

Balancing Search and Stability: Interdependencies Among Elements of Organizational Design

Jan W. Rivkin • Nicolaj Siggelkow

239 Morgan Hall, Harvard Business School, Boston, Massachusetts 02163

The Wharton School, 2017 Steinberg Hall—Dietrich Hall, Philadelphia, Pennsylvania 19104

jrivkin@hbs.edu • siggelkow@wharton.upenn.edu

We examine how and why elements of organizational design depend on one another. An agent-based simulation allows us to model three design elements and two contextual variables that have rarely been analyzed jointly: a vertical hierarchy that reviews proposals from subordinates, an incentive system that rewards subordinates for departmental or firm-wide performance, the decomposition of an organization's many decisions into departments, the underlying pattern of interactions among decisions, and limits on the ability of managers to process information. Interdependencies arise among these features because of a basic, general tension. To be successful, an organization must broadly search for good sets of decisions, but it must also stabilize around good decisions once discovered. An effective organization balances search and stability. We identify sets of design elements that encourage broad search and others that promote stability. The adoption of elements that encourage broad search typically raises the marginal benefit of other elements that provide offsetting stability. Hence, the need to balance search and stability generates interdependencies among the design elements. We pay special attention to interdependencies that involve the vertical hierarchy. Our findings confirm many aspects of conventional wisdom about vertical hierarchies, but challenge or put boundary conditions on others. We place limits, for instance, on the received wisdom that firm-wide incentives and capable subordinates make top-level oversight less valuable. We also identify circumstances in which vertical hierarchies can lead to inferior long-term performance.

(Organizational Search; Stability; Interactions; Organizational Design; Organizational Structure; Vertical Hierarchy; Agent-Based Simulation)

1. Introduction

In the diverse literature on organizational design, at least one proposition has gained widespread acceptance: the many formal and informal structures, systems, and processes that make up an organization's design affect one another (e.g., Khandwalla 1973, Mintzberg 1979). Organizations are typically seen as "complex entities . . . composed of tightly interdependent and mutually supportive elements" (Miller and Friesen 1984, p. 1) and as "highly integrated system[s]

whose performance is determined by the degree of alignment among the major elements" (Nadler and Tushman 1997, p. 23). The marginal costs and benefits associated with any design element depend on the configuration of others. For instance, the efficacy of decentralized decision making may hinge on the incentives, information, and training given to middle management. Management debacles are often interpreted as failures to appreciate the systemic nature of organizational design.

Prior studies pinpoint specific interdependencies among elements of design, but they do not explain *why interdependencies arise in general*. In this paper, we use an agent-based model of organizational decision making to identify a general, underlying tension that gives rise to interdependencies. Our results show that successful firms balance two opposing needs. On one hand, to be effective, a firm must search broadly for good combinations of decisions; it must not lock in prematurely on the first decent set of choices it discovers. On the other hand, a successful firm must halt its search efforts and stabilize its decisions once it finds an outstanding set of choices. We identify specific elements of organizational design that drive a firm toward broad search and others that encourage stability. The need to balance search and stability creates interdependencies among the elements. Often, a firm that adopts an element that pushes it toward broad search benefits from a second element that pulls it toward stability. Prior formal models of organizational search tend to overlook these interdependencies because they often grant “stability for free”; that is, they assume that firms that discover good decisions through search can lock in on those decisions forever. Contrary to this assumption, we illustrate that organizational elements that enable discovery may undermine lock-in.

Our model allows us to examine the relationships among three prominent elements of formal organizational design: a *vertical hierarchy* embodied in a CEO, subordinate managers, and a flow of information among them; an *incentive system* that influences whether managers act for the good of the overall firm or pursue the parochial interests of their departments; and the *decomposition* of a firm’s many decisions into discrete departments. We also pay careful attention to two contextual variables: The *underlying pattern of interaction* among a firm’s decisions, and the *limits on the ability of managers* to process information and consider alternatives. These five features surely do not form an exhaustive list of the design elements and contextual variables that organizational designers have explored. However, as we discuss in §2, they do cover the important classes of considerations in the literature on formal organizational design. Moreover,

this limited list of considerations is more than sufficient to illustrate the need for balance between search and stability.

We model the five features using an agent-based simulation derived from research on complex adaptive systems (see §3). This approach enables us to look simultaneously at all five features, distinguishing our work from prior models that have examined only one or two at a time. Our effort joins a growing set of agent-based simulations of human organizations (e.g., Carley and Lin 1997, Carley and Svoboda 1996, Levinthal 1997, Anderson et al. 1999, Axelrod and Cohen 1999, Chang and Harrington 2000). In our model, firms with different organizational designs face a long series of multidimensional decision problems. Decisions within each problem interact with one another in a manner controlled by the modeler. For each decision problem, each firm attempts to find a good solution; that is, an effective set of choices. The management team of each firm consists of a simple hierarchy: a CEO and two subordinate managers. Each subordinate manager has purview over a subset of the organization’s decisions, a “department.” Starting from a particular configuration of choices, each subordinate considers altering the decisions under his command, evaluates the alternatives in light of an incentive system, and makes recommendations to the CEO. The CEO reviews the proposals and accepts the pair of proposals—one from each subordinate—that will best serve the firm. Alternatively, she may overrule her subordinates and maintain the status quo for either or both departments.

Modeled firms differ in their designs: how active a role the CEO takes and how much information she receives from subordinates; whether subordinates are rewarded for departmental success or the performance of the firm as a whole; and how firms decompose decisions into departments. Organizations also differ in the cognitive abilities of their CEOs and managers. By comparing the performance levels of firms with different designs across a large number of decision problems, we can isolate how the distinct design elements depend on one another.

