07-106

Coupled Search Processes: Why is it so difficult to find that organizational design matters?

Nicolaj Siggelkow
Jan W. Rivkin

Coupled Search Processes:
Why is it so difficult to find that organizational design matters?

Nicolaj Siggelkow
2211 Steinberg-Dietrich Hall
Wharton School
Philadelphia, PA 19104
V: (215) 573-7137
F: (215) 898-0401
siggelkow@wharton.upenn.edu

Jan W. Rivkin
239 Morgan Hall
Harvard Business School
Boston, MA 02163
V: (617) 495-6690
F: (617) 495-0355
jrivkin@hbs.edu

Special thanks to Howard Brenner for programming help. The authors would like to thank participants at the Organization Science Winter Conference 2007 and seminar participants at Harvard Business School and at 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.

Coupled Search Processes:
Why is it so difficult to find that organizational design matters?

Abstract:

Organizational design affects performance via coupled search processes. At low frequency, managers search for appropriate organizational designs. At higher frequency, managers use designs to search for high-performing operational choices. The two searches are coupled: organizational design molds the choice among operational alternatives, and performance feedback from operational choices shapes design. Our simulation model shows how coupled search processes can dramatically obscure the true impact of design on performance, confounding empirical research. We identify research strategies for tackling this difficulty; discuss population-level advantages of coupled search processes; and highlight implications for analogous coupled search processes that shape networks, cognition, and capabilities.

Organizational performance is often driven by coupled search processes that operate at different time scales. Consider, for instance, the impact of a firm's formal organizational design on its economic performance. Intermittently, at a relatively low frequency, managers search for an effective set of design choices: an allocation of decision rights, an incentive system, a system of information collection and flow, and so forth. Such choices tend to be stable over time, persisting for years. These design choices have a profound effect, because they shape a second, higher-frequency search process: the process of making day-to-day operational choices concerning pricing, sales calls, production levels, shipping, procurement, etc. Such operational choices have large and immediate impacts on the costs a firm incurs and the buyer value it creates. The outcomes of the low-frequency search for an organizational design mold the high-frequency search for operational choices, which in turn shape economic performance.

Coupled search processes like these are common. At a low frequency, a firm searches for partners and develops its position in a network of allied firms; at a faster cadence, the position influences the firm's ability to tap into knowledge and to search for innovation (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). Gradually, a nation might develop and adapt its legal Constitution; the Constitution then guides the formation of policies and laws at a faster pace, and those policies and laws – not the Constitution – have the direct impact on the country's prosperity (North and Weingast, 1989; Glaeser *et al.*, 2004).

An enduring goal of organizational research is to examine the low-frequency objects of search and to isolate their (indirect) impact on performance. How, for instance, do formal organizational design choices, a firm's position in an alliance network, managers' cognitive frames, characteristics of the resource allocation process, or dynamic capabilities affect organizational performance? How do

Constitutional provisions influence national prosperity? Clearly, such questions are of first-order interest to students of organizations. In this paper, we show that to address these questions, researchers must grasp the consequences of the two-level, asynchronous search processes that stand between the low-frequency objects and performance. An understanding promises two related benefits. First, it can shed light on the mixed results that have plagued past research. For instance, empirical studies of the relationship between organizational design and performance have resulted in a vexingly large number of ambiguous findings despite strong theoretical arguments that such a relationship should exist (Donaldson, 2001). Why are clean results so hard to obtain? What forces dampen relationships that should, in theory, exist? Second, an understanding can guide empirical work by identifying research settings in which those dampening forces are particularly prevalent or muted.

The arguments we make below apply to a wide variety of organizational contexts. For the sake of concreteness, however, we focus on the challenge of organizational design. Specifically, we develop an agent-based simulation model in which firms are engaged in a two-level search process: a search for how to configure a particular aspect of organizational design (the degree of information flow between hierarchical levels) and, subject to the chosen design, a search for a configuration of concrete activities that directly determine performance. Examination of this multi-level search process identifies two mechanisms that compress performance differences among firms with different design choices and thereby make the design/performance relationship difficult to observe. The first mechanism is a *survivor* effect. To retain a design that typically produces poor performance, a firm must be faring relatively well in the search for effective activities. Thus, firms that persist with “poor” designs must have levels of performance that are much higher than the typical level for that design. The second mechanism is a *wrong-attribution* effect. Firms that attain high-performing configurations of decisions with a “good” design can subsequently experiment with other designs. If one observes firms during such an experimentation stage, one might wrongly attribute current performance to the firm’s current design, even though the firm’s current design had little to do with the firm’s current performance. In addition to identifying factors that dampen the design/performance relationship, we pinpoint the settings in which those factors are likely to loom especially large. In particular, we show that dampening is far worse in static settings than in dynamic settings. Our concluding section generalizes these findings to questions other than the relationship between organizational design and economic performance.

PRIOR LITERATURE AND PROPOSITIONS

Like many of the coupled search processes mentioned above, the process related to organizational design has produced a mixed literature. On the one hand, theoretical work gives us strong reasons to expect organizational design to have a meaningful impact on performance. On the other, that impact has been difficult to detect empirically. We consider the reasons that design should affect patterns in performance; the mixed empirical results; and prior explanations for the mixed results.

Expected Impact

An organization's design is "the sum total of the ways in which it divides its labor into distinct tasks and then achieves coordination among them" (Mintzberg, 1979: 2). Empirical work has focused primarily on formal aspects of design, such as the allocation of decision rights, the flow of information and material, reporting relationships, decision-making processes, and incentive systems; it has paid less attention to informal components such as social networks, corporate culture, and norms. The notion that the formal design of an organization's work affects its performance dates back at least to the work of Taylor (1911) and Fayol (1916 [1949]) on the scientific organization of labor. While Taylor emphasized that organization could boost performance by codifying and spreading efficient practices, later authors highlighted other causal links between organization and performance. Barnard (1938: 256), for instance, argued that organizational design could influence performance by affecting coordination among workers – essential since "the quality of coordination is the crucial factor in the survival of organization." In contrast, Galbraith (1973) saw organizations as information processors; well-designed organizations would process information in the manner demanded by environmental conditions and would thrive as a result. Agency theorists from Jensen and Meckling (1976) argued that organizational design, especially the structure of incentives, would influence performance by affecting the motivation of actors to exert effort rather than shirk. Formal design might also shape the informal design and social conditions of a workplace, thereby affecting performance (Nadler and Tushman, 1997). Regardless of the causal path, a dominant theme in prior literature is:

PROPOSITION 1: Appropriate organizational designs are associated with high relative performance.

Two aspects of this proposition deserve attention. First, a great deal of the literature in this stream of work has emphasized how the definition of "appropriate" design is contingent upon environmental conditions. Appropriate design may hinge, for instance, on the degree of environmental uncertainty (Lawrence and Lorsch, 1967), the pace of change (Burns and Stalker, 1961), or the nature of a firm's technology (Woodward, 1965). Second, whether organizational design acts on performance via efficient work practices, coordination, information processing, or motivation, the effect of design is indirect: it affects performance by shaping operational actions and decisions that, in turn, influence profitability. Thus we can envision two coupled search processes at work simultaneously, with the search for a good design molding the search for operational actions and decisions. Of the two search processes, the search for an appropriate design is likely to unfold at a slower pace, reflecting the inertia of organizations' structural features (Hannan and Freeman, 1977).

Two streams of prior research – on organizational learning (e.g., March and Simon, 1958) and organizational ecology (Hannan and Freeman, 1989) – have examined how the search for appropriate design will influence the relative prevalence of different designs among a population of firms. Within the organizational learning paradigm, it is postulated that managers within firms actively search for designs that are appropriate. Search is triggered by firms' performing under their aspiration level, which may capture either some absolute level of desired performance or some level relative to that of competitors (Cyert and March, 1963). The search process herds firms toward appropriate designs, leading to "a substantive prediction of equilibrium whereby the advantaged forms should prevail (often referred to in the sociological literature as isomorphism)" (Carroll and Harrison, 1994: 722). March and Olsen (1989) have referred to the assumption of convergence of firms onto an appropriate form as "historical efficiency."

While the organizational learning paradigm emphasizes firm change as the driving force towards historical efficiency, the organizational ecology paradigm stresses the role that births and deaths of firms play in shaping the diversity of organizational forms. Two key mechanisms are assumed to influence births and deaths, legitimation and competition, both of which are density-dependent. While the assumed drivers behind changes in the distribution of forms within a population are different, the dynamics produced by ecological reasoning also tend to point to a gradual move of a population of firms towards more appropriate designs. In either case, historical efficiency would imply:

PROPOSITION 2: Over time within a population of firms, appropriate designs become more frequent than inappropriate designs.

Mixed Evidence

Despite the strong and diverse theoretical arguments for Propositions 1 and 2, empirical support for them has been difficult to obtain. Dalton *et al.* published a critical review of the evidence in 1980, arguably the time at which researcher interest in this arena peaked. Surveying scores of studies on the direct impact of organizational structure on performance, Dalton and his co-authors found robust support for only one of the nine relationships they examined. (The positive relationship between subunit size and performance, if performance was measured by the degree of absenteeism, was robust across several studies.) Table 1 quotes the authors' summary of a number of relationships and reflects their general sense of exasperation (Dalton *et al.*, 1980).

