of firms exists, besides survivors and wanderers, that end up with P = 1: firms that started with P = 1 and found their sticking points when they had a P = 1 design, yet not within the first 50 periods. In general, these firms have a relatively small effect on dampening; hence we focus on survivors and wanderers.

design. The degree of dampening, however, is still very high. About 73 percent of the performance difference between the best and the poorest design is washed out. Again the dampening is driven by the higher performance of the P = 1 firms. The same two mechanisms explain this result. First, there is still the same set of high-performing survivors that early on found a sticking point with P =1 and afterwards were unable to find any better performance despite having better designs. Second, a wanderer effect remains at work. Firms that reach a sticking point with P = 1 may have had a different, better design at some earlier time, which allowed the firm to achieve relatively high performance. If that earlier search did not end in a sticking point but search with the P = 1 design did, the P = 1 design would receive credit for the performance generated by the entire sequence of a firm's organizational designs, leading to a seemingly high performance associated with the P = 1 design. Dampening would thus arise even if one were able to track down the design that was responsible for reaching the final performance level.

With a time-driven incremental search process, significant compression of performance differences would make it difficult for a researcher to detect the relationship between design (in this case, P) and performance. Suppose one hypothesizes that firms with P = 5 outperform firms with P = 1. How large a sample would a researcher need in order to find a statistically significant ($p < .05$) difference between the performance of these two types of firms? If firms were endowed with fixed designs (the implicit assumption generally made when theorizing about why certain designs outperform others), one would only need 21 firms of each type. But if one observes these firms only after firms have adjusted and experimented with their designs, one would need 612 firms of each type to find statistically significant support for one's hypothesis. Tongue in cheek, one can conclude that it is nearly 30 times harder to find a significant difference empirically than to theorize about it. (The number of firms of each type required to find significant differences is reported in the last row of table 1.)

Performance-driven incremental search process. So far we have assumed that firms change their designs every 50 periods, regardless of their performance. In the next simulation, we assume instead that the probability of a firm's changing its design is an inversely correlated function of its current performance. Columns 6 and 7 in table 1 report the results of this simulation. In contrast to the time-driven process, we can observe a slightly larger drift toward P = 5 designs, with 25.3 percent of all firms ending up with P = 5 and a lower degree of dampening. There is still significant performance compression, however, with about 57 percent of the performance difference between the best and worst designs being dampened away. Thus performance-driven search processes also show performance dampening across organizational designs. In this case, one would need 77 firms of each design to find a significant difference between firms with P = 1 and P = 5.

To gauge the validity of H1a and H1b, we compute again the degree to which survivors and wanderers dampen performance.

As suggested by H1b, performance-driven search, which allows high-performing firms to retain their designs, creates a larger survivor effect (27 percent) than does time-driven search (8 percent). This difference is significant, with $p < .001$.[3] With a performance-driven adaptation trigger, about 18 percent of all original P = 1 firms survive, reaching their final performance in the first 50 periods, never finding a higher performance later, and ending the simulation with a P = 1 design. About half of these firms never change their design at all and have a P = 1 design throughout their lives. Moreover, these survivors have a high average performance of .929. Likewise, consistent with H1a, the wandering effect is statistically significantly larger for the time-driven adaptation trigger (83 percent) than for the performance-driven search (46 percent) ($p < .001$).

Mimicry as adaptation mode. In simulations testing the effects of mimicry, firms imitate the designs of other high-performing firms. The results of the time-driven mimicry search process appear in columns 8 and 9 of table 1. In aggregate, there is hardly any drift toward P = 5 firms. At the end of the simulation, 17.0 percent of all observed firms had a P = 1 design, while 22.8 percent of firms had a P = 5 design. This apparent diversity of designs persists in part because of diversity across landscapes. The average number of designs one can observe in any given landscape at the end of each simulation is 2.3. Thus the mimetic process does lead to a reduction of variety in any given landscape, but in different landscapes, different designs survive, creating an overall diversity of designs. This outcome has been empirically observed by Djelic and Ainamo (1999), who studied the evolution of organizational forms in the fashion industry. They found the emergence of one dominant organizational design in France, Italy, and the U.S., yet this design was different in each country.

With respect to dampening, we continue to see a large performance compression. The performance difference between firms with P = 1 and P = 5 is dampened by 94 percent.[4] The mechanisms behind the dampening effect remain the same as with the incremental search process above. As before, the dampening effect is driven mainly by the performance improvement of P = 1 firms, and the reason for this high performance is still twofold. Some P = 1 firms are high-performing survivors, and others are wanderers, firms whose current designs have nothing to do with the designs that generated their performance.

Columns 10 and 11 of table 1 report the results when firms engage in a performance-driven mimicry search process. This case produces a larger drift away from poor designs toward better designs. Overall, the number of P = 1 firms is more than halved to 7.6 percent, while 44.0 percent of all firms end with a P = 5 design. The degree of dampening, however, remains extremely high. Still 94 percent of the performance difference is washed away, requiring 4,119 observations of each type to detect a statistically significant performance difference.

In comparing the degree to which survivors and wanderers contribute to dampening in time-driven and performance-driven

3
To compute the significance levels, we proceeded as follows. For instance, let s_1 denote the survivor effect for time-driven incremental search and s_2 the survivor effect for performance-driven incremental search. Given the large number of simulation runs on which the means, which enter the measures of the survivor and wanderer effects, are based (see footnotes 1 and 2), these means are normally distributed. Using these distributions, we employ Monte-Carlo simulations to construct the distribution of, for instance, $s_2 - s_1$. The results of the Monte-Carlo simulations allow us to compute the likelihood that one would wrongly accept the hypothesis that $s_2 > s_1$.