To structure our analysis, we focus on the interdependencies between the vertical hierarchy and the other organizational elements. The results of our

modeling effort (see §4) show that a hierarchy that actively reviews subordinates' proposals is often helpful, but circumstances exist under which thorough decentralization produces superior performance. The findings also pinpoint design elements and contextual variables that amplify or dampen the value of an active hierarchy. The results confirm many aspects of the conventional wisdom about hierarchies produced by prior observers of organizations, but challenge or show boundary conditions for other aspects. For instance, consistent with received wisdom, we find that an active vertical hierarchy tends to be more valuable when interactions among decisions are pervasive. However, this benefit arises only if the information flow in the hierarchy is rich enough. Without rich information flow, active hierarchies can lead firms to lock in on suboptimal solutions prematurely, leading to worse performance than a firm with a purely passive hierarchy would achieve. Similarly, we bound conventional wisdom by showing that firm-wide incentives and capable managers are complements to, not substitutes for, an active hierarchy when interactions among decisions are sufficiently pervasive.

To interpret our results (see §5), we use a landscape conceptualization that has become popular in certain formal models of organizational search (Kauffman 1995, Levinthal 1997, Levinthal and Warglien 1999, McKelvey 1999, Gavetti and Levinthal 2000, Ghemawat and Levinthal 2000, Rivkin 2000, 2001, Kauffman et al. 2000, Fleming and Sorenson 2001). A mapping from firm decisions to payoffs creates a landscape in the space of decisions. Firms can be conceived of as trying to attain and sustain a high spot on such a landscape—a combination of decisions that, together, yield a high payoff. Organizational design, we argue, affects firm performance by altering firms' search behavior on the landscapes they face. A firm typically gravitates on its landscape toward a "sticking point"—a configuration of choices from which it will not change. Organizational design affects long-term performance by two primary channels. First, it alters the nature of a firm's sticking points—the number of such points and the payoffs associated with them. Second, it influences the likelihood that a firm will actually reach such a stable

configuration of choices. Organizations with the most effective designs, ones that balance search and stability, find good points and stick to those points.

2. Organizational Design Elements and Their Interdependencies

A rich heritage of qualitative studies identifies a host of design elements and contextual features that a formal model of organizational design might encompass. In this section, we first explain why we focus our modeling effort on five specific considerations. We then turn to conventional wisdom, drawn from the qualitative studies, about how these five interact. This wisdom, focused on interdependencies between an active vertical hierarchy and each of the other considerations, provides a backdrop for the results in §4. Finally, we discuss prior formal models of organizational design. In contrast to the qualitative studies, only a few prior models have emphasized interdependencies among design elements.

Common Considerations in Qualitative Studies.

The qualitative literature on organizational design is extensive and diverse, encompassing grounded theoretical work, field studies, and numerous syntheses (e.g., Mintzberg 1979, Gibson et al. 2000). Nonetheless, the literature is unified in what it perceives as the central challenge of organizational design: to divide the tasks of a firm into manageable, specialized jobs, yet coordinate the tasks so that the firm reaps the benefits of harmonious action. Implicit in this challenge are two important assumptions. First, coordination is valuable because the tasks of a firm interact with one another; that is, a decision made concerning one task affects the efficacy of performing another task one way or another. Without such interactions, coordination would be unnecessary. Accordingly, the literature on organizational design (Thompson 1967, Galbraith 1973, Mintzberg 1979 (especially Chapter 7), Nadler and Tushman 1997) has repeatedly returned to the *underlying pattern of interaction* among a firm's tasks, the first consideration we include in our model. The second assumption is that the firm's situation creates a demand for information processing that exceeds the deliberative capacity of any individual manager

(Simon 1957). If this were not so, a single über-manager could evaluate all of a firm's alternatives and dictate the best, coordinated course of action. Hence, our model incorporates, as a second critical feature, *limits on managerial ability*.

How can an organization in which the complexity of decision problems exceeds the cognitive capacity of any single decision maker achieve coordinated action? Below, we focus on three design elements that have played a particularly prominent role in the rich discussions found in the qualitative literature: vertical hierarchies, incentive systems, and decision decomposition.¹ We emphasize these particular elements both because they are ubiquitous in real organizations and because they exemplify the three major classes of organizational elements identified by Nadler and Tushman (1997, p. 67): structural links (formal relationships among decision makers separated by structural boundaries), systems and processes (formal guides to decision making), and grouping (the aggregation of tasks into work units).

A *vertical hierarchy* is perhaps the most common mechanism employed to coordinate the decisions of separate decision makers. A CEO, for instance, may sit above a set of department heads, review the proposals of the departments, and try to integrate them in a way that achieves coordination. Other formal connections across groups exist, e.g., lateral linkages such as liaison positions (Galbraith 1973), but for the sake of parsimony, and because it exists in virtually every organization, we focus our modeling effort on the vertical hierarchy.

Coordination can also be enhanced by systems and processes that span group borders. Designed well, such systems and processes can make managers

aware of and responsive to what happens beyond their own domains. Nadler and Tushman (1997) identify a rich variety of coordinating systems and processes: strategic planning efforts, resource allocation programs, information management systems, and so forth. Among these, we choose to model the *incentive system*, the system that arguably has received the greatest attention. In particular, we explore a system in which managers may be rewarded on the basis of overall firm performance rather than on the performance of their individual departments.

Besides vertical hierarchies and incentive systems, treatises on organizational design consider the way tasks and decisions are grouped together to be "a fundamental means to coordinate work in the organization" (Mintzberg 1979, p. 106). Theoretical analyses emphasize how interactions among tasks influence the way in which decisions should be grouped. Thompson (1967), for instance, argues that decisions should be grouped so that the most intensive interactions are internalized, while Simon (1973) stresses that interactions across decision makers should be minimized. Consequently, groups should be formed so that, as nearly as possible, the firm is decomposed (Simon 1962) into independent entities. Given the prominence of grouping in the qualitative literature on organizational design, our model includes, as its fifth feature, the notion of *decomposition*. That is, modeled firms are able to try to achieve coordination by assigning decision rights in a way that places related decisions under a single manager.