A separate set of studies examined the impact of design on performance contingent on environmental conditions. A number of early and prominent tests of such contingency hypotheses found little support. Mohr (1971), for instance, hypothesized that autocratic supervision of work groups would be most effective for routine jobs while democratic supervision would succeed for nonroutine jobs; yet he found no support. Similarly, Pennings (1975) found that an organization's fit between environmental and structural conditions had little relationship to its effectiveness. In defense of contingency theory, Donaldson (2001) surveyed a set of papers that are more supportive. He concluded that, “More studies support the positive relationship of fit and performance than fail to support it. Moreover, the later and more sophisticated research provides more support for contingency theory than the earlier research” (p. 226). This begs the question of why it has taken such effort and sophistication to uncover relationships that have such strong theoretical foundations.

<INSERT TABLE 1 ABOUT HERE>

Explanations

We consider two potential answers to this question. The first is that the relationships are not valid despite the theoretical foundations. For example, if there is widespread equifinality (Doty, Glick and Huber, 1993; Gresov and Drazin, 1997) and many different organizational designs will produce the same levels of performance (Levinthal, 1997), then it becomes hard to specify what design is “appropriate,” and Propositions 1 and 2 lose their meaning. Similarly, a multiplicity of environmental contingencies may create conflicting structural requirements and undermine the notion of a single appropriate design (Galunic and Eisenhardt, 1994). Moreover, if managers choose organizational designs that they view as legitimate (Meyer and Scott, 1983) rather than ones that are effective, Proposition 2 loses its force. Herd behavior may lead a population of firms to lock in on a design that does not optimize performance (Staw and Epstein, 2000). Lastly, if the benefits of changing to another organizational structure are outweighed by the direct costs of this adaptation, a firm may stick with its “inappropriate” structure.

The second possible answer is that the predicted relationships – though fully valid – are hard to detect. Detection requires that we measure constructs such as autonomy, uncertainty, and performance, each of which is difficult to monitor on its own. The constructs interact with one another, and results may depend delicately on fairly arbitrary choices of functional forms (Schoonhoven, 1981). If one does not control for interacting elements, the underlying interdependencies among elements of design 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." Moreover, if Proposition 1 is valid and design affects performance, then selection pressure and firm adaptation may drive out heterogeneity in organizational design. Without variation in design, of course, it is impossible to detect the impact of design on performance. Any variation in design that persists may reflect heterogeneity across firms in unobserved factors, and any differences in performance that remain may reflect the unobserved heterogeneity, not a true impact of design on performance.

We fully accept this set of explanations for the mixed empirical evidence, but we will argue that the set is incomplete in an important way. In particular, none of these explanations reflects the coupled, asynchronous search processes that generate the underlying relationship between design and performance. These processes alone, without any of the explanations offered so far, can make it difficult to detect the performance impact of organizational design. To clarify how they can do so, we construct a simulation model in which none of the alternative explanations are valid: in our model, there is one, truly appropriate design; managers pursue performance and not legitimacy; change of organizational structure is costless; relevant constructs can be measured perfectly; interactions among design elements do not interfere with measurement; and selection and adaptation do not drive out variation in design. Even in this setting, coupled search for good designs and for high-performance operational choices can obscure the link between design and performance.

Modeling Precedents

A number of prior studies have used a simulation methodology to assess population-wide changes in the distribution of organizational forms. Unlike the present study in which firms are searching, and thus are changing their structures, Carroll and Harrison (1994) treat "organizational change as a selection process, meaning that changes in the number of organizations with a particular form occur through fluctuations in founding and failure rates rather than through transformations of existing organizations having other forms" (p. 724). They find that inferior forms may outlive more appropriate forms, if firms with inferior forms are given a head-start. In our model below, we will even the playing field completely, by having an equal number of firms with different designs start the evolutionary process at the same time.

Carley and Svoboda (1996) allow organizations to change their designs over time along three dimensions: the number of employees, the reporting relationships between these employees, and the task allocation among the employees. Organizational change is modeled as a stochastic, simulated annealing process, i.e., as a process that is more likely to occur for lower-performing organizations than for higher-performing organizations. Unlike our model, Carley and Svoboda's model assumes that organizations can evaluate the benefits of proposed organizational changes in an "off-line" manner (Gavetti and Levinthal, 2000) by correctly anticipating whether a change will be performance-enhancing. Moreover, given the more complex structure of the modeled organizations, interactions among the various design elements play a key role.

The model by Strang and Macy (2001) comes closer to the model presented in this study, as firms are allowed to adopt an innovation both through own experimentation and through imitation of other firms. In their model, the innovation that is being adopted has a direct impact on a firm's performance. Since we model organizational design specifically, we do not need to assume that different designs have a direct performance effect, but only that designs influence the second search process that our firms are engaged in, i.e., the search for good activity configurations.

Lastly, also using an agent-based simulation model, Siggelkow and Levinthal (2003; 2005) analyze the performance effects of changing organizational designs. In their model, firms are, however, exogenously forced to change their designs. Thus, their model does not capture a coupled search process in which firms endogenously choose when and how to change their designs.

MODEL

Our goal is to study the consequences of coupled, multi-level search processes on the ability to detect underlying relationships between the objects over which firms conduct their lower-frequency search and ensuing performance. To make the analysis more concrete, we focus on the context of organizational design. We model firms that are constantly engaged in search for higher-performing choices of activities. This search process is intimately influenced by the firm's current design. At a slower, intermittent pace, firms are also searching for designs that may allow them to improve performance even further.

Agent-based simulation modeling is an attractive methodology to make headways on these issues. Broadly, "[s]imulation 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) – all of which are true in the context of the impact of design on performance. More specifically, agent-based simulation allows us to model the two search processes directly and let the consequences of these processes emerge from the simulation. At the same time, the model allows us to control a host of other issues that would be particularly difficult to control for in empirical studies. First, we can study a situation in which there exists a true underlying relationship between design and performance. Thus, if we find that this relationship is hard to detect once the coupled, multi-level search processes operate, it is not due to the fact that this relationships does not exist. Second, we can make the organizational design problem as simple as possible in order to avoid issues of equifinality or interactions among various elements of design. In particular, we focus on only one element of organizational design, which we model along a single dimension: the degree to which, in a hierarchical firm, lower-level managers are allowed to narrow the information that flows to superiors, i.e., the degree of their 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 single-dimensional variable, even more problems for analysis are likely to arise when designs are more complex. Third, the model allows us to control for other reasons that may lead to a weak design/performance relationship. For instance, firms are all endowed with the same starting conditions and capabilities for change. Lastly, the model gives us control over the environments in which firms operate, thereby providing us with the opportunity to investigate for which environmental contingencies our results are more likely to hold.

In sum, our model contains three parts: environments in which firms operate; high-paced search for good configurations of activities; and lower-frequency search for better organizational designs. We discuss each part in turn.

Environments

We conceptualize a firm's management team as facing a system of interdependent decisions (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, whether to pursue basic R&D or not, etc. Moreover, a number of these 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 set of interdependent choices to its performance.

In our model, each simulated firm has to make decisions concerning N choices, 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 ($a_2 = 0$). In total, thus, a firm has 2^N possible configurations of choices.

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; Ethiraj and Levinthal, 2004). 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 choices $\mathbf{a} = (a_1, a_2, \ldots, a_N)$ to a performance value $V(\mathbf{a})$.

We create two types of performance landscapes 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). In *simple environments*, all decisions are independent from each other, i.e., the contribution of each individual choice a_i to firm payoff V is affected only by the state (0 or 1) of the choice itself. In contrast, in *complex environments*, decisions are highly interdependent. In this case, the contribution of each individual choice a_i depends on the states of all other choices $\mathbf{a_{-i}}$.

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 given set of choices **a** is then given by the average of the N contributions: $V(\mathbf{a}) = [c_1(a_1, \mathbf{a_{-1}}) + c_2(a_2, \mathbf{a_{-2}}) + \ldots + c_N(a_N, a_{\mathbf{-N}})]/N$. Throughout the simulations, we set N to 8. (A larger value for N does not change the qualitative results reported in the paper.)

The second distinction we make is between environmental *stability* and *turbulence*. Firms in stable environments operate on the same landscape for their entire life-histories (200 periods). In turbulent environments, the landscape undergoes "correlated" shocks in periodic intervals. In particular, once a landscape is created, every 10 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.

Fast-paced Search for Good Configurations of Choices

In the first period of each simulation, firms are placed on a randomly chosen point of the landscape, i.e., endowed with a randomly chosen set of choices. In subsequent periods, each firm tries to find higher-performing sets of choices, i.e., higher locales on the performance landscape.

In modeling how firms search a performance landscape, we focus on a hierarchical, two-level design, in which lower-level department managers search for improvements in their respective departments and send proposals to upper-level managers, who in turn, coordinate low-level 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 setting for three reasons. First, the degree of lower-level 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, given our particular model set-up, a clear relationship between richness of information flow and performance exists. In particular, in stable, complex environments a firm benefits from a rich information flow (and hence, relatively low autonomy), while in dynamic, simple environments a narrow vertical information flow (i.e., a higher degree of autonomy) is more appropriate (Siggelkow and Rivkin, 2005).