4
The dampening effect is not driven by the fact that P = 1 and P = 5 firms survive in different landscapes. Even if we constrain our analysis to those landscapes in which both P = 1 and P = 5 firms survive (163 out of 1,000 landscapes), the dampening effect is 98 percent.

search processes, we find further support for H1a and H1b. Again, the performance-driven trigger leads to a significantly larger survivor effect than does the time-driven trigger (44 percent vs. 7 percent, $p < .001$), while the time-driven trigger leads to a larger wanderer effect (91 percent vs. 56 percent, $p < .001$).

Hypotheses 2a and 2b compared the effects of incremental change vs. mimicry on the wanderer and survivor effects. H2b posited that mimicry leads to a larger survivor effect than does incremental experimentation. The logic was that mimicry would bring about a larger expected improvement so that only really lucky firms would not benefit from mimicry. Accordingly, firms with P = 1 that survive would be firms with very strong performance. The simulations with the performance-driven adaptation trigger support this argument. With mimicry, the degree of dampening generated by survivors is 44 percent, while with incremental change, the degree of dampening generated by survivors is only 27 percent ($p < .001$). Likewise, the performance of the survivors under mimicry is higher (.956) than under incremental change (.929). When adaptation is purely triggered by time, we do not find support for H2b. Given our understanding of the effects of different triggers on dampening (results from H1b), this result is not very surprising. Once adaptation is triggered by time, survivors do not play an important role in any case. With this trigger, the degree of dampening generated by survivors is fairly constant at a low 7–8 percent regardless of whether change is incremental or by mimicry.

Hypothesis 2a posited that mimicry would lead to a smaller wanderer effect than would incremental adaptation. The intuition was that only firms with good designs would be high performers and therefore the targets of mimicry. As a consequence, firms with good performance would not wander into P = 1. Our results clearly disconfirm this intuition. Regardless of whether the trigger is time-driven or performance-driven, mimicry actually leads to similar or even larger wanderer effects than does incremental change (91 percent vs. 83 percent, $p < .05$; and 56 percent vs. 46 percent, $p < .001$). This raises the question of why poor designs are imitated by firms with better designs. Conceptually, we can identify three reasons why firms might imitate the poor designs of other firms: firms may imitate the lucky, they may imitate the fast, or they may imitate wanderers.

Imagine a firm with P = 1 that for some lucky reason achieved a very high performance. If firms with different, more appropriate designs reached relatively high performance but not as high as the outlier P = 1 firm, these firms might copy the P = 1 design while retaining their high performance. Consequently, one might observe a number of high-performing firms with a P = 1 design and ascribe the high performance of all firms to the P = 1 design, even though the P = 1 design was only responsible for the good performance of one, mimicked firm. In a sense, the process of mimicry magnifies the luck of an outlier P = 1 firm, thereby leading to large wanderer and dampening effects.

Besides luck, P = 1 firms can appear to be appropriate targets for imitation because they improve their performance quickly.

Firms with P = 1 tend to improve faster than firms with P = 5 because senior management in a P = 5 firm, being inundated with proposals, can create a bottleneck in that firm. As a result, in period 50, the long-run superiority of P = 5 designs has not always come to the fore, and some P = 1 firms still outperform P = 5 firms. Hence, a number of P = 1 firms will serve as templates to be copied.

Once firms copy firms with P = 1, the number of firms with P different from 1 decreases, thereby reducing the likelihood that firms in later rounds will find attractive targets that have a P different from 1. Thus, once the imitative process starts, the population of firms can tip toward the "wrong" design. Interestingly, while such seemingly suboptimal outcomes due to path dependencies and positive feedback effects have generally been bemoaned (David, 1985; Arthur, 1989), the population-wide performance outcome when the wrong design wins is not that bad in our context. Wrong designs can win only if they performed fairly well, and other firms may attain relatively high performance before they adopt the wrong design; hence the harm done by eventually adopting the wrong design is muted.

Although we can distinguish conceptually three reasons that firms might imitate the poor designs of other firms (imitate the lucky, the fast, or a wanderer), parsing out the effects empirically is more difficult. For instance, one would have to be able to distinguish whether a high-performing firm with P = 1 was lucky or fast. The following observations provide, nevertheless, some indication of the sizes of the various effects. In about 14 percent of all landscapes, P = 1 firms outperform P = 5 firms even in period 1,000. One might call these firms lucky, as they have been able to find a very good configuration of activity choices despite their low level of P. In an additional 5 percent of all landscapes, P = 1 firms outperform P = 5 firms in period 50 but do not continue to do so in period 1,000. One might call these firms fast improvers, as they were able to achieve high performance quickly, yet they were eventually overtaken by P = 5 firms. Lastly, by period 1,000, for the time-driven mimicry search process, 75 percent of all firms with P = 1 adopted this design by having copied a wanderer. Likewise, for the performance-driven mimicry process, 54 percent of all firms with P = 1 adopted this design by having copied a wanderer. Thus mimicking wanderers appears to play a large role in creating the possibility for P = 1 firms to become numerous in a given landscape.