Conventional Wisdom About Interdependencies.

The qualitative literature not only emphasizes that design elements in general have profound effects on each other, but it also identifies some patterns in these interdependencies. Many of these patterns address a critical question about vertical hierarchies: Under what circumstances and in which combination with other design elements should a firm employ senior managers who actively review subordinates' proposals and retain the right to veto changes? Conversely, when should a firm delegate the right to make decisions to subordinate managers? We summarize the received wisdom on these questions as follows:

- *Active vs. passive vertical hierarchy and degree of interaction.* In general, qualitative studies have argued

¹ We focus on formal design elements rather than informal elements such as casual communication systems (Mintzberg 1979) or corporate culture (Camerer and Vepsäläinen 1988). We do so not because informal elements are unimportant or lack interesting interdependencies, but rather for other reasons. The formal elements alone are more than sufficient to fuel a complex analysis as the rest of the paper shows. Moreover, the formal elements can be modeled more precisely than the informal. Precise modeling of each individual element is especially important if one wants, as we do, to examine the interdependencies among them. That said, we consider the analysis of informal elements of organizational design to be an important topic for future research.

that the value of active vertical oversight grows as the degree of interaction among decisions increases. Khandwalla (1973, p. 521), for example, notes that "the greater the interface between functionally organized departments, the greater is the need for coordinative mechanisms such as... a common boss." However, the value of active vertical oversight is crucially limited by the specter that senior managers may become overloaded by the information flowing up to them, causing decision making to grind to a halt. As Child (1984, p. 148) points out: "If top executives are overloaded then the effective control they can exercise will be diminished and they will tend to sit on decisions which may require speedy attention." Decentralized decision making allows a firm to respond more quickly to (changed) local conditions (Mintzberg 1979) and is therefore more appropriate in volatile environments (Mintzberg 1981).

- *Active vertical hierarchy and managerial ability.* The presence of an active hierarchy and the ability of subordinate managers are usually seen as substitutes. Child (1984, p. 71) notes: "The greater the competence of subordinates, the less closely they need to be supervised and the less often does their work require review. Therefore as the competence of subordinates rises... so it becomes feasible to widen spans of control and to reduce levels of management."

- *Active vertical hierarchy and firm-wide incentives.* An incentive system that rewards overall firm goals, rather than departmental goals, has also been considered a partial substitute for an active, coordinating hierarchy. For instance, Galbraith (1973, p. 14) points out: "Goal setting helps coordinate interdependent subtasks and still allows discretion at the local subtask level," while Child (1984, p. 149) observes that "attention to developing a strong identification with top management objectives... permits delegated decisions to be made."

- *Active vertical hierarchy and decision decomposition.* If decisions can be decomposed—parsed out to departments—in such a way that few cross-departmental interactions remain, the value of an active vertical hierarchy declines. "[G]roups that are only minimally interdependent have relatively little need for coordination" and therefore do not require active oversight by a hierarchy (Nadler and Tushman

1997, p. 92). Moreover, decisions that interact should be grouped together as much as possible and assigned to a single decision maker, regardless of the absence or presence of a vertical hierarchy (Thompson 1967, Simon 1973).

Below, we revisit each of these pieces of received wisdom in light of our simulation findings.

Prior Formal Modeling Efforts. Formal modeling of organizational design has burgeoned in the past decade as economists have sought to pry open the "black box" of the firm, yet few studies have considered multiple organizational elements. As a result, interdependencies among elements have received little modeling attention. Insofar as interdependencies have been noted (often as a by-product rather than the focus of analysis), results involving vertical hierarchies confirm the conventional wisdom outlined above. For instance, prior models have noted that firm-wide incentives and an active hierarchy may serve as substitutes. "Firms with... compensation schemes that reward company-wide performance... are more likely to have highly decentralized organization structures" (Harris and Raviv 2002, p. 864). Further, in line with conventional wisdom, Aghion and Tirole (1997) observe that the value of an active hierarchy increases when decisions are not fully decomposed. They show that a CEO is more likely to intervene when a division manager's preferred decisions are likely to be suboptimal for the firm. This situation tends to arise "when there are substantial externalities on other divisions, on future managers of the division, or on the firm as whole" (p. 14).

3. A Model of Organizational Design and Search

Our goal is to test and extend conventional wisdom about interdependencies among organizational design elements by probing the roots of such connections. To do so, we develop a simulation model in which the modeler dictates the underlying pattern of interaction among decisions; a computer generates a set of particular decision problems that follow that


pattern; and large numbers of modeled firms with different organizational characteristics tackle these decision problems. In this section, we describe each of these steps in detail.

Though a simulation model does not yield “exact,” closed-form solutions as an algebraic approach might, the number and nature of the features we feel are important in a model of interdependencies (varying degrees of interactions among decisions, limited ability of decisions makers, a vertical hierarchy with information flows, changeable incentive systems, and different types of decision decompositions) make an algebraic approach infeasible. Take, for instance, interactions among decisions. In our model, as in reality, pairs of decisions may interact as complements or as substitutes (Porter and Siggelkow 2002). As a result, the closed-form analysis of supermodular functions, which has been used to study complementary interactions (e.g., Milgrom and Roberts 1990, 1995, Holmström and Milgrom 1994), cannot be readily applied. Moreover, given the limited cognitive abilities of our agents, the long-term outcomes of the model can be fully understood only by taking the search process, i.e., the dynamics of the model, into account. Characterizing such paths is straightforward with simulations and extremely difficult with analytical models, which tend to focus on equilibrium outcomes and tend to ignore how, indeed whether, such equilibria can be reached.