It is important to note that we do not claim that this particular hierarchical structure is any way "optimal" – other structures including ones that stress more horizontal information exchange may outperform this one (Siggelkow and Rivkin, 2005). However, to simplify the organizational search process as much as possible, we focus on adaptation of design within the hierarchical structure.

Search for higher-performing sets of choices proceeds as follows. We assume that the firm's eight decisions are split between two department managers. 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 her department and ranks them from most attractive to least attractive for her department. (An alternative is "local" if it only differs in 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.)

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 upper management. A low level 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 level of P reflects a firm in which senior managers want to review many alternatives themselves. Thus, P captures the richness of the vertical information exchange. P is bounded by 1, since department managers need to send up at least one proposal, and by 5, since department managers can at most send up 4 local alternatives and the status-quo.

Upper management focuses on coordinating the actions of the firm's two departments. In particular, from all possible combinations of departmental proposals and status quo choices, upper management selects one composite alternative at random, assesses it in light of the interests of the firm as a whole (that is, evaluates the overall V(**a**)), compares it to the status quo, and implements the option if it yields a higher payoff for the firm than the status quo does. The configuration implemented by upper management forms the starting point for further search at the departmental level in the next period. One should note that the random selection of composite alternatives by upper management creates a source of heterogeneity between firms with the same value of P. A firm whose upper management happens to have picked a good combination of proposals may improve its performance faster than a firm with a less lucky upper management.[1]

Lower-frequency Search for Better Organizational Designs

While firms search for better activity configurations in every period, the search for better organizational structures happens at a lower frequency and is “on-line” (Gavetti and Levinthal, 2000), i.e., firms need to implement a new design and experience it before they can ascertain what performance they can achieve with this design. In particular, we assume that firms consider design changes every 10 periods. (In the later robustness section, we analyze firms that consider design changes every 5 or 25 periods.) When a firm considers a design change, it first has the opportunity to revert to its most recent design should the performance achieved by the current design be worse than the performance under its former design. 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.)[2] If a firm does not revert, it proceeds to search for a better design. The literature on organizational change identifies a number of different drivers for change and a number of ways of how firms may change. We focus on two frequently described drivers (***time-driven***

[1] The random choice of composite alternatives to be considered by upper management is the only source of heterogeneity, since department managers always evaluate all local alternatives and always send up their most preferred P alternatives. Thus in two firms with the same choice configuration and the same level of P, the same sets of alternatives will be sent up to upper management.

[2] In more detail, every 10 periods, a firm will 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 that the firm had under its most recent design, the firm will revert to the old design, which completes the adaptation process for this 10-period cycle. If the performance is better (or equal), the firm will consider the possibility of changing its design as described below. In a sense, if a firm was successful in a past experiment with a new design, the firm will continue to experiment with new designs, since the firm can always return to its old design.

and *performance-driven*) and two commonly noted modes of changing (*incremental* and *mimicry*).

We say that a firm engages in *time-driven* change when a firm that decides to change its design (rather than revert to the old one) will always change its design, regardless of its performance. This is akin to a "time-paced" evolutionary path "in which change is keyed to the passage of time" (Brown and Eisenhardt, 1997: 25).

A firm is said to follow *performance-driven* change if the firm adjusts its design with a probability that is inversely correlated to its performance. This driver captures an element of organizational behavior suggested by March and Simon (1958) and subsequently developed by many others: firms with low performance are likely to be below their aspiration level of performance and are thus triggered to search for improvements; firms with high performance are more likely to be above their aspiration level and thus are less likely to search for improvements. 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 Π_{min} be the minimum performance. Then the focal firm will reorganize with probability $p = (\Pi_{max} - \Pi)/(\Pi_{max} - \Pi_{min})$.[3] This implies, for instance, that the firm with the best performance would never reorganize and the firm with the worst would always reorganize.[4] One should note that with performance-driven change the two search processes that the firm is engaged in are coupled in both directions. The organizational design influences the fast-paced search for better choice configurations; and the degree to which the fast-paced search for better choice configurations is successful influences the frequency of the search over organizational designs.

Once a firm decides to reorganize, it has to determine how to change its design. In our case, the only element that a firm can change is P, the richness of information exchange. We distinguish two change modes: *incremental* changes and changes employing *mimicry*. Incremental change captures internal, small-scale experimentation. Incremental and local search, in which "members of organization begin by looking at close alternatives to improve performance" (Barnett and Sorenson, 2002: 291), has often been

[3] In the special case in which all firms have the same performance, i.e., $\Pi_{max} = \Pi_{min}$, it is assumed that reorganization occurs for sure.
[4] This mechanism only requires firms to know their performance relative to other firms, but not relative to the highest performance possible, as assumed, e.g., by Strang & Macy (2001).

used in formal models as a reflection of managerial behavior (March and Simon, 1958). Empirical support for this assumption, in particular with respect to changing the parameter of our interest, is provided by Colombo and Delmastro (1999). In a sample of 438 Italian metalworking plants, organizational change with respect to decision autonomy at different levels in the organization (which is inversely correlated with our parameter P) was "characterized by a process of marginal adaptation instead of radical modification" (p. 264). In our model, with incremental experimentation, a firm either increases or decreases P by one unit with equal probability.[5]

If a firm employs *mimicry*, the firm will imitate the level of P of another firm. Following prior conceptual (DiMaggio and Powell, 1983), empirical (Haveman, 1993), and modeling work (Strang and Macy, 2001), we assume 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 an organizational feature that appears to have good performance consequences. In particular, if Π_i is the performance of any given firm *i*, then a focal firm *m* will copy another firm *j*, $j \neq m$, with probability $p = \Pi_j / \Sigma_{i \neq m} \Pi_i$.[6] As per timing, each firm chooses the design it intends to mimic before any firm actually implements its changes. Thus firms do not accidentally mimic a just-adopted design of other firms.

Combining the different drivers and modes of organizational search captures a range of organizational search processes: time-driven and performance-driven incremental adaptation, and time-driven and performance-driven mimetic adaptation.

[5] Should P reach its upper or lower bound, the firm will move away from the boundary with probability 0.5 and stick at the boundary with probability 0.5.
[6] 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; Firm 5: 0.8. If Firm 5 reorganizes using the mimicry direction rule, it has a 0.3 / (0.3 + 0.7 + 0.4 + 0.6) = 15% chance of copying Firm 1's design, a 0.7 / (0.3 + 0.7 + 0.4 + 0.6) = 35% chance of copying Firm 2's design, and so forth.

RESULTS

To focus on the effect of organizational design, we eliminate all other heterogeneity within each simulation: In each simulation we place all firms on the same, randomly chosen starting point on the landscape. Thus firms start off with exactly the same choice configuration (and performance). The capabilities of managers in all firms are the same; i.e., they all evaluate the same number of alternatives, as described in the previous section. Likewise, all firms start at the same time. At the beginning of each simulation, firms differ 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 employ the same search process for new organizational designs. For each landscape we run the simulation for 200 periods, and we repeat each simulation on 2,000 landscapes. (Each landscape is created by randomly re-drawing all contribution values c_i.) The reported results are thus always averages over 2,000 observations. Performance values are given as percentages of the highest performance possible in each landscape. Each firm considers a reorganization every 10 periods.

For most of the results, we report the final design that a firm adopts by the end of the simulation, i.e., by the end of 200 periods. We determine the final design of each firm as follows: If the firm would have reverted in period 201 to a prior design, we count this prior design as its final design. Otherwise, the final design is the design observed in period 200.

In discussing the results, we find the notion of a sticking point to be helpful. Sticking points are choice configurations from which a firm will not move (Rivkin and Siggelkow, 2002). That is, at a sticking point, there exists 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 environment does not change afterwards).

We discuss, in turn, results for static, complex environments, results for turbulent, simple environments, and robustness with respect to various model parameters.

Results in Stable, Complex Environments

The first set of analyses considers firms that operate in stable, yet complex environments. Column 1 in Table 2 reports an important benchmark: the average performance in period 200 of firms that cannot change their designs. The results show that in complex, stable environments a rich information flow is very valuable: performance increases monotonically with P. For instance, firms with P = 5 have an average performance of 0.920, whereas firms with P = 1 have an average performance of only 0.843, a statistically highly different value. The intuition behind this benchmark result is that highly complex environments are very "rugged" with many sticking points that create the possibility firms will get stuck with suboptimal choices. Rich information flow increases the degree of exploration that a firm engages in, thereby increasing its long-run performance. Given the relatively large performance difference between P = 1 and P = 5 firms and the absence of confounding circumstances, the set-up seems ideal for finding support for Propositions 1 and 2: At the end of the simulation, one would expect firms with more appropriate designs to outperform firms with less appropriate designs (Proposition 1). Moreover, one would expect firms to find the more appropriate design so that better designs appear more frequently in the population of firms than do less appropriate designs (Proposition 2).