The Effects of Environmental Turbulence

To test hypotheses 3a and 3b, we observed how firms with different coupled search processes perform in turbulent environments. In column 1 of table 2, we report the baseline performances of firms that do not change their designs in turbulent environments. In this environment, we find again that a high level of P is beneficial. As the other columns in table 2 reveal, regardless of whether adaptation is driven by time or by performance, or whether change occurs incrementally or by mimicry, we observe very little dampening in turbulent environments. In some situations, we even observe a slight increase in performance differences. Turbulence

Table 2

Results of Simulations for Firms in Turbulent Environments

Final design	(1) Baseline	(2) Time-driven incremental	(3)	(4) Perf.-driven incremental	(5)	(6) Time-driven mimicry	(7)	(8) Perf.-driven mimicry	(9)
P = 1	.862	.860	12.9%	.859	11.6%	.865	1.4%	.854	0.7%
P = 2	.884	.883	16.1%	.882	16.2%	.887	6.5%	.883	4.0%
P = 3	.900	.900	20.2%	.898	21.1%	.900	15.1%	.898	13.5%
P = 4	.912	.910	24.0%	.909	24.4%	.910	29.4%	.912	31.8%
P = 5	.919	.920	26.8%	.919	26.8%	.919	47.5%	.919	50.1%
Dampening		–7%		–7%		3%		–17%	
Degree of dampening due to P = 1:									
Survivors		0%		0%		0%		0%	
Wanderers		–1%		–1%		–1%		0%	
Number of observations required to find significant differences	31	28		27		27		22	

negates the two mechanisms for dampening that we identified, supporting hypotheses 3a and 3b. As the table shows, both the survivor and the wanderer effects are essentially eliminated. Consider, for instance, a firm with P = 1 that by some luck achieved high performance by period 50. When the environment changes in period 51 and the relationship between choices and performance is reset, this luck disappears. The firm finds itself in a new situation and its prior good performance is of little help. Moreover, the firm's inappropriate design now hurts it. Hence, the performance of firms that end with a P = 1 design is very similar to those firms with P = 1 that never change their design, yielding only a small dampening effect. With environmental shocks that negate prior advantages, survivor effects vanish.

For similar reasons, the wanderer effect disappears. A high-performing firm that changes its design from a more appropriate design (that led to its high performance) to a less appropriate design will pay a price in a turbulent environment for this misguided organizational change. Again, the firm cannot leverage its high performance from its prior design because the current less appropriate design is now actively employed in reacting to the new environment in which the firm finds itself. As a result, poor designs are correlated with poor performance.

In turbulent settings, we also observe a larger degree of drift toward the P = 5 design in the populations of firms. For each of the four cases, we find a larger portion of P = 5 firms and a smaller portion of P = 1 firms in the turbulent environment than in the corresponding case in the stable environment. For instance, in the stable environment with firms that engaged in time-driven mimicry, we found that in period 1,000, 17.0 percent of all firms had P = 1 and 22.8 percent had P = 5. In contrast, in the turbulent environment, only 1.4 percent of firms have P = 1 and 47.5 percent have P = 5 in period 1,000.

We thus find strong support for hypotheses 3a and 3b. In turbulent environments, the survivor and wanderer effects are diminished. As a consequence, dampening is much smaller, and drift is much more pronounced, in turbulent environments than in stable environments.

Robustness

We tested the robustness of our results with respect to a number of assumptions of the model: the reorganization frequency, the degree of correlation between landscapes in the turbulent setting, the capability of upper management, different optimal organizational structures, and the inclusion of a selection model in which poorly performing firms die and new firms are born. Under all assumptions and conditions, results were consistent with the ones reported here: dampening is pervasive and significant in stable environments and much reduced in more turbulent environments. The Appendix provides more detail.

DISCUSSION AND CONCLUSION

Though cast in the context of organizational design, our analysis has an underlying logic that applies to diverse other contexts in which coupled search processes operate. We identify the conditions that define the boundaries of these contexts and describe the empirical patterns we expect to arise within those boundaries. Our findings pose tough challenges both to researchers who seek empirical evidence of the performance impact of high-level choices embedded in coupled search processes and to managers who work amidst such processes. At the same time, our research reveals that coupled search processes may improve the robustness of populations of organizations.

Boundary Conditions, Diverse Contexts, and Empirical Patterns

The mechanisms that underlie performance dampening apply to a wide range of coupled search processes. Specifically, we expect dampening to occur in populations of organizations in which three conditions are jointly met: (1) Some set of high-level choices affects performance not directly and deterministically but by shaping the search for a set of lower-level choices; (2) organizations explore alternative configurations of the set of high-level choices over time (e.g., in our simulation, reorganize); and (3) some stickiness in the low-level set of choices acts, at least partially, to preserve high performance once it has been discovered. This last element arises in our simulation because senior management refuses to accept changes in activity choices that cause performance to decline. In populations in which these conditions hold, organizations that deploy good high-level choices to discover effective low-level choices can subsequently meander in the space of high-level choices without undermining performance. This creates the wanderer effect. Moreover, organizations that persist with poor high-level choices tend to be those that, by luck, discovered unusually effective low-level choices despite their high-level choices. This drives the survivor effect.

These three are boundary conditions. If any of them does not hold, our results disappear. If the first condition fails, then performance is determined directly by high-level choices and, perforce, the performance effect of those choices is not dampened. If the second condition does not hold, then no movement occurs in the space of high-level choices, and each set of high-level choices is accurately associated with the performance it produces on average; again, dampening does not occur. This corresponds to the baseline results in our simulations. If the third condition fails, then low-level choices adjust when high-level choices change, and organizations that wander away from good high-level choices that produced effective low-level choices no longer do so with impunity. Their low-level choices degenerate, performance declines, and the link between poor high-level choices and poor performance appears without dampening.

These boundary conditions are met for a wide range of organizational processes, including not only a firm's search for an organizational design and the choice of attributes in a top management team (the two examples we have mentioned most often), but also the creation of a firm's position in a network of innovating firms, the formation of managers' cognitive frames, the development of a firm's resource allocation process, and a firm's investments in dynamic capabilities, for example. Each of these processes is characterized by multiple loosely coupled levels of choices, exploration of alternative high-level choices, and stickiness among low-level choices.