3.1. Setting the Pattern of Interaction

The management team of each modeled firm must make N binary decisions about how to configure its activities. N reflects the fact that a real firm must make numerous decisions. It must choose, for instance, whether to have its own sales force or to sell through third parties, whether to field a broad product line or a narrow one, whether to pursue basic R&D or not, and so forth. An N -digit string of zeroes and ones summarizes all the decisions a firm makes that affect its performance. We represent this “choice configuration” as $\mathbf{d} = d_1 d_2 \dots d_N$ with each d_i either 0 or 1. Two firms that arrive at the same configuration of choices are presumed to achieve the same level of performance even if different organizational structures guided them to this common set of choices.

Figure 1 Examples of Influence Matrices ($N = 6$)



Put differently, organizational arrangements have no direct costs or benefits. They influence performance only through the operational choices they evoke.²

The efficacy of each decision is affected not only by the choice (0 or 1) made concerning that decision, but also by the choices regarding other decisions. In the model, each decision i makes a contribution C_i to overall firm performance. C_i depends not only on d_i , but also on how other decisions are resolved: $C_i = C_i(d_i; \text{other } d_j\text{'s})$. An $N \times N$ “influence matrix,” $\mathbf{I}$, records the relationships among decisions. Figure 1 gives some examples of influence matrices for $N = 6$. The (i, j) th entry of $\mathbf{I}$ is marked by an “x” if column decision j influences the contribution of row decision i and is blank otherwise.³

In the simulations reported below, we set influence matrices in two ways. In some cases, we fully

² The model also does not incorporate interplay among firms. That is, a firm’s payoff from a configuration is independent of other firms’ configurations.

³ Our influence matrix is closely related to the design and task structure matrices pioneered by Steward (1981), Eppinger et al. (1994), and Baldwin and Clark (2000) in the context of product design.

specify the matrix by hand (Ghemawat and Levinthal 2000). We might, for instance, dictate a block-diagonal matrix of influences as shown in Figure 1C. With such a pattern of influence, one can compare, say, the performance of a firm that allocates decisions 1–3 to one subordinate and decisions 4–6 to the other to the performance of a firm that divides responsibility in some other way. In other simulations, we simply specify K , the number of decisions that influence each decision (Kauffman 1993). The computer then randomly determines the identity of the K other decisions that affect each focal decision. With $N = 6$, $K = 2$, the influence matrix might appear as shown in Figure 1D. There, the contribution of each decision to firm performance is influenced by the resolution of that decision itself and the choices made concerning two randomly assigned other decisions. (Thus, each row contains two off-diagonal x 's.) More generally, K can range from 0 to $N - 1$. $K = 0$ implies that decisions are independent. $K = N - 1$ captures a situation in which each decision is influenced by all others. The parameter K allows us to tune the degree of interaction from full independence (see Figure 1A) to full interaction (see Figure 1B) without specifying particular patterns of influence narrowly.

3.2. Generating Decision Problems

Once the pattern of interaction is set, the computer generates a decision problem. That is, it assigns a payoff to each of the 2^N possible configurations of choices. Recall that the contribution C_i of each decision to overall firm value is affected by other decisions: $C_i = C_i(d_i; \text{other } d_j\text{'s})$. For each possible realization of d_i and the relevant other d_j 's, a contribution is drawn at random from a uniform $U[0, 1]$ distribution. The overall payoff associated with a configuration is the average over the N contributions

$$P(\mathbf{d}) = \frac{\sum_{i=1}^N C_i(d_i; \text{other } d_j\text{'s})}{N}.$$

This procedure for generating payoff functions—stochastically, but with well-controlled patterns of interaction—is adapted from Kauffman's (1993) NK model, a model originally developed in the context of

evolutionary biology. Numerous management scholars have used the procedure in recent years to generate payoff functions that can be employed to examine organizational search (Levinthal 1997, Gavetti and Levinthal 2000, Ghemawat and Levinthal 2000, Rivkin 2000, Marengo et al. 2000, McKelvey 1999 and references therein). It is common to interpret such payoff functions in terms of high-dimensional landscapes. Each of the N decisions constitutes a "horizontal" axis in a high-dimensional space, and each decision offers different options. Resulting from each combination of choices is a payoff for the firm, which is plotted on the vertical axis. The goal of organizational search is to find and occupy a high spot on this landscape, i.e., to select a combination of choices that, together, are highly successful (Siggelkow 2001). Interactions among decisions cause the landscape to become rugged and multi-peaked, making the search for a high peak profoundly more difficult (Kauffman 1993, Rivkin 2000).

3.3. Searching the Landscapes

Having fixed a pattern of interaction, we use the procedure described above to generate many—typically 10,000—landscapes with the same underlying pattern of interaction. Onto each landscape (or equivalently, decision problem), we send a set of firms. Each firm in a set searches for a good configuration of choices. All firms in a particular set start with the same initial configuration of choices. Firms in a set differ, however, in their organizational designs and the capabilities of their management teams. For instance, Firm 1 in a set might have highly capable subordinates while Firm 2 has subordinates of limited ability.

Decomposition: Allocation of Decisions. Each firm has a management team consisting of a CEO, subordinate Manager A, and subordinate Manager B. Manager A has primary responsibility for a subset of the N decisions, and Manager B has responsibility for the complementary subset. We use a string of a 's and b 's to designate a particular allocation of decisions. In a simulation with $N = 6$, for instance, the allocation *abbbba* would indicate that Manager A has responsibility for decisions 1 and 6 while Manager B controls decisions 2–5. We think of decisions 2–5 as Manager

B's department. The more that related decisions are assigned to a single manager—that is, the more often that x 's in the influence matrix correspond to pairs of decisions under one manager—the better an allocation decomposes the decisions.