<INSERT TABLE 2 ABOUT HERE>

Time-driven incremental search process. In the first simulation, we consider the time-driven incremental search process. Thus, every 10 periods each firm will reorganize by adjusting its P up or down by 1 (or revert to its prior design). Column 3 in Table 2 shows the percentage of firms that end up with different levels of P, while column 2 shows the average performance of these firms. For these firms and for the benchmark firms in column 1, Figure 1 graphs average performance as a function of P.

Two results 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% each), and by the end, we find 19.4% firms with P =1 and 20.3% firms with P = 5. This difference is statistically significant, but its magnitude is clearly underwhelming.[7] For firms following a time-driven incremental search process, we do not find much support for Proposition 2.

Prior studies have noted that non-optimal features might survive in a population of firms, but they have argued that such non-optimal features will undermine performance (Proposition 1). To the contrary, our second main result is that the performance penalty associated with a poor design may be drastically smaller than a benchmark would lead one to expect (column 2 vs. column 1; Figure 1). Whereas the performance difference between the best (P = 5) and worst (P = 1) designs was 0.077 when firms could not change their designs, now the performance difference between firms that have the best and worst designs is seven times as small (0.011). Thus, 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 $\Pi_0(P1)$ and $\Pi_0(P5)$ be the performance of firms with P = 1 and P = 5 when they cannot change their design, and let $\Pi(P1)$ and $\Pi(P5)$ be the performance of firms that have P = 1 and P = 5 at the end of each simulation. Then *dampening* $= 1 - [\Pi(P5) - \Pi(P1)]/[\Pi_0(P5) - \Pi_0(P1)]$. For instance, a *dampening* of 1 would imply that the entire performance difference is dampened away.

As reported in the first row of the bottom panel of Table 2, about 85% of the performance difference between firms with best and worst designs is dampened away by the coupled search processes. One should note that this performance dampening is driven almost exclusively by higher performance of firms with poor designs, rather than lower performance of firms with good designs. (This result will hold for other drivers and change modes as well.) Consequently, to understand performance dampening, we need to understand why firms with seemingly poor designs can have high performance.

<INSERT FIGURE 1 ABOUT HERE>

Two mechanisms dampen performance differences. First, there is a *survivor effect*: Structures that survive an evolutionary adaptive process despite being inappropriate must have been present in firms that

[7] On average, each of the 2,000 landscapes contains 5.08 firms with P = 5 and 4.85 firms with P = 1. Using the standard deviation of number of firms of each type across landscapes, we can compute the significance of the difference in the mean values. The p-value is below 0.001.

had some other advantage (for instance, lucky upper management who happened to evaluate some very good alternatives).[8] Thus if one observes a firm with an inappropriate design, that firm must have had relatively high performance (compared to an "average" firm with this design), thereby belying the generally poor outcome associated with this design. To gauge how much of the dampening effect is driven by P = 1 survivors, we identify survivors by tagging firms whose performance at the end of period 10 equals the performance in period 200 and whose starting design equals its final design. These firms found a sticking point with their initial design, never found a better configuration of choices, and ended up with their original design. Given that firms using the time-driven, incremental search process change their design every 10 periods, we may expect relatively few survivors to exist. Indeed, only about 10% of all P = 1 firms are survivors. As expected, the performance of these survivors, 0.894, is significantly higher than the performance of the average non-changing P = 1 firm of 0.843. Overall, survivors contribute 8% of the performance dampening effect (see the second row in the bottom panel of Table 2).[9]

The second mechanism that leads to high performance among firms with poor designs is a *wrong-attribution effect*: A firm might start with a good design and find a good configuration of choices; after this, the firm may continue to experiment with its design and may adopt a poor design. This poor design may prevent the firm from improving its performance, but if it does not pull the firm from its current set of choices, the firm will retain its current, relatively high performance. If one correlates final performance with final design, one would attribute the firm's high performance to the final design; yet the final design had little to do with the firm's final performance. For instance, about 24% 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. These firms' average performance is a relatively high 0.917 and is being attributed to the P = 1 design, since that is the design observed at the end of the simulation

[8] Note, in our model firms do not die, but structures may, as over time a particular structure may fail to be adopted by any firm. In this sense, some structures survive, while others do not.

[9] Let $\Pi_0(P1)$ and $\Pi_0(P5)$ be the performance of firms with P = 1 and P = 5 when they cannot change their design. Let $\Pi_S(P1)$ be the performance of P = 1 survivors and α the fraction of final P = 1 firms that are survivors. Then, the degree to which survivors contribute to the dampening effect is given by $\alpha/\textit{dampening} * [(\Pi_S(P1) - \Pi_0(P1)]/[\Pi_0(P5) - \Pi_0(P1)]$.

for these firms.

If firms follow a time-driven incremental search process, such experimentors – firms that achieved their final performance with a design that is different from the current design – are rampant. About 85% of all firms with P = 1 at the end of the simulation reached their sticking points (i.e., final-period performance) with a different design. Overall, these experimentors contribute 86% of the dampening effect (see the third row in the bottom panel of Table 1).[10]

Given that in this case the dampening effect is mainly caused by the wrong-attribution effect, one might argue that a researcher has been "lazy" not to track down for each firm the design that really led to the firm's final performance (or the researcher might simply have been unable to do so). How would the results look if one did not simply attribute final performance to the final design, but if one attributed final performance to the design that actually led the firm to this performance? As reported in column 5 in Table 2, tracking down the design that leads to final performance reveals a larger drift towards the appropriate design. On average, only 9.7% of all firms reach their final sticking point with a P = 1 design, while 27.7% of all firms reach their final sticking point with a P = 5 design. The degree of dampening, however, is still very high. About 76% of the performance difference between the best and the poorest design is washed out. Again the dampening is driven by higher performance of the P = 1 firms. What could explain this result? The same two mechanisms are actually still at play. First, there is still the same set of high-performing survivors, i.e., firms 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 wrong-attribution effect is also still 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 reach a relatively high performance. If that earlier search did not end in a sticking point, the P = 1 design would get credit for finishing off the search process. Thus, the P = 1 design, which leads the firm onto its final sticking point, still receives credit for the performance

[10] There is a third category of firms, besides survivors and experimentors that end up with P = 1: firms that started with P = 1 and that found their sticking point when they had a P = 1 design, yet not within the first 10 periods. The performance of these firms makes up the remaining dampening effect. In general, this effect is relatively small; hence we focus on survivors and experimentors.

generated by the entire sequence of organizational designs that the firm ever had, leading to a seemingly high performance associated with having a P = 1 design. In other words, the dampening effect would not be eradicated 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, we thus see a significant compression of performance differences, which would make a test of Proposition 1 more difficult. Suppose one hypothesizes that firms with P = 5 outperform firms with P = 1. How large a sample would a researcher need to find a statistically significant ($p < 0.05$) difference between the performance of these two types of firms? If firms were endowed with fixed designs that they could not change (the implicit assumption generally made when theorizing about why certain designs outperform other designs), one would only need 21 firms of each type. (This calculation takes into account the differences in the means reported in column 1 of Table 2, as well as the standard deviations of the individual means.) If one observes these firms, however, only after a period during which firms have adjusted and experimented with their designs, one would need 672 firms of each type to find statistically significant support for one's hypothesis. In this sense, it is more than 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 are reported in the last row of Table 2.)

Performance-driven incremental search process. So far we have assumed that firms would change their design every 10 periods, regardless of their performance. In the next simulation, we relax this assumption by using the performance-driven search process, i.e., by making the probability of a firm changing its design an (inversely correlated) function of its current performance. Columns 6 and 7 in Table 2 contain the results of this simulation. In contrast to the time-driven process, we can observe a slightly larger drift towards P = 5 designs, with 24.4% of all firms ending up with P = 5, and a lower degree of dampening. However, there is still a significant performance compression, with about 63% of the performance difference between the best and worst designs being dampened away. In the presence of a performance-driven search process, which allows high-performing firms to retain their designs, survivors are playing a larger role, as one might expect. In this case about 19% of all original P = 1 firms survive, i.e., reach their final performance in the first 10 periods, never find a high performance later, and end the simulation with P = 1 design. (About half of these firms never change their design at all, i.e., have a P = 1 design throughout their lives.) Moreover, these survivors have a high average performance of 0.926. Overall, in this case, survivors contribute 42% of the dampening effect. Thus, even with performance-driven search processes we find performance dampening across organizational designs. In this case, one would need 106 firms of each design to find a significant difference between P = 1 and P = 5 firms.

Time-driven mimicry search process. We now turn our attention to the effect of mimicry on both the distribution of designs that are observed at the end of the coupled search process and the degree of performance dampening. In these simulations, when firms decide to change their design, they do not change their design incrementally. Rather, they imitate the design of high-performing firms.

With mimicry one might expect more drift towards P = 5, since better designs should be copied over time, thereby crowding out poor designs. The results of the time-driven mimicry search process can be found in columns 8 and 9 of Table 2. First, we observe that in aggregate there is hardly any drift towards P = 5 firms. At the end of the simulation, 19.5% of all observed firms had a P = 1 design, while 20.1% of firms had a P = 5 design. Why is there such a large degree of diversity of designs? In part, the apparent diversity of designs is generated by diversity across landscapes. The average number of designs one can observe in any given landscape at the end of each simulation is 2.2. Thus, the mimetic process does indeed lead to a reduction of variety in any given landscape. However, in different landscapes different designs survive, creating 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 find the emergence of one dominant organizational design within France, Italy, and the US – yet this design was different in each country.