These examples of coupled search processes are diverse, but we expect them to exhibit similar empirical patterns that are important to both researchers and managers. Figure 4 provides an overview of the patterns. Both wanderer and survivor effects should be stronger in stable environments than they are in turbulent environments (H3a and H3b). Combining H3a and H3b gives the strongest prediction of our simulations: dampening should be strongest and the performance impact of high-level choices hardest to detect in stable environments. In stable environments, wanderer effects should be stronger when the adaptation of high-level choices is time-driven (H1a) or mimetic (contrary to H2a) than when high-level adaptation is performance-driven and incremental. In stable environments, we also expect survivor effects to loom larger when adaptation is sparked by poor performance (H1b) and is mimetic (H2b) than when adaptation is time-driven or incremental.

Implications for Research

For theory, a better understanding of coupled search processes reveals the micro-structure of how a disconnect can emerge between high-level choices and performance outcomes, creating an apparent loose coupling (Weick, 1982) between high-level choices and performance. The most important implications of our work, however, are for empirical researchers who examine the performance impact of high-level choices such as organizational designs, management

Figure 4. Predicted degree of dampening, survivor, and wanderer effects.

<table>
<tr><th></th><th></th><th colspan="2">Stable Environment</th><th colspan="2">Turbulent Environment</th></tr>
<tr><td rowspan="2">Adaptation Mode</td><td>Incremental</td><td>Very stong dampening due mainly to strong wanderer effect</td><td>Moderate dampening due to wanderer effect and weak survivor effect</td><td colspan="2" rowspan="2">Weak dampening due to weak wanderer and survivor effects</td></tr>
<tr><td>Mimetic</td><td>Extremely strong dampening due mainly to strong wanderer effect</td><td>Extremely strong dampening due to strong wanderer and survivor effects</td></tr>
<tr><td></td><td></td><td>Time</td><td>Performance</td><td>Time</td><td>Performance</td></tr>
<tr><td></td><td></td><td colspan="4">Adaptation Trigger</td></tr>
</table>

traits, alliance networks, cognitive frames, or resource allocation processes. We offer these researchers sobering news, that the nature of coupled search processes may make it difficult to detect the impact of even very consequential high-level choices. In light of our findings, empirical researchers should begin their work by asking whether they face a coupled search process that includes multiple loosely coupled levels of choices, organizations that explore alternative high-level choices, and stickiness in low-level choices. If so, then they should be wary as they seek the performance impact of high-level choices. Especially in stable settings, much of the performance difference between firms with appropriate and inappropriate high-level choices may be dampened by wanderer and survivor effects. As a result, researchers may conclude, for instance, that there is equifinality (a range of high-level choices have the same outcome) even when there is not. For instance, the fact that firms with different designs or management attributes display similar performance at a point in time does not necessarily imply that these designs or attributes are equally effective. Hence our findings reveal a major challenge for testing claims about the appropriateness of high-level choices such as organizational designs or different top management team attributes. In this type of work, one usually makes a conceptual argument about why a particular design or management team should perform well, given certain contingencies. One then assembles a data set and tests the relationship between high-level choices and performance. Failure to find such a relationship is generally seen as a failure of the conceptual argument (for an example of such a test, see Doty, Glick, and Huber, 1993). Our simulations show, however, that even when the conceptual argument is correct and high-level choices create substantial performance differences, coupled search processes can dampen these differences significantly and throw doubt on the hypothesized relationship between high-level choices and performance.

Consider, for instance, social networks. In social networks, coupled search processes are common. Actors (firms or individuals) search at a low frequency for what they believe to be advantageous network positions; more frequently, they use their network positions to obtain helpful information (e.g., innovation leads) that directly affects performance. A long-standing debate in the social network literature concerns which type of network position is beneficial. Those in one camp have argued that actors benefit from being embedded in densely connected networks (Coleman, 1988), while those in another camp have argued that a focal actor benefits from bridging holes in disconnected and sparse networks (Burt, 1992). Most empirical studies that try to adjudicate between these theories have measured network position at a point in time and related this position to performance. The results of these studies have been very mixed, leading researchers to propose contingency factors that may distinguish conditions under which each type of position is more beneficial (Podolny and Baron, 1997). Our simulation results point to another reason why it might be difficult to find systematic evidence for the performance benefits of one or the other network position. If networks change over time and the performance of actors is sticky, differences in performance that might be generated by different network positions can be dampened away. For instance, occupying a structural hole may provide a firm with a long-lasting performance advantage, perhaps because it allows a firm to establish a reputation in its industry. But if the network structure changes over time because actors in the network create more ties to high-reputation firms, and if one observes firms only at the end of this process, one might conclude that the high performance of a focal firm is generated by the dense network, even though the current network structure has little to do with the current high performance of the focal firm.

On a more hopeful note, our work suggests four empirical strategies that may help researchers detect the performance effects of high-level choices. The first strategy is to conduct tests in turbulent settings, in which high-level choices must rise repeatedly to the challenge of finding good low-level choices. In such an environment, a firm cannot hide poor high-level choices by getting lucky once with its low-level choices. In light of this empirical strategy, we can reinterpret prior empirical work that has explored performance relationships in environments of differing turbulence. For instance, Cannella, Park, and Lee (2008) studied the effect of top management teams' intrapersonal functional diversity (whether members of the top management team were functional specialists or generalists) on firm performance. For industries that were stable and had low environmental uncertainty, the authors found hardly any effect of intrapersonal functional diversity on performance. In contrast, for turbulent environments with high environmental uncertainty, they found a strong positive effect. Given these results (and the theory they presented to motivate their study), they concluded that intrapersonal diversity is more important in uncertain environments. Although this contingency certainly could exist, our simulation results raise the concern that the "missing" result for stable environments was simply much harder to find because it was dampened away.