Subordinate Capability. Search proceeds in a series of periods. In each period, each subordinate manager reconsiders the configuration of choices in his department. Specifically, he compares the status quo to some number, ALT_{SUB} , alternatives. Continuing with the $N = 6$ example mentioned above, suppose that $ALT_{SUB} = 5$ and the current configuration of firm choices is 100111. This means that the current configuration of choices in Manager B's department is 0011. He considers five *local* alternatives to 0011. These include all four of the adjacent alternatives (1011, 0111, 0001, and 0010) and one of the six alternatives that involve changing two decisions. (One of the six is chosen at random.) ALT_{SUB} reflects the cognitive abilities of the subordinate manager. A manager with a higher level of ALT_{SUB} is "smarter"—able to consider more alternatives and able to assess the ramifications of changing more choices at once within his department.

Incentives: Assessment of Alternatives. Each manager ranks the ALT_{SUB} alternatives from most preferred to least. In assessing alternatives, he puts primary weight on the performance of his department, but he may also consider the effects of his changes beyond his domain. $INCENT$, a parameter that ranges from 0 to 1, captures the degree to which the subordinate cares about the ramifications of his actions on the other department. $INCENT = 0$ implies that each manager considers only effects within his department; this may reflect, for instance, a firm in which managers are paid strictly on the basis of local business unit profitability. $INCENT = 1$ implies that each manager is equally concerned with effects outside his department and genuinely wants to maximize firm-wide payoff; this may reflect a firm in which divisional officers are rewarded for overall corporate performance. Continuing with the $N = 6$ example: In assessing any alternative d , subordinate Manager B will consider

$$P'(d) = \{[C_2(d) + C_3(d) + C_4(d) + C_5(d)] \\ + INCENT * [C_1(d) + C_6(d)]\} / 6.$$

In evaluating alternatives, each subordinate assumes that choices in the other department will not change.⁴

Vertical Hierarchy and the Ability of the CEO.

Each subordinate considers the ALT_{SUB} alternatives and the status quo in his department, and sends up to the CEO the P proposals that he most prefers. 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. Note that the term "CEO" need not be taken literally. We use the term as a shorthand for any vertical coordinating mechanism, such as an executive committee, that fulfills functions similar to those outlined below.

We consider two types of CEOs: rubberstamping and active. The first type simply approves Manager A's favorite proposal and Manager B's favorite without review. A firm with a rubberstamping CEO is equivalent to one with no CEO at all. In such firms, decision making has been completely decentralized and subordinate managers have full autonomy over decisions in their departments. In contrast, an active CEO exercises discretion. From roughly P^2 combinations of proposals [$(P$ from Manager A) $\times$ (P from Manager B)], she selects ALT_{CEO} at random, assesses them in light of the interests of the firm as a whole, compares them to the status quo, and selects the option that yields the best payoff for the firm.⁵ ALT_{CEO} reflects the cognitive power of the CEO, or

⁴ Our formulation requires that managers know the total contribution of each department or, equivalent, the performance of the entire firm and the total contribution of one department. Experts in accounting and performance measurement have developed sophisticated techniques to isolate the contributions of individual divisions, product lines, and functional departments (e.g., Kaplan and Atkinson 1998). These techniques may, of course, err in their measurements of contributions. In related work, we are exploring what happens when managers in firms with various organizational designs misperceive performance.

⁵ Each subordinate sends up P options, which may or may not include the departmental status quo. The CEO always considers the departmental status quo as an option, even if the subordinate does not submit it. Hence, the CEO may have as many as $(P + 1)^2 - 1$ new combinations at her disposal (if neither manager submits the status quo as one of his proposals) or as few as $P^2 - 1$ (if both managers submit the status quo as one of their proposals).

more generally, the processing capacity of the coordinating unit.

In the $N = 6$ example with decision allocation *abbbba* and current choice configuration 100111, suppose that $P = 2$ and $ALTCEO = 2$. Manager B might send up for review the alternatives 1011 and 0010 for choices 2–5 while Manager A might propose 11 and 10 for decisions 1 and 6. From the possibilities, the CEO might select configurations 110110 and 100110 for comparison to the status quo 100111. Whichever of these three yields the highest payoff for the firm is selected and implemented. The new configuration becomes the launching point for further search in the next period. Subordinates and CEOs, thus, differ in the type of knowledge they possess. Subordinate managers have local knowledge that allows them to generate proposals for their departments. CEOs possess global knowledge that enables them to assess the full ramifications of departmental choices on overall firm performance.⁶

In sum, firms differ in their organizational arrangements: the grouping of decisions into departments, the amount of information conveyed to senior management (P), the degree to which the CEO acts upon that information (rubberstamping vs. active), and the incentives that managers have to consider effects beyond their domains ($INCENT$). Firms also differ in the abilities of their subordinates (ALT_{SUB}) and their CEOs (ALT_{CEO}). Overall, the organization we envision resembles the one examined by Bower's (1970) classic study of the resource allocation process. Senior management lays out some basic structural elements of the firm: the allocation of decision rights and the incentive system, for instance. Subject to those "rules of the game" (Jensen et al. 1999), lower level managers

select and promote proposals that they find attractive. Senior management then exercises some discretion in selecting among, and integrating across, the proposals that "bubble up."

3.4. Sticking Points

Firms continue to search for many periods. In many (but not all) cases, firms reach "sticking points" after a number of periods. That is, they reach configurations of choices from which they do not move. From a sticking point, there is no alternative configuration of the N choices within the search radius of the firm that meets the approval of enough actors within the firm that the alternative can be adopted. ALT_{SUB} influences how broad the search radius is. Organizational arrangements dictate the standards that an alternative set of N choices must meet to be accepted. For instance, when the CEO exercises discretion, one such standard is that the alternative must yield a higher payoff for the firm as a whole than the status quo achieves. The same standard does not apply when the CEO simply rubberstamps proposals and $INCENT$ is low. Then an alternative that is in the interest of just one department may be implemented.