With respect to dampening, we continue to see a large performance compression. The performance difference between P = 1 and P = 5 firms is dampened by 86%.[11] The mechanisms behind the dampening effect remain the same as with the incremental search processes 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 two-fold. Some P = 1 firms are high-performing survivors, and others are experimentors, i.e., firms for which their current design has nothing to do with the design that generated the performance for the firm.

With mimicry one might be surprised by the large degree that wrong-attribution plays in contributing to the dampening effect (about 85%). In contrast to the incremental search processes, with mimicry experimentation is not random: the design of high-performing firms are imitated. Thus, for the wrong-attribution mechanism to work, “poor” designs must be imitated by firms with “better” designs. Why does this happen? We can identify three reasons why firms might imitate poor designs of other firms. Firms may imitate the lucky; firms may imitate the fast; or firms may imitate other imitators, i.e., experimentors.

Imagine a P = 1 firm that for some (lucky) reason reached a very high performance. If firms with different, more appropriate design, reached relatively high performance but not as high as the outlier P = 1 firm, these firms might copy the P = 1 design. 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 firm. In a sense, the process of mimicry magnifies the luck of an outlier P = 1 firm, thereby leading to a large dampening effect.

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 (a fact we will exploit later in the simulations with dynamic environments). As a result, in period 10, the long-run superiority of

[11] The dampening effect is not driven by the fact that P = 1 and P = 5 firms survive in different landscapes, i.e., the dampening effect is not an across-landscape phenomenon. Even if we constrain our analysis to those landscapes in which both P = 1 and P = 5 firms survive (this happens in 330 out of 2,000 landscapes), the dampening effect is 95%.

P = 5 designs has not come to the fore, and many P = 1 firms actually 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 towards the “wrong” design. Interestingly, while such “sub-optimal” outcomes due to path dependencies and positive feedback effects have generally been bemoaned (David, 1985; Arthur, 1989), we see that the (population-wide) performance outcome should the “wrong” design win, is not that bad. Wrong designs can win only if they performed fairly well; and other firms may obtain relatively high performance before they adopt the wrong design; hence the harm done by eventually adopting the wrong design is muted in this case.

Performance-driven mimicry search process. If firms follow performance-driven mimicry search processes, we finally see a larger drift away from poor designs toward better designs. Overall, the number of P = 1 firms is halved (less than 10% end with a P = 1 design), while 35% of all firms end with a P = 5 design. The degree of dampening, however, increases even more. Now, 93% of the performance difference is washed away, requiring 2,668 firms of each type to detect a statistically significant performance difference. While the wrong-attribution mechanism still accounts for a large fraction of the dampening effect (48%), survivors now contribute significantly (34%). The number of P = 1 survivors is still relatively small (only 11% of P = 1 firms are survivors), but their performance of 0.950 is extremely high. We thus observe an interesting dynamic: The stronger the selection mechanism that operates on poor designs, the better the survivors have to be in order to survive with poor designs, leading to a greater dampening effect, i.e., a worsening of the correlation between design and performance.

Results in Turbulent, Simple Environments

We now turn our attention to a different environmental setting – turbulent and simple – that calls for a different appropriate organizational design. In column 1 of Table 3, we report the benchmark performances of firms that do not change their designs. In this environment, narrow information exchange, i.e., a low P, is beneficial. Firms with P = 1 and P = 2 have the highest performance, while firms with P = 5 have the lowest performance. (Firms with P = 1 and P = 2 have statistically indistinguishable performance. All other performance differences among the firms are statistically highly significant.) The intuition behind this result is as follows: in the simple environment, no interdependencies between the two departments of the firm exist. As a result, alternatives that are optimal for individual departments are also optimal for the firm. Consequently, upper management should give departmental managers high autonomy in filtering alternatives and only request one proposal. In contrast, if upper management insists on many proposals, upper management can become a bottleneck, not getting to the best proposals sent up by the departments, thereby slowing down the improvement process. This reduction in the speed of improvement is particularly harmful in turbulent environments, in which quick realization of performance improvements is necessary for high performance.

<INSERT TABLE 3 ABOUT HERE>

Incremental search processes. When firms follow time-based incremental search processes, we again see relatively little drift towards appropriate designs. 40% of firms started with the appropriate design (P = 1 or P = 2), and 45% ended with appropriate designs. Likewise, while 20% started with the worst design, P = 5, 15% of all firms still have this design by the end of the simulation. In sharp contrast to the results from the stable environments, however, is the low degree of dampening. In the turbulent environment, we see only 10% performance dampening.

Turbulence basically negates the two mechanisms for dampening that we had identified. Consider, e.g., a firm with P = 5, that by some luck survived the adaptation process. As the environment changes in period 191, i.e., as the relationship between choices and performance is re-set, this luck is discarded. The firm finds itself in a new situation and its prior good performance is of little help. Moreover, the firm’s inappropriate design now does hurt the firm. Hence, we see that the performance of firms that end with a P = 5 design is very similar to those P = 5 firms that never change their design, yielding only a small dampening effect. In sum, with environmental shocks that negate prior advantages, survivor effects vanish.

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

Similar results can be seen when firms follow performance-driven, incremental search processes (columns 4 and 5 in Table 3). We see a slightly larger drift towards high-performing designs and away from P = 5 designs. At the same time, absolutely no performance dampening is present.

Mimicry search processes. With time-driven mimicry search processes, we see both a larger drift towards appropriate designs and the re-emergence of, at least, moderate dampening. 63% of all firms end up with P = 1 and P = 2 designs, but the performance difference between P = 1 and P = 5 firms is dampened away by about 24%. Here, a subtle survivor effect arises again. Recall, we identify the "final" design of a firm as either the firm's design in period 200, or the design the firm would revert to the next time it considers restructuring. A firm with P = 5 in the last period is thus counted as having a final P = 5 design only if its performance is higher than the performance the firm had with its prior design. Thus, some poor performing P = 5 firms, that would revert to some other design in the next cycle of restructuring, are not classified as P = 5 firms, boosting the average performance of the (remaining) P = 5 firms. Indeed, if we simply coded as final design the design firms have in period 200, we would observe more P = 5 firms (4.7% rather than 3.9%), with an overall lower average performance and hence a smaller dampening effect (6%).[12]

[12] This survivor effect also accounts for the small dampening effect seen in the time-driven, incremental adaptation case above (columns 2 and 3 in Table 3). In that case the survivor effect is, however, smaller. With incremental adaptation, a P = 5 firm could only have been a P = 4 firm in its prior organization cycle. The performance

Lastly, in the case of the performance-driven mimicry search processes, we finally see a very strong drift towards appropriate designs (about 80% of all firms end up with P = 1 or P = 2), and a small degree of dampening.

Robustness

We tested the robustness of our results with respect to a number of assumptions of the model: the capability of upper management, the number of firms per starting design, and the reorganization frequency. We also investigated two other environmental conditions (stable and simple; and turbulent and complex). Under all assumptions and conditions, results could be explained by how changes to the simulation altered the strength of the two underlying mechanisms – the survivor effect and the wrong-attribution effect. (See the Appendix for more detail.)

DISCUSSION AND CONCLUSION

Our analysis identifies tough challenges for researchers who seek empirical evidence of the impact that design has on organizational performance. At the same time, it implicitly reveals some good news for the health of populations of organizations. Both the empirical challenges and the implicit good news may be echoed in diverse contexts, beyond organizational design, where coupled search processes are at work.

Empirical challenge and response. Our findings indicate that it can be very difficult to test empirically the appropriateness of particular organizational designs, for two related reasons. First, contrary to Proposition 2, there may be very little drift in a population of firms toward an optimal design. This is particularly likely if:

(a) reorganizations are sparked by the passage of time, not by poor performance, so that high-performers with appropriate designs can gradually wander away from good designs; and

---

difference between P = 4 and P = 5 firms tends to be slight; hence, relatively few P = 5 firms would actually revert to P = 4 firms. With mimicry, however, P = 5 firms might have been P = 1 or P = 2 firms. In order for a P = 5 firm not to revert, it must have had better performance than it had with these much better designs. Clearly, if firms survived this selection criteria (and are only observed and counted when they survive), their performance will be abnormally high, thereby creating a larger dampening effect.

(b) the environment is stable. In a turbulent setting, a firm that wanders away from a good design soon learns that its altered design cannot cope well with new conditions, and it reverts to the good design. In a stable setting, a firm may retain good operational choices, and associated high performance, even as it wanders away from an appropriate design.

Thus, under these conditions, it may not be fruitful to assess the optimality of designs by observing the frequency distribution of designs.

An empirical researcher can ameliorate this problem by tracing back each firm's history to identify the design that first delivered the firm's current performance. One can then examine whether good designs are disproportionately likely to have led a firm to arrive at its current performance. We find, however, that this may not uncover appropriate designs fully. A firm's current performance can result from an entire sequence of prior designs, so the design that finally delivered current performance might not be the design most responsible for that performance. Our results suggest that individual designs may not be the appropriate unit of analysis; rather, the right unit of analysis may be different *sequences* of designs.