Another example is provided by research on the performance effects of the comprehensiveness of strategic decision-making approaches. Miller (2008) summarized a set of findings in this arena, noting that studies in turbulent or mixed-turbulence industries generally support the comprehensiveness approach, whereas studies in stable environments do not. Miller proposed that this puzzle may be resolved by a more complex contingency theory that posits differently shaped relationships between comprehensiveness and turbulence. Our results suggest another possible solution to the puzzle: comprehensiveness affects performance not directly but through a coupled search process; managers decide how comprehensive to be, they search for activity choices in light of their comprehensiveness, and the activity choices determine performance. If managers adjust their degree of comprehensiveness over time, wanderer and survivor effects can make it far harder to detect the performance effect of comprehensiveness in a stable setting than in a turbulent context. More broadly, our results imply that the test of environmental turbulence as a contingency factor is tricky because environmental turbulence can also influence the likelihood of finding a significant relationship. Non-findings in stable environments and significant results in turbulent environments do not necessarily imply that an environmental turbulence contingency exists.

A second empirical strategy, intended specifically to overcome the wanderer effect, is to trace back each firm's history to identify the high-level choice that first delivered the firm's current performance. One can then examine whether good high-level choices are disproportionately likely to have led firms to achieve strong performance. But even this approach may not uncover appropriate high-level choices fully. A firm's current performance can result from an entire sequence of prior high-level choices, so the high-level choice that was present when the current performance level was first achieved might not be the high-level choice most responsible for this performance. Our results suggest that individual high-level choices, such as designs, may not be the appropriate unit of analysis. Rather, the right unit of analysis may be different sequences of high-level choices (Siggelkow and Levinthal, 2003, 2005).

A third empirical strategy to combat dampening is to conduct research in settings in which firms cannot alter high-level choices easily, because dampening is driven by the potential for firms to move among high-level choices. This comes in two forms: (1) the potential for firms with good high-level choices and high performance to wander toward poor high-level choices without performance consequences, and (2) the potential for firms with poor high-level choices and low performance to move toward better high-level choices, leaving only lucky high performers among those with poor high-level choices. Conditions that limit such movement may

Figure 5. Dampening and the number of reorganization opportunities.

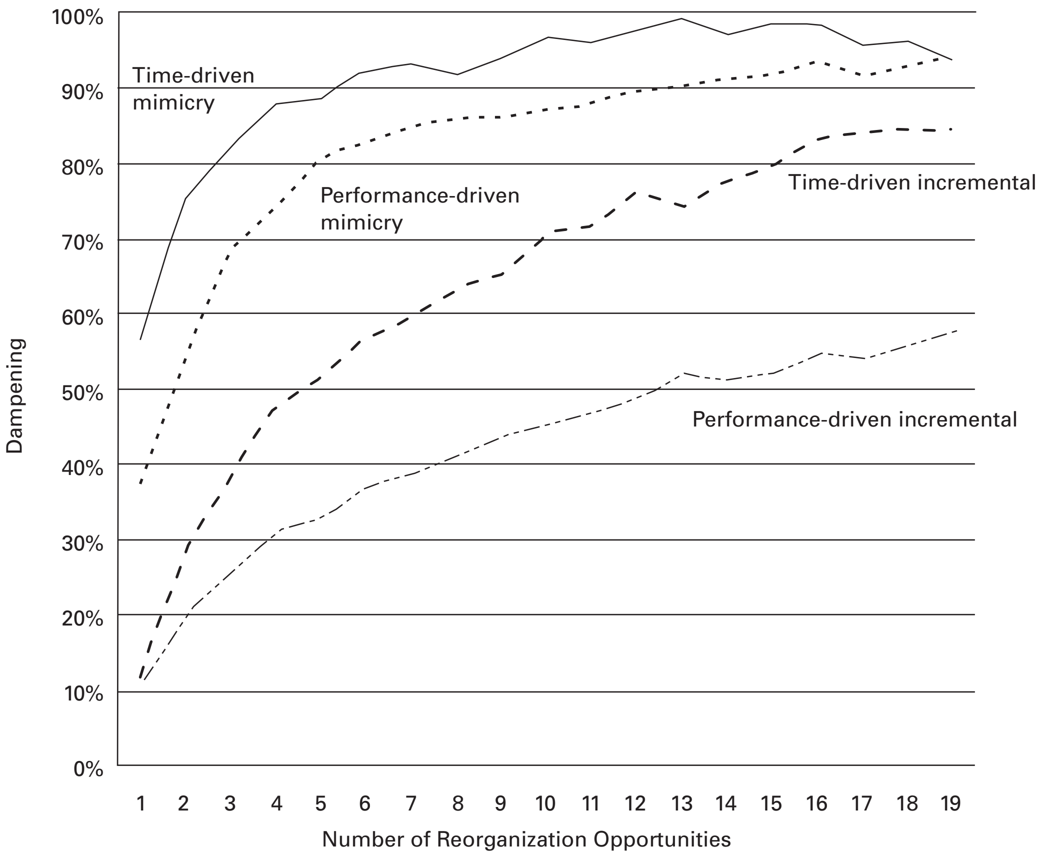

reduce the degree of dampening. For instance, one may suspect that if firms are allowed fewer opportunities to change high-level choices, dampening may disappear. Simulations, however, temper this hope. Figure 5 shows the degree of dampening that arises in our simulations in a stable environment after each round in which reorganization can happen (i.e., after every 50 periods). Dampening worsens with more reorganizations, but it is salient after firms have had only a few opportunities for reorganization. This is especially true when firms can mimic each other and leap to a new design rather than tweak a design incrementally. Only severe limits on the number of reorganization opportunities will make dampening far less damaging to empirical research.