In conceptions of organizational search, it has been common to think of firms as getting stuck on "local peaks" (e.g., Alchian 1950, Levinthal 1997). A local peak is a configuration of choices for which a change in any single choice leads to worse firm performance, even though simultaneous changes in several choices may improve firm performance. In contrast, our more richly modeled firms may well get stuck at points other than local peaks (Rivkin and Siggelkow 2002). The set of sticking points is neither a sub- nor a superset of the set of local peaks. Suppose, for instance, that a firm with low $INCENT$ and a rubberstamping CEO sits atop a local peak. It is quite possible that a subordinate will discover and implement a move that is beneficial for his department but detrimental for the firm, causing the firm to descend from the peak. Thus, a local peak might not be a sticking point. Similarly, a sticking point need not be a local peak. Consider a situation in which a firm is one decision away from a local peak but the change required to attain the peak is not in the interest of the manager who controls the relevant decision. In such an instance, the manager

⁶ For parsimony, our formulation suppresses a number of considerations worthy of future research. For example, our managers have no cognitive representations (Gavetti and Levinthal 2000), our CEO has no agenda of her own, our firms implement plans without error (Siggelkow 2002), and our subordinate managers do not laterally communicate. Note, however, that our CEO is functionally equivalent to a form of lateral communication in which subordinates rank departmental options, convene in a conference room, consider composite alternatives, and pick the composite that is best for the firm. P then reflects the number of options that subordinates bring to the conference room, and ALT_{CEO} reflects the limited time that managers can afford to spend there.

may never propose the needed change, and a firm may get stuck on the “hillside” below a peak. At such a sticking point, it is also possible that a subordinate will want to make some incremental change, but the CEO will veto it. In that case, the firm is stuck even though it is not on a local peak of the overall firm’s landscape nor is it on a local peak of the lower dimensional “subspace” defined by departmental choices and payoffs.

Likewise, sticking points are related to, but different from, the concept of Nash equilibrium outcomes. With a rubberstamping CEO, all sticking points are indeed Nash equilibrium outcomes in a game played between the two subordinates: At a sticking point, each subordinate is picking the best alternative (for him) given the other subordinate’s current decisions. However, once the CEO becomes active, the sets of sticking points and Nash equilibrium outcomes diverge and become neither sub- nor supersets of each other. At a Nash equilibrium outcome, each player in a game must be taking the best possible action given the strategies of the other players. Unlike players at a Nash outcome, our subordinate managers do not anticipate the CEO’s reaction when they propose alternatives. As a result, they may forego self-beneficial opportunities. For instance, they always send up their most preferred proposals, even when those proposals are subsequently rejected by the CEO. A forward-looking subordinate might opt to make a proposal that he prefers less but the CEO will accept. This failure to consider the CEO’s strategy drives a wedge between the sets of sticking points and Nash equilibria when an active hierarchy is in place. In an appendix available on this journal’s website (mansci.pubs.informs.org), we discuss in detail the subtle relationship between sticking points and Nash equilibrium outcomes.

4. Results

We conducted a comprehensive set of analyses involving each of the five features—vertical hierarchy, incentives, decomposition, degree of decision interaction, and managerial ability—and all the relations among them. A few themes recur in the results. First, certain sets of design elements encourage firms

to search and evaluate a broad array of options while others lead firms to stabilize and cease their search. Second, firms that perform well typically balance search and stability. Third, an organizational design that promotes search is especially effective when underlying decisions intensely interact with one another.

To illustrate these themes, we focus on a series of interdependencies that involve the vertical hierarchy⁷ (see Table 1 for a summary). Mirroring the presentation of conventional wisdom in §2, we first examine the effects of an active CEO in isolation. We confirm that an active CEO can slow down decision making, but we also identify another hazard of an active CEO: She acts as a strong force for stability and can prematurely channel her firm toward a low sticking point before adequate search is undertaken. This hazard is alleviated, however, when the active CEO is coupled with features that encourage broad search—rich information flow, capable subordinates, firm-wide incentives, and incomplete decision decomposition. The stabilizing influence of the active CEO and the search induced by other factors create the interdependencies we report here. Some of the interdependencies confirm conventional wisdom, but others bound or challenge it.

Often in this section, we report that one type of firm achieves a higher level of performance on average than another. In each instance, the difference in mean performance is statistically significant with $p < 0.001$, assuring that reported differences are not simply chance occurrences caused by the stochastic nature of the landscape generation. We report firm performance as a portion of the highest performance attainable on each landscape that was explored.

Robustness. We have observed the qualitative patterns described in this section under a far broader range of parameter values than reported here, and we are happy to share our simulation software with any researcher who wants to probe the robustness of particular results in depth. That said, our primary goal in

⁷ Interdependencies also exist that do not involve the vertical hierarchy, but the interdependencies we examine here are sufficient to illustrate our general themes. Full results that show other interdependencies are available from the authors.

Table 1 Summary of Results

	Conventional wisdom	Results
Active vs. passive vertical hierarchy and degree of interaction	- An active vertical hierarchy can be overloaded, which slows down decision making - An active vertical hierarchy is more valuable when decisions richly interact	- Confirmed when interactions are sparse. When interactions are dense, however, the coordinating benefit of the hierarchy outweighs the cost of overloading. - Confirmed for short-run performance. However, an active hierarchy can undermine long-term performance when interactions are dense and information flow is limited, by rapidly locking the firm into a poor set of decisions.
Active vertical hierarchy and managerial ability	An active hierarchy is less valuable when subordinates are highly capable	An active hierarchy is <i>more</i> valuable when subordinates are highly capable and interactions are dense. An active hierarchy is required to stabilize excessive search of smart managers.
Active vertical hierarchy and firm-wide incentives	An active hierarchy is less valuable when subordinates are rewarded for firm-wide performance	An active hierarchy is <i>more</i> valuable when subordinates are rewarded for firm-level performance, especially when subordinates are highly capable. Firm-wide incentives alone may not achieve stability, because they coordinate only intentions, not actions. The active hierarchy provides necessary stability.
Active vertical hierarchy and decision decomposition	- An active hierarchy is less valuable when decisions can be decomposed so that few cross-departmental interactions remain - Decisions should be allotted to subordinates so that cross-departmental interactions are minimized	- Confirmed. When decisions are completely decomposed, there is no benefit of coordination across departments, so the value of hierarchy is zero. - Confirmed for a rubberstamping vertical hierarchy. With an active hierarchy, however, leaving “unnecessary” interactions between departments can yield higher performance. Incomplete decomposition creates additional search that can be exploited by an active, stabilizing hierarchy.