The second reason empirical testing may be difficult is that, in many instances, designs expected to perform poorly can appear to perform quite well. Especially in stable and complex settings, much of the performance difference between firms with appropriate and inappropriate designs may be compressed by survivor and wrong-attribution effects. This finding reveals a major challenge for testing conceptual claims about the appropriateness of organizational designs. In this type of work, one usually makes a conceptual argument about why a certain design should perform particularly well (given environmental conditions, etc.). One then assembles a data set and tests the relationship between designs and performance (Proposition 1). A failure to find such a relationship is generally seen as a failure of the conceptual framework. (For such a test, see for instance Doty, Glick and Huber (1993)). Our simulations showed, however, that even when a substantial performance difference among designs exists (i.e., the conceptual framework is correct), coupled search processes can dampen performance differences significantly and throw doubt on the hypothesized design-performance relationship. Similarly, one may conclude that equifinality is present among a range of designs even when it is not; the fact that firms with different designs display similar performance at a point in time does not necessarily imply that these designs are equally effective.

An empirical researcher can combat dampening, we find, by conducting tests in turbulent settings, where designs must rise repeatedly to the challenge of finding good operational choices. In such an environment, a firm cannot hide its poor design by getting lucky once with its operational choices. Environmental turbulence, however, brings new problems for the empirical researcher; performance data may become noisier in turbulent settings, for instance, and patterns are harder to discern.

A researcher may also combat dampening by seeking settings where firms cannot reorganize easily. After all, dampening is driven by the ability of firms to move among designs – that is, to reorganize. This ability comes in two forms: (a) the ability of firms with good designs and high performance to move toward poor designs without performance consequences, and (b) the ability of firms with poor designs and low performance to move toward good designs, leaving only lucky, high performers among those with poor designs. Conditions that limit these abilities may reduce the degree of dampening. For instance, one may suspect that if firms are allowed fewer opportunities to reorganize, dampening may disappear. Our simulations, however, temper this hope. Figure 2 shows the degree of dampening that arises in our simulations in the stable, complex environment after each round of reorganization (i.e., after every ten periods). Dampening indeed worsens with more reorganizations, but it is salient after only a few reorganizations. This is especially true when firms can mimic each other and leap rather than tweak in the space of possible designs. Only severe limits on the number of reorganizations will make dampening far less damaging to empirical research.

In sum, we identify challenges to empirical research on the performance impact of design, and we suggest approaches to tackle these challenges. None of the approaches, however, fully resolves the difficulties we uncover.

<INSERT FIGURE 2 ABOUT HERE>

Design diversity and population health. While the lack of a strong drift towards more effective designs makes empirical testing of appropriateness of designs based on frequency counts difficult, it may actually be a helpful property for populations of firms. As Hannan and Freeman (1989: 7) point 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 under some conditions coupled search processes might achieve both long-term adaptability *and* strong short-term performance. "Poor" designs may be retained and serve as diverse, helpful seeds should the environment change. At the same time, firms with "poor" designs show fairly high performance, either because they arose from "mutations" of better designs or because poorly designed but lucky firms discovered good operational choices. In our model, at least in stable environments, a change from a good design to a poorer one is often performance neutral; performance is not negatively affected, although performance improvement is stunted. Such "wandering with impunity" in organizational design space has a biological counterpart: "neutral mutations" have been proposed to be a key element in the context of biological evolution (Kimura, 1983) and to serve as a springboard for future adaptations should the environment require it (Huynen, Stadler and Fontana, 1996).

A second piece of good news was seen in simulations that allowed firms to imitate each other. In these instances, even if a poor design prevailed in the evolutionary process (i.e., was the sole surviving form or among the few survivors within a particular landscape), the overall population performance tended to be high. Firms within a landscape could get locked in to a suboptimal design, but the conditions under which this happened made it likely that overall performance was relatively high. Otherwise, the poor design could not have won the evolutionary process.

In sum, the same processes that complicate the work of the empirical researcher harbor potential benefits for the population of firms as a whole: organizational variety may be preserved without incurring a performance penalty.

Beyond organizational design. While our simulation focused on organizational design as an influence on performance, the argument that underlies performance dampening applies to a wide range of coupled search processes. The elements required are (a) some set of choices that affects performance not directly and deterministically but by shaping a second set of choices; (b) a tendency among organizations to explore alternative configurations of the first set of choices over time (e.g., in our simulation, to reorganize); and (c) some stickiness in the second set of choices that acts to (at least partially) preserve high performance once it is discovered. (This last element arises in our simulation because upper management refuses to accept changes in operational choices that cause performance to decline.) Beyond these elements, our results show that dampening is especially pronounced when the choices in the second set are intertwined and the mapping from the second set of choices to performance is stable. As noted in the introduction, the required elements seem present in a wide range of organizational phenomena: the creation of alliance networks, the development of cognitive frames among managers, firms' resource allocation processes, and investment in dynamic capabilities, for example. In each setting, we see the potential for the performance effects of the low-frequency choices to be dampened and hard to detect.

Consider, for instance, the long-standing debate in the social network literature concerning which type of network design is beneficial. One camp has argued that actors (e.g., firms or individuals) benefit from being embedded in densely connected networks (Coleman, 1988), while another camp has 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 measure network design at a point in time and relate this design 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 network 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 design. If networks change over time, and performance of actors can be sticky, differences in performance that might be generated by the different networks can be dampened away. For instance, let us assume that structural holes provide firms with a long-lasting performance advantage as they may, e.g., allow a firm to establish a reputation in its industry. If the network structure changes over time because actors in the network create more ties to each other, 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.

Perhaps the most prevalent coupled search process in large corporations involves the search for a management team. At a low frequency, a firm's board of directors searches for a team with a winning combination of traits. At a higher frequency, the team seeks to find and implement an effective array of strategic choices. We would expect these coupled processes to produce data in which it would be hard to discern the performance effect of managerial traits. The empirically observable relationship would be obscured by survivor and wrong-attribution effects. The survivor effect would arise because the firms that persisted with poorly-traited management teams would be those that happened by luck upon good strategic choices despite their management. The wrong-attribution effect would arise because firms that secure good strategic choices under strong management teams may persist in those good choices even when management changes for the worse.

In the contexts of management selection, network formation, and other coupled search processes, we expect not only to see performance dampening but also to find that the empirical responses outlined above are effective. Researcher who aim to discern relationships despite dampening would do well to examine whole histories of low-frequency choices, not just final sets of such choices; to focus on turbulent settings; and to look for contexts in which low-frequency choices are relatively immobile. Moreover, we would expect such settings to display diversity in low-frequency choices without large, observable performance penalties for seemingly “wrong” choices. Coupled search processes are able to 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 seeks 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, hierarchical architecture. Why aren’t the search processes unitary, with—for instance—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.

TABLE 1: SUMMARY OF DALTON *ET AL.* (1980) REVIEW OF EMPIRICAL EVIDENCE

Structural dimension that may affect performance	Number of studies	Conclusion
Span of control	4	"Conflicting reports and a paucity of empirical work in the area make it difficult to summarize this research…. It is probably safe to say that there is *no* evidence concerning the relationship of span of control and performance of blue-collar, nonmanagerial, or non-professional employees."
Flat or tall hierarchy	4	"Summarization of the vertical span relationship is problematic. It is difficult to generalize across findings with professionals, laboratory studies, and white collar employees, with both positive and negative associations reported."
Administrative intensity	7	"Once again, a definitive summarization is not possible. There are those who report positive and negative associations. The relationship between administrative intensity and performance remains undetermined."
Specialization and complexity	6	"Although the preponderance of evidence suggests a positive relationship between specialization / complexity and performance, the lack of hard performance criteria, coupled with reports of no association, leads to the conclusion that the association between specialization and performance has not been clearly demonstrated."
Formalization and standardization	8	"We can conclude that an association between levels of formalization and performance has not been convincingly demonstrated."
Centralization	16	"[T]he limited evidence tends to support a negative relationship between centralization and performance for managers and professionals in studies using hard performance criteria. Otherwise, little is known of the association between centralization and performance."