A fourth empirical strategy is to assess the attractiveness of each high-level choice based not on the average performance of firms with that choice but on the choice's relative ability to enable a firm to improve its performance. Consider adaptation by time-driven mimicry. As shown in figure 5, by period 300 (after 5 reorganization opportunities), considerable performance dampening has already occurred. Among the top 20 percent of all performers, 17 percent are P = 1 firms and 24 percent are P = 5 firms. Thus it is difficult to discern which high-level choices are effective by asking which high-level choices stand-out performers use. In contrast, if one observes whether a firm has been able to improve between periods

250 and 300, it is much easier to discern good high-level choices. By that time, no firm that adopted a P = 1 design was able to improve, while 71 percent of all firms that did improve had P = 5 designs. Thus empirical researchers might fruitfully focus on firms that are improving performance, rather than ones that simply have steady high performance, when they try to pinpoint beneficial high-level choices.

Implications for Managers

Coupled search processes are especially salient for the most senior managers, who typically have special responsibility for high-level choices. For instance, Bower (1970) argued that the top layer of management influences performance in large part by making high-level choices about structural context, which in turn molds the resource allocation decisions that affect performance. Similarly, boards of directors may exercise their greatest influence on firm performance not by making specific decisions but by shaping the character of the management team, which then makes low-level decisions.

To make good high-level choices, senior managers must first discern the roots of performance differences. Our findings suggest that this will be difficult in contexts in which coupled search processes operate. For instance, a manager involved in a coupled search process who hopes to find new high-level choices to improve performance may tend to imitate the high-level choices of successful rivals. Given survivor and wanderer effects, however, this approach may fail. Successful firms that persist with a particular high-level choice may be atypical among firms that tried that choice (the survivor effect), and current high-level choices may be unrelated to current high performance (the wanderer effect).

The empirical strategies we described above for researchers have analogues for managers. Managers can more safely imitate the high-level choices of stand-out performers when they operate in turbulent environments or in settings that inhibited changes in rivals' high-level choices. They may wisely pay more attention to the high-level choices that first delivered strong performance for successful rivals than to the current choices of those firms. And they might target for high-level imitation those firms that are improving performance rather than firms that have high performance today.

For clarity of exposition, we assumed throughout our simulation effort that firms can mimic the high-level choices of successful firms but cannot imitate the low-level choices that directly drive performance. We feel this is a reasonable first approximation of reality, given that low-level choices often are numerous, detailed, shifting, and hidden deep within an organization. In practice, however, managers who want to imitate rivals must make an allocation decision: how much effort will they devote to understanding the high-level choices of successful rivals and how much to grasping low-level choices? Our research suggests that for imitators in turbulent environments, this allocation should shift toward high-level choices. Turbulence makes it easier to discern the high-level choices that deliver superior performance, and in turbulent environments, low-level choices must shift often, making information on competitors at that level obsolete quickly. In contrast, in stable environments, an imitator should devote more attention to understanding the low-level choices of stand-out performers. In such settings, dampening makes it hard to pinpoint good high-level choices, and an effective set of low-level choices—once copied—can remain useful for a long time.

Implications for Populations of Organizations

Although performance dampening and a failure of organizations to drift toward more effective high-level choices create problems for researchers and managers, lack of drift may be a helpful property for populations of organizations. As Hannan and Freeman (1989: 7) pointed out,

> Questions about the diversity of organizations in society might seem to have only academic interest. In fact, these issues bear directly on important social issues. Perhaps the most important is the capacity of a society to respond to uncertain future changes. Organizational diversity within any realm of activity . . . constitutes a repository of alternative solutions to the problem of producing sets of collective outcomes. These solutions are embedded in organizational designs and strategies.

Evolutionary processes that lead to a reduction of forms thus can weaken long-term population-wide performance. Traditionally, one might have thought of this as a clear tradeoff: an evolutionary process that leads to quick convergence on the optimal form (given current conditions) has a short-term advantage but may lead to a homogeneous population that cannot adapt should environmental conditions change. Conversely, a process that lets inferior forms survive leads to lower short-term performance for the population as a whole but retains a greater degree of diversity and associated adaptability. Our simulation results indicate that coupled search processes might achieve both diversity in high-level choices, which enhances long-term adaptability, and strong short-term performance. Poor high-level choices may be retained and serve as diverse, helpful seeds should the environment change. At the same time, firms with poor high-level choices show fairly high performance, either because they arose from "mutations" of firms with a history of better high-level choices or because firms with poor high-level choices were lucky to discover good low-level choices.

Coupled search processes may complicate the work of empirical researchers and managers, but they harbor potential benefits for a population of organizations as a whole: they preserve diversity in high-level choices without large performance penalties for seemingly wrong choices. Coupled search processes can both sustain population-level diversity and generate strong short-term performance. This is a powerful combination, and it may explain why coupled search processes are so prevalent among organizational phenomena. In studying organizations, we often take for granted that hierarchical, coupled search processes exist: firms search for structures and then for strategies within those structures, boards seek managers and then managers search for strategies, firms look for alliance partners and then for strong performance within their alliance networks, and so on. We rarely ask why search should so commonly take on a coupled architecture, rather than being a unitary process, for instance, with board members searching for strategies directly. Our analysis points to one potential answer. At a population level, such a coupled architecture breaks the usual tradeoff between short-term performance and long-term adaptability.

REFERENCES

Ahuja, G.
2000 "Collaboration networks, structural holes, and innovation: A longitudinal study." Administrative Science Quarterly, 45: 425–455.