this paper is not to prove the generality of any single, fine-grained result. Rather, we aim to illustrate particular ways in which the elements of organizational design can relate to one another, and to identify broad drivers of those relationships across a wide range of results.

4.1. Active vs. Passive Vertical Hierarchy and Degree of Interaction

Before considering the interdependency between an active CEO and other design elements, one must carefully examine the effect of an active CEO in isolation. Recall that the active CEO vets the proposals of her subordinates, weighing interactions that local managers ignore and accepting only those changes that serve the firm as a whole. Viewed in this light, the CEO seems to be an unalloyed source of benefit. The coordinative value of an active CEO should be particularly high when interactions are pervasive, received wisdom tells us. Conventional wisdom, however, also points to a potential downside of active

vertical hierarchies: Active CEOs may slow down decision making. In this subsection, we confirm these two pieces of conventional wisdom, but put boundary conditions on both.

We analyze firms that face decision problems with $N = 6$. In each firm, Manager A controls the first three decisions and Manager B the last three; i.e., the decision allocation is *aaabbb*. Each manager considers only one alternative per period ($ALT_{SUB} = 1$). We examine all possible degrees of interaction among the decisions ($K = 0, 1, 2, 3, 4$, or 5). The firms we analyze differ in the degree of CEO activity (i.e., whether the CEO does anything more than rubberstamp proposals), the ability of the CEO to consider multiple options (ALT_{CEO}), and the number of proposals that subordinates submit (P).

Specifically, we consider the four firms described on the top panel of Table 2. In Firm 2A, the CEO rubberstamps proposals. In Firm 2B, each manager sends up his preferred option to a CEO who considers one composite alternative per period. Firm 2C differs

Table 2 Effect of Active vs. Rubberstamping CEO on Performance

Firm	2A	2B	2C	2D
ALTCEO	R	1	1	3
P	—	1	2	1
Performance in period 4				
K = 0	0.951	0.939	0.886	0.950
K = 1	0.884	0.893	0.857	0.905
K = 2	0.842	0.862	0.838	0.877
K = 3	0.814	0.846	0.827	0.861
K = 4	0.797	0.835	0.825	0.850
K = 5	0.780	0.828	0.822	0.842
Performance in period 100				
K = 0	1.000	1.000	1.000	1.000
K = 1	0.943	0.952	0.987	0.954
K = 2	0.910	0.918	0.970	0.921
K = 3	0.894	0.897	0.956	0.900
K = 4	0.886	0.880	0.945	0.884
K = 5	0.882	0.865	0.938	0.869
For K = 5:				
Benefit of active CEO in period 100	—	−0.017*	0.056*	−0.013*
Portion of landscapes on which a firm still wanders after period 80	3.9%	0.0%	2.1%	0.0%
Number of sticking points	3.9	13.0	4.1	13.0
Average height of sticking points	0.869	0.827	0.918	0.827

Note. $N = 6$, $aaabbb$, $INCENT = 0$, $ALT_{SUB} = 1$. R indicates rubberstamping CEO. Performance is an average across 10,000 landscapes. Sticking point results are an average across 500 landscapes.

Differences denoted by * are significant with $p < 0.001$.

from Firm 2B in that each manager sends up two proposals. In Firm 2D, each manager sends up one proposal, but the CEO can consider all new composite alternatives (i.e., as many as three) in each period. We examine the performance of each firm in period 4 and period 100, representing the short run and long run. (We choose periods 4 and 100 simply because most firms take considerably more than 4 and less than 100 periods to reach their long-run levels of performance. Qualitatively, the results are not sensitive to the choice of periods 4 and 100.)

Short-Run Performance. The first six rows of Table 2 compare the ability of different firms to quickly scramble uphill, for varying levels of K . The results for $K = 0$ confirm the conventional wisdom that an active but overloaded CEO can be a liability. Firm 2A, with its rubberstamping CEO and

decentralized decision making, performs as well as or better than any of the firms with active CEOs. An active CEO undermines performance because she is overwhelmed with proposals and becomes a bottleneck, standing between good proposals from the departments and implementation of those proposals. In Firm 2B, for instance, the CEO has up to three composite alternatives to evaluate each period, but can assess only one. Overloaded, she may ignore good proposals that are sent to her by the subordinates. As one would expect, the effect is exacerbated when the number of proposals rises (Firm 2C) and mitigated when the CEO can process more options per period (Firm 2D).

As interactions across departmental boundaries proliferate, the active CEO changes from liability to asset. At high K , Firm 2B fares better than Firm 2A in period 4. On rugged landscapes, the benefits of a CEO who takes account of interactions across departments outweigh the danger of CEO overload. Put differently, when K is high, local managers do not have a good idea of what proposals are valuable for the firm as a whole. Hence, even if the CEO does not get around to looking at the proposal most preferred by a manager, not much harm is done. This finding puts a boundary condition on the conventional wisdom about overloaded CEOs: If interactions across departments are dense enough, an active CEO is beneficial in the short run despite the hazard that she may delay acceptance of good proposals.