TABLE 2: RESULTS FROM STABLE, COMPLEX ENVIRONMENTS

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Final design	Bench-mark	Time-driven incremental		Time-driven incremental*		Perf.-driven incremental		Time-driven mimicry		Perf.-driven mimicry	
P = 1	0.843	0.906	19.4%	0.900	9.7%	0.891	15.8%	0.909	19.5%	0.913	9.7%
P = 2	0.868	0.906	19.9%	0.902	16.9%	0.900	17.7%	0.908	20.8%	0.911	14.3%
P = 3	0.884	0.910	20.0%	0.908	21.1%	0.907	19.7%	0.910	19.4%	0.911	18.2%
P = 4	0.903	0.915	20.4%	0.915	24.6%	0.915	22.4%	0.912	20.2%	0.915	23.3%
P = 5	0.920	0.917	20.3%	0.919	27.7%	0.920	24.4%	0.920	20.1%	0.918	34.5%
Dampening		85%		76%		63%		86%		93%	
Percent of dampening due to:											
Survivors		8%				42%		8%		34%	
Wrong-attribution		86%				54%		85%		48%	
Number of obs. required to find significant differences	21	672		262		106		718		2668	

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

TABLE 3: RESULTS FROM TURBULENT, SIMPLE ENVIRONMENTS

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Final design	Bench-mark	Time-driven incremental		Perf.-driven incremental		Time-driven mimicry		Perf.-driven mimicry	
P = 1	1.000	1.000	21.9%	1.000	20.9%	1.000	32.4%	1.000	42.6%
P = 2	0.999	0.999	23.3%	0.999	26.4%	0.999	30.9%	0.999	37.1%
P = 3	0.992	0.993	21.7%	0.991	21.8%	0.993	22.9%	0.991	13.5%
P = 4	0.974	0.979	18.5%	0.973	17.6%	0.979	9.9%	0.973	4.7%
P = 5	0.952	0.956	14.6%	0.951	13.2%	0.963	3.9%	0.953	2.1%
Dampening		10%		0%		24%		4%	
Number of designs						3.58		3.67	
Number of obs. required to find significant differences	11	11		10		13		10	

FIGURE 1: PERFORMANCE BY ORGANIZATIONAL DESIGN IN STABLE, COMPLEX ENVIRONMENTS

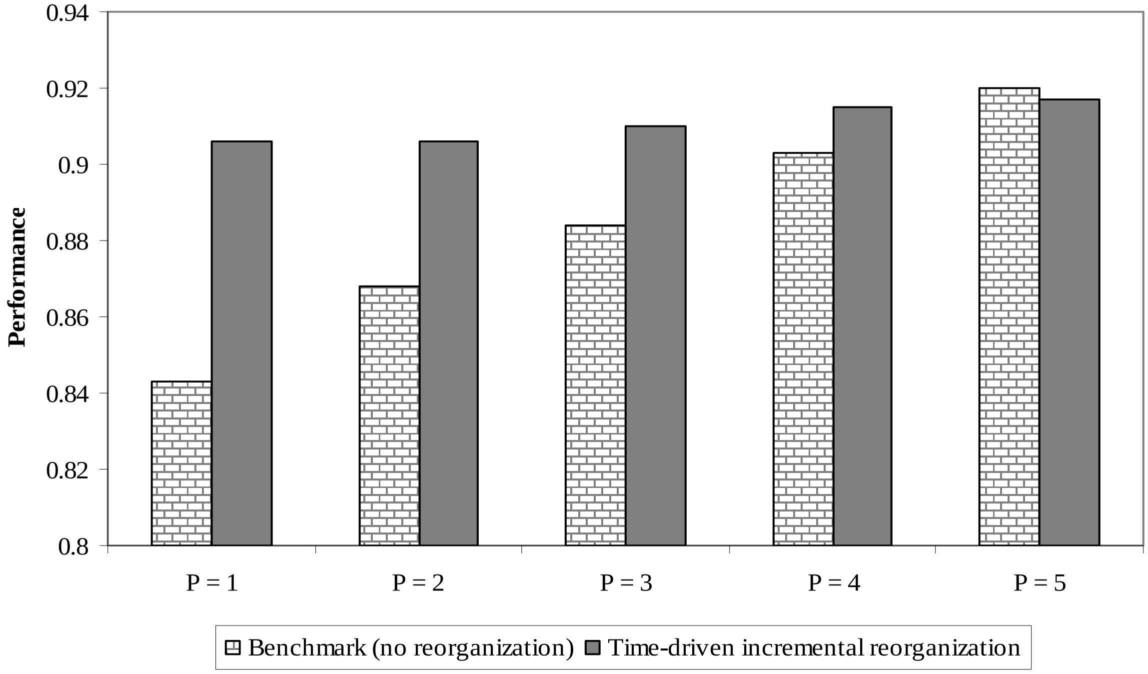

FIGURE 2: DAMPENING AND THE NUMBER OF REORGANIZATIONS

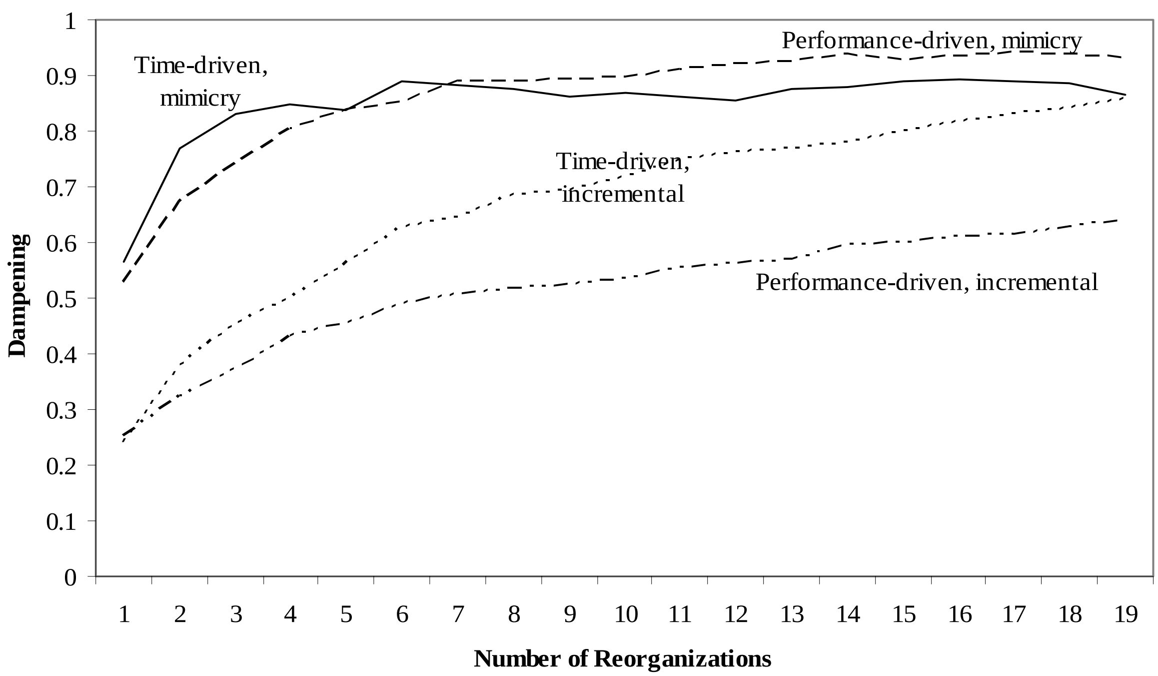

Appendix: Robustness results

We investigated the degree of dampening and the drift towards appropriate structures in a broad range of circumstances by modifying assumptions and parameters within our model. In particular, we changed the capability of upper management, the number of firms per starting design, and the reorganization frequency. Lastly, we analyzed two other environmental conditions (stable and simple; and turbulent and complex). We discuss only those results that showed interesting differences from the main results reported in the body of the paper.

Robustness results for stable, complex environments. An increase in the capability of upper management is modeled by increasing the number of composite alternatives that upper management can evaluate. Rather than assuming that upper management can evaluate only one composite alternative in each period, we assume that all possible combinations can be evaluated (up to 25). By increasing upper management capabilities, upper management is not a source of "luck" anymore, since all possible composite alternatives are always evaluated. Moreover, given that upper management is now "supercharged," we remove the possibility that upper management is a bottleneck. As a result, P = 1 firms lose their short-run performance advantage over P = 5, i.e., P = 1 firms are not faster anymore in improving their performance. Consequently, when firms are mimicking high-performing firms, fewer P = 1 firms should present high-performance targets for copying attempts. Indeed, with more capable upper management we observe by the end of the simulation 8% fewer P = 1 firms and 8% more P = 5 firms than with less capable upper management (time-driven search process). The observed dampening increases from 86% to 97%.

When we reduce the number of firms per starting design to 1 (rather than 5), mimicry processes have fewer firms to operate on. As a result, the diversity of designs decreases in almost any landscape down to only one form. Which design wins, appears still to be random, as overall we do not observe any significant drift to P = 5 design. The dampening effect, however, is weakened to 65%. In this case, when the P = 1 design "wins," firms with other designs did not have much time to reach high performance. Given that the P = 1 design does not allow firms to significantly improve their performance afterwards, if a landscape ends up with firms having only P = 1 designs, their overall performance tends to be low. Thus, in these cases lock-in into the wrong design does have a detrimental effect on population-wide performance.

Lastly, we change the reorganization frequency. Rather than considering reorganization every 10 periods, firms consider reorganizing every 5 or every 25 periods. Two qualitative effects arise from different reorganization frequencies. First, the shorter the reorganization frequency, the larger the early short-run performance advantage of P = 1 firms over P = 5 firms. Second, the shorter the reorganization frequency, the larger the total number of reorganizations that occur over 200 periods. The greater short-run advantage of P = 1 firms leads to a larger fraction of eventual P = 1 firms when firms are mimicking, especially when firms are using performance-driven search processes (since in that case high-performing P = 1 firms can retain their design). In this case, 26% more P = 1 firms can be seen at the end of the simulation, as compared to the case when reorganization happens every 10 periods. The larger number of total reorganizations leads to an eventual decrease of number of different designs that can be observed. For instance, with the time-driven, mimicry search process, only 1.48 different designs survive till the end of the simulation (as compared to 2.24 when reorganization occurs every 10 periods). Dampening effects for all search processes are fairly similar to the case of 10-period reorganization.