Arthur, W. B.
1989 "On competing technologies and lock-in by historical events: The dynamics of allocation under increasing returns." Economic Journal, 99: 116–131.

Barnett, W. P., and O. Sorenson
2002 "The Red Queen in organizational creation and development." Industrial and Corporate Change, 11: 289–325.

Bower, J. L.
1970 Managing the Resource Allocation Process. Boston: Harvard Business School Press.

Brown, S. L., and K. M. Eisenhardt
1997 "The art of continuous change: Linking complexity theory and time-paced evolution in relentlessly shifting organizations." Administrative Science Quarterly, 42: 1–34.

Burt, R. S.
1992 Structural Holes. Cambridge, MA: Harvard University Press.

Cannella, A. A., J.-H. Park, and H.-U. Lee
2008 "Top management team functional background diversity and firm performance: Examining the roles of team member colocation and environmental uncertainty." Academy of Management Journal, 51: 768–784.

Carley, K. M., and D. M. Svoboda
1996 "Modeling organizational adaptation as a simulated annealing process." Sociological Methods and Research, 25: 138–168.

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

Coleman, J. S.
1988 "Social capital in the creation of human capital." American Journal of Sociology, 94: S95–S120.

Colombo, M. G., and M. Delmastro
1999 "Some stylized facts on organization and its evolution." Journal of Economic Behavior and Organization, 40: 255–274.

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

Dalton, D. R., W. D. Todor, M. J. Spendolini, G. J. Fielding, and L. W. Porter
1980 "Organizational structure and performance: A critical review." Academy of Management Review, 5: 49–64.

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

Davis, J. P., K. M. Eisenhardt, and C. B. Bingham
2007 "Developing theory through simulation methods." Academy of Management Review, 32: 480–499.

Djelic, M.-L., and A. Ainamo
1999 "The coevolution of new organizational forms in the fashion industry: A historical and comparative study of France, Italy, and the United States." Organization Science, 10: 622–637.

Donaldson, L.
2001 The Contingency Theory of Organizations. Thousand Oaks, CA: Sage.

Doty, D. H., W. H. Glick, and G. P. Huber
1993 "Fit, equifinality, and organizational effectiveness: A test of two configurational theories." Academy of Management Journal, 36: 1196–1250.

Galbraith, J. R.
1973 Designing Complex Organizations. Reading, MA: Addison-Wesley.

Galunic, D. C., and K. M. Eisenhardt
1994 "Renewing the strategy-structure-performance paradigm." In B. M. Staw and L. L. Cummings (eds.), Research in Organizational Behavior, 16: 215–255. Greenwich, CT: JAI Press.

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

Gresov, C., and R. Drazin
1997 "Equifinality: Functional equivalence in organizational design." Academy of Management Review, 22: 403–428.

Hambrick, D. C., and P. A. Mason
1984 "Upper echelons: The organization as a reflection of its top managers." Academy of Management Review, 9: 193–206.

Hannan, M. T., and J. Freeman
1977 "The population ecology of organizations." American Journal of Sociology, 82: 929–964.
1989 Organizational Ecology. Cambridge, MA: Harvard University Press.

Haunschild, P. R., and A. S. Miner
1997 "Modes of interorganizational imitation: The effects of outcome saliency and uncertainty." Administrative Science Quarterly, 42: 472–500.

Haveman, H. A.
1993 "Follow the leader: Mimetic isomorphism and entry into new markets." Administrative Science Quarterly, 38: 593–627.

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

Lawrence, P. R., and J. W. Lorsch
1967 Organization and Environment. Boston: Harvard Business School Press.

Lenox, M. J., S. F. Rockart, and A. Y. Lewin
2006 "Interdependency, competition and the distribution of firm and industry profits." Management Science, 52: 757–772.

Levinthal, D. A.
1997 "Adaptation on rugged landscapes." Management Science, 43: 934–950.

Levinthal, D. A., and H. E. Posen
2007 "Myopia of selection: Does organizational adaptation limit the efficacy of population selection?" Administrative Science Quarterly, 52: 586–629.

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

March, J. G., and H. A. Simon
1958 Organizations. New York: Wiley.

Meyer, J., and W. R. Scott
1983 Organizational Environments: Ritual and Rationality. Beverly Hills, CA: Sage.

Miller, C. C.
2008 "Decisional comprehensiveness and firm performance: Toward a more complete understanding." Journal of Behavioral Decision Making, 21: 598–620.

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.

Mintzberg, H.
1979 The Structuring of Organizations. Englewood Cliffs, NJ: Prentice-Hall.

Podolny, J. M., and J. N. Baron
1997 "Resources and relationships: Social networks and mobility in the workplace." American Sociological Review, 62: 673–693.

Porter, M. E.
1996 "What is strategy?" Harvard Business Review, 74 (6): 61–78.

Priem, R. L., D. W. Lyon, and G. G. Dess
1999 "Inherent limitations of demographic proxies in top management team heterogeneity research." Journal of Management, 25: 935–953.

Rivkin, J. W.
2000 "Imitation of complex strategies." Management Science, 46: 824–844.

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

Schoonhoven, C. B.
1981 "Problems with contingency theory: Testing assumptions hidden within the language of contingency 'theory'." Administrative Science Quarterly, 26: 349–377.

Siggelkow, N.
2002 "Evolution toward fit." Administrative Science Quarterly, 47: 125–159.

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.
2005 "Escaping real (non-benign) competency traps: Linking the dynamics of organizational structure to the dynamics of search." Strategic Organization, 3: 85–115.

Siggelkow, N., and J. W. Rivkin
2005 "Speed and search: Designing organizations for turbulence and complexity." Organization Science, 16: 101–122.

Simon, H. A.
1962 "The architecture of complexity." Proceedings of the American Philosophical Society, 106: 467–482.

Sorenson, O.
2002 "Interorganizational complexity and computation." In J. A. C. Baum (ed.), Companion to Organizations: 664–685. Oxford: Blackwell.

Staw, B. M., and L. D. Epstein
2000 "What bandwagons bring: Effects of popular management techniques on corporate performance, reputation, and CEO pay." Administrative Science Quarterly, 45: 523–556.

Strang, D., and M. W. Macy
2001 "In search of excellence: Fads, success stories, and adaptive emulation." American Journal of Sociology, 107: 147–182.

Teece, D. J., G. Pisano, and A. Shuen
1997 "Dynamic capabilities and strategic management." Strategic Management Journal, 18: 509–533.

Thompson, J. D.
1967 Organizations in Action. New York: McGraw-Hill.

Weick, K. E.
1982 "Management of organizational change among loosely coupled elements." In P. S. Goodman (ed.), Change in Organizations: 375–408. San Francisco: Jossey-Bass.

West, C. T., and C. R. Schwenk
1996 "Top management team strategic consensus, demographic homogeneity and firm performance: A report of resounding non-findings." Strategic Management Journal, 17: 571–576.

APPENDIX: Robustness Results

To test the robustness of our results, we changed a number of key variables and assumptions of the model. First, we increased the reorganization frequency by allowing firms to reorganize every 10 periods rather than every 50 periods. As reported in table A.1, results are very similar to those reported above. Dampening is significant, and the different adaptation triggers and adaptation modes influence the magnitude of the survivor and wanderer effects in similar ways.

Second, we varied the degree of correlation between successive landscapes in our turbulent setting. With turbulence, every 50 periods, each contribution value c_i is replaced by $\alpha * c_i + (1 - \alpha) * u$, where u is a new draw from a uniform distribution over the unit interval. In the main analysis, for turbulent environments, we use $\alpha = 0.2$ (the stable environment would correspond to $\alpha = 1$). As we vary α, the degree of dampening changes monotonically. The smaller is α, the smaller is the degree of dampening, because both survivor and wanderer effects decrease the more different the new environments are from the old environments.

Table A.1

Robustness Results

	Time-driven incremental	Perf.-driven incremental	Time-driven mimicry	Perf.-driven mimicry
1. *Reorganization every 10 periods; stable environment*				
Dampening	85%	63%	86%	93%
Survivor effect	7%	26%	7%	32%
Wanderer effect	73%	34%	73%	45%
2. *Reorganization every 10 periods; turbulent environment changing every 10 periods; no interdependencies among low-level choices*				
Dampening	10%	0%	24%	4%
3. *Reorganization every 50 periods; stable environment; high-bandwidth senior management*				
Dampening	91%	64%	98%	109%
Survivor effect	5%	14%	6%	28%
Wanderer effect	73%	45%	79%	65%
4. *Reorganization every 50 periods; turbulent environment changing every 50 periods; high-bandwidth senior management*				
Dampening	1%	-4%	23%	27%
5. *Reorganization every 50 periods; stable environment; high-bandwidth senior management; selection model*				
Dampening	28%	32%	20%	39%

Third, we changed parameters in a way such that P = 5 was no longer the highest-performing design. In particular, we studied the case of K = 0, i.e., no interactions exist among the activity choices. If such an environment is turbulent, fast improvement, rather than broad search, becomes the key driver of performance. Fast improvement, in turn, is facilitated by low levels of P. For instance, if the environment changes every 10 periods, firms with P = 1 or P = 2 have the highest performance. As reported in the table, we find again that in turbulent environments dampening is greatly reduced. Thus our results appear not to depend on P = 5 being the optimal design.

In the next set of tests, we increased the capability of senior management. In our main model, in each period, senior management evaluates one randomly chosen composite alternative. Thus it is possible that in one period, senior management could have examined one alternative and rejected it, and in the next period senior management could have considered this alternative again. In our robustness check, we endowed senior management with vastly higher capabilities (high-bandwidth senior management). In particular, we assumed that in each period senior management could evaluate all possible combinations of the alternatives proposed by department managers. For example, when P = 5, senior management would in each period consider 24 alternatives plus the status quo. In this way, no good combination from those possible, given department managers' recommendations, is left unconsidered. As reported in the table, the results are very consistent with our main results. For stable environments, performance dampening is significant and even higher than in our main analysis, and the sizes of the survivor and wanderer effects follow a similar pattern as before. Likewise, in turbulent environments, dampening is reduced significantly.

Lastly, we created an explicit selection model, in which firms, not just organizational designs, die and new firms are born. In particular, we assumed that every 50 periods, the three firms with the lowest performance are eliminated and replaced by three new firms. The new firms start at the same location as all firms did in period 0. Each new firm adopts the level of P of high-performing firms. In particular, the probability that a new firm will adopt the level of P of firm i is $\text{performance}_i/\Sigma_j \text{performance}_j$, where the sum is over all existing firms j. This model, by itself, generates large performance dampening through a strong survivor effect. If a firm is not selected out despite its poor structure, it must have been lucky to have found a good combination of activities. In particular, we found a strong selection effect across landscapes. For instance, in landscapes in which the starting point is very close to the global peak of the landscape, firms with low levels of P perform very well. In this setup, we find the selection model by itself creates a dampening effect of 78 percent relative to the baseline of no selection. On top of this explicit selection model, we allowed our firms to adapt their structures as in the main analysis of the paper. As reported in the table, an additional dampening effect can still exist. Though this dampening effect appears to be smaller, it is measured relative to a baseline that is already dampened by 78 percent. Thus even in a model that contains explicit selection, an additional dampening effect created by the coupled search process can arise.