When the reorganization frequency is changed to 25 periods, the largest changes can be observed for the incremental search processes. Here dampening decreases for the time-driven rule from 85% to 64% and for the performance-driven rule from 63% to 44%. With fewer reorganization opportunities, and only incremental changes to P, fewer high-P firms end up tweaking their way all the way to P = 1 designs. As a result, observed P = 1 firms contain fewer experimentors, and the wrong-attribution effect is smaller.

Robustness results for turbulent, simple environments. In the simulations reported above, we assumed that the frequencies of environmental turbulence and reorganization were equal (every 10 periods). To check for robustness, we also studied the effects of reorganizations that happen more frequently (every 5 periods) than environmental changes. Again, shorter reorganization periods pronounce the short-run performance advantage of P = 1 firms, which in this case is the appropriate design. As a consequence, with more frequent reorganizations, we generally find a larger drift toward the efficient design. This is in particular the case when firms use the performance-based, mimicry search process. In this case, hardly any P = 5 firms remain (0.1%) and almost all firms find the most appropriate design, P = 1 (81.3%). The few P = 5 firms that do remain, however, have high performance, so that the performance dampening is considerable (72%). Thus, even in turbulent environments, we can observe considerable dampening. In this case, however, given the very few firms with “wrong” designs, one might be naturally more cautious to overinterpret the high performance of firms with seemingly inappropriate designs.

Results for stable, simple or turbulent, complex environments. In stable, simple environments, all firms, regardless of their value of P reach the global peak relatively quickly. As a result, in the benchmark scenario with no organizational change, there are no performance differences among the firms. Hence, no optimal value of P exists, making the exercise of assessing whether firms can find an appropriate design not very sensible. (Consistent with the results in the simple, turbulent environment, if firms are allowed to mimic, we do find a drift towards P = 1 and P = 2 firms, reflecting their speed advantage in the first periods. Still, all firms end up on the global peak, with performance 1.0.)

In turbulent, complex environments, we see the expected tradeoff between speed and breadth of search that different values of P provide. Firms with low P are fast but get stuck quickly; firms with high P are slow but are able to explore a little more before the landscape changes again. For a reorganization frequency of 10 periods, we find a curvilinear benchmark result. Firms with P = 2 and P = 3 significantly outperform firms with P = 1 and P = 5. The internal optimum of P, given these environmental conditions, does not create a clear gradient for improvement along the P-dimension. Consequently, we find results that are intermediate to the results of the other two environments. Generally, we find more drift towards efficient designs than in the complex, stable environment (depending on the reorganization rules, we find 18 - 108% more firms with P = 2 or P = 3 than with P = 1 or P = 5), yet less so than in the simple, turbulent environment. Likewise, the degree of dampening is at a level intermediate to that found in the other two environments.

REFERENCES

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

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

Barnard, Chester I.
1938 The functions of the executive. Cambridge, MA: Harvard University Press.

Barnett, William P., and Olav Sorenson
2002 "The Red Queen in organizational creation and development." Industrial and Corporate Change, 11: 289-325.

Bower, Joseph L.
1970 Managing the resource allocation process. Boston: Harvard Business School Press.

Brown, Shona L., and Kathleen 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.

Burns, Tom, and G. M. Stalker
1961 The management of innovation. London: Tavistock.

Burt, Ronald S.
1992 Structural holes. Cambridge, MA: Harvard University Press.

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

Carroll, Glenn R., and J. Richard 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, Massimo G., and Marco Delmastro
1999 "Some stylized facts on organization and its evolution." Journal of Economic Behavior and Organization, 40: 255-274.

Cyert, Richard M. , and James G. March
1963 A behavioral theory of the firm. Englewood Cliffs, NJ: Prentice Hall.

Dalton, Dan R., William D. Todor, Michael J. Spendolini, Gordon J. Fielding, and Lyman W. Porter
1980 "Organizational structure and performance: A critical review." Academy of Management Review, 5: 49-64.

David, Paul A.
1985 “Clio and the economics of QWERTY.” American Economic Review, 75: 332-337.

Davis, Jason P., Kathleen M. Eisenhardt, and Christopher B. Bingham
2007 “Developing theory through simulation methods.” Academy of Management Review, 32: 480-499.

DiMaggio, Paul J., and Walter W. Powell
1983 “The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields.” American Sociological Review, 48: 147-160.

Djelic, Marie-Laure, and Antti 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 Publications.

Doty, D. Harold, William H. Glick, and George P. Huber
1993 “Fit, equifinality, and organizational effectiveness: A test of two configurational theories.” Academy of Management Journal, 36: 1196-1250.

Ethiraj, S. K., and D. A. Levinthal
2004 “Bounded rationality and the search for organizational architecture: An evolutionary perspective on the emergence of organizational forms and their evolvability.” Administrative Science Quarterly, 49: 404-437.

Fayol, H.
1916 [1949] General and industrial management. London: Pitman.

Galbraith, Jay R.
1973 Designing complex organizations. Reading, MA: Addison-Wesley.

Galunic, D. Charles, and Kathleen M. Eisenhardt
1994 “Renewing the strategy-structure-performance paradigm.” In B. M. Staw, and L. L. Cummings (eds.), Research in Organizational Behavior: 215-255. Greenwich, CT: JAI Press.

Gavetti, Giovanni, and Daniel Levinthal
2000 “Looking forward and looking backward: Cognitive and experiential search.” Administrative Science Quarterly, 45: 113-137.

Glaeser, Edward L., Rafael La Porta, Florencio Lopez-de-Silanes, and Andrei Shleifer
2004 “Do institutions cause growth?” Journal of Economic Growth, 9: 271-303.

Gresov, Christopher, and Robert Drazin
1997 “Equifinality: Functional equivalence in organizational design.” Academy of Management Journal, 22: 403-428.

Hannan, Michael T., and John Freeman
1977 “The population ecology of organization.” 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 inter-organizational 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.

Huynen, M.A., P.F. Stadler, and W. Fontana
1996 "Smoothness within ruggedness: The role of neutrality in adaptation." Proceedings of the National Academy of Sciences USA, 93: 397-401.

Jensen, Michael C., and William H. Meckling
1976 "Theory of the firm: Managerial behavior, agency costs and ownership structure." Journal of Financial Economics, 3: 305-360.

Kauffman, Stuart A.
1993 The origins of order: Self-organization and selection in evolution. New York: Oxford University Press.

Kimura, M
1983 The neutral theory of molecular evolution. Cambridge, UK: Cambridge University Press.

Lawrence, Paul R., and Jay W. Lorsch
1967 Organization and environment. Boston: Harvard Business School Press.

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

March, J. G., and J. P. Olsen
1989 Rediscovering institutions: The organizational basis of politics. New York: Free Press.

March, James G., and Herbert A. Simon
1958 Organizations. New York: John Wiley & Sons.

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

Mintzberg, Henry
1979 The structuring of organizations. Englewood Cliffs, NJ: Prentice Hall.

Mohr, L.
1971 "Organization technology and organizational structure." Administrative Science Quarterly, 16: 444-459.

Nadler, David A., and Michael L. Tushman
1997 Competing by design. New York: Oxford University Press.

North, Douglass C., and Barry R. Weingast
1989 "Constitutions and commitment: The evolution of institutions governing public choice in seventeenth-century England." Journal of Economic History: 803-832.

Pennings, Johannes M.
1975 “The relevance of the structural-contingency model for organizational effectiveness.” Administrative Science Quarterly, 20: 393-410.

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, Michael E.
1996 “What is strategy?” Harvard Business Review, 74: 61-78.

Rivkin, Jan W.
2000 “Imitation of complex strategies.” Management Science, 46: 824-844.

Rivkin, Jan W., and Nicolaj Siggelkow
2002 “Organizational sticking points on NK-landscapes.” Complexity, 7: 31-43.

Schoonhoven, Claudia Bird
1981 “Problems with contingency theory: Testing assumptions hidden within the language of contingency “theory”.” Administrative Science Quarterly, 26: 349-377.

Siggelkow, Nicolaj
2002 “Evolution toward fit.” Administrative Science Quarterly, 47: 125-159.

Siggelkow, Nicolaj, and Daniel 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, Nicolaj, and Jan W. Rivkin
2005 “Speed and search: Designing organizations for turbulence and complexity.” Organization Science, 16: 101-122.

Sorenson, Olav
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, David, and Michael W. Macy
2001 “In search of excellence: Fads, success stories, and adaptive emulation.” American Journal of Sociology, 107: 147-182.

Taylor, F. W.
1911 The principles of scientific management. New York: Harper.

Teece, David J., Gary Pisano, and Amy Shuen
1997 “Dynamic capabilities and strategic management.” Strategic Management Journal, 18: 509-533.

Woodward, J.
1965 Industrial organization: Theory and practice. London: Oxford University Press.