Long-Run Performance. Intuition suggests that the CEO-overload problem should fade in the long run. If a good proposal is sent up enough times, at some point the CEO will consider it, accept it, and cease to be a bottleneck. This is particularly salient on a smooth landscape (i.e., $K = 0$). Given enough time, every firm—even one with a very overloaded CEO—should reach the global peak of such a landscape. In the following results, we confirm this intuition, but identify another persistent drawback of an active CEO.

The middle panel of Table 2 reports the performance of each firm in period 100. When $K = 0$, we expect and find that all firms reach the global peak in the long run. For high values of K , however, Firm 2A, which has a rubberstamping CEO, performs better

than Firm 2B, which has an active CEO. Why is the active CEO a detriment even in the long run for Firm 2B? Firms with active CEOs never move downhill on a landscape; the CEO vetoes such a maneuver. As a result, these firms quickly reach sticking points before considering a wide range of alternatives. They run the risk of excessive stability. In contrast, firms with rubberstamping CEOs will sometimes implement alternatives that cause overall performance to decline temporarily. This promotes a wider search of possibilities and can lead, in the long run, to higher performance. The final two rows of Table 2 support this interpretation for the case of $K = 5$. Firm 2B perceives far more sticking points than Firm 2A and is much more likely to get stranded on a low sticking point, i.e., a poor compromise among the subordinates and the CEO. In sum, Firm 2B has more stability and undertakes less search than is optimal when interactions are pervasive.

Employing a smarter CEO, as Firm 2D does, does not alleviate the problem that an active CEO causes when K is high; an ability at the top of the organization to assess more alternatives does not help if too few alternatives are being proposed. On the other hand, it is helpful to insist on a greater flow of information from subordinates to the CEO as Firm 2C does. With more proposals (P) coming in, the danger of premature lock-in diminishes and the benefits of an active CEO—especially her ability to keep subordinates from acting in ways that undermine overall performance—reassert themselves. The search promoted by higher P productively balances the stability provided by the CEO.

The active CEO, thus, has quite different drawbacks in the short run and long run. In the short run, especially if interactions are sparse, the CEO can pose a bottleneck, blocking the rapid progress that a decentralized firm would enjoy. The problem is exacerbated by a greater flow of information and alleviated by the hiring of a smarter CEO. In the long run, an active CEO can funnel a firm prematurely to a mediocre sticking point. The problem is alleviated by a greater flow of information, but not by an increase in the CEO's processing power. The risk of premature lock-in is especially acute when interactions are pervasive. Interactions make the underlying landscape

rugged and multipeaked, which provides a "natural" source of stability. It is in this setting that the additional stability provided by an active, low- P CEO can harm firm performance. Thus, we pinpoint a boundary condition on the conventional wisdom that active hierarchies are more valuable when interactions are pervasive: This appears to be true only if the CEO receives a rich flow of information. If interactions are pervasive and the CEO gets little information, she serves the firm better by rubberstamping subordinates' proposals than by actively exercising oversight.

The long-run results in Table 2 illustrate a general pattern: As interactions among decisions become more pervasive, design elements that encourage broad search grow more important. Comparing Firms 2B and 2C, for instance, we see that the incremental benefit of a richer flow of information, which promotes more search, rises steadily from 0.000 when $K = 0$ to +0.073 when $K = 5$.

The findings in this section, taken as a whole, also suggest that very different vertical hierarchies may be suitable for volatile and stable environments. In volatile settings, firms essentially face a series of short-run problems. Attaining decent results quickly requires either a passive CEO who lets subordinates make final choices (when K is low) or a very smart CEO who acts on the basis of limited information (when K is high). In stable settings, firms can focus on long-run performance, which is best delivered by an active, not-necessarily-brilliant CEO who receives and reviews numerous proposals. This conclusion is reminiscent of Burns and Stalker's (1961) findings on organic and mechanistic organizations and Brown and Eisenhardt's (1997) observations of successful organizational structures in highly dynamic environments. This speculative interpretation deserves further research.

4.2. Active Vertical Hierarchy and Managerial Ability

In the last subsection, we identified the stabilizing effect of the active CEO in the long run, noted that the effect may be harmful when interactions provide natural stability, and illustrated how the search induced by richer information flow may counterbalance the CEO's stability. In the next three subsections, we examine how other design elements and context-

Table 3 Interdependency Between Active CEO and Subordinate Capability with INCENT = 0

Firm	3A	3B	3C	3D	3E	3F	3G	3H	3I	3J	3K	3L
ALTCEO	R	1	1	3	R	1	1	3	R	1	1	3
P	—	1	2	1	—	1	2	1	—	1	2	1
ALTSUB	1	1	1	1	4	4	4	4	7	7	7	7
Performance in period 100 for:												
K = 0	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
K = 1	0.943	0.952	0.987	0.954	0.955	0.964	0.981	0.967	0.951	0.960	0.979	0.964
K = 2	0.910	0.918	0.970	0.921	0.921	0.938	0.965	0.943	0.897	0.929	0.962	0.938
K = 3	0.894	0.897	0.956	0.900	0.900	0.920	0.953	0.926	0.854	0.912	0.948	0.922
K = 4	0.886	0.880	0.945	0.884	0.884	0.909	0.943	0.917	0.825	0.900	0.937	0.915
K = 5	0.882	0.865	0.938	0.869	0.870	0.900	0.936	0.910	0.810	0.896	0.931	0.912
For K = 5:												
Benefit of active CEO	—	−0.017*	0.056*	−0.013*	—	0.030*	0.066*	0.040*	—	0.086*	0.121*	0.102*
Portion of landscapes on which a firm still wanders after period 80	3.9%	0.0%	2.1%	0.0%	36.7%	0.0%	0.3%	0.0%	60.8%	0.0%	0.0%	0.0%
Number of sticking points	3.9	13.0	4.1	13.0	1.2	7.4	2.9	7.4	1.0	9.2	5.1	9.2
Average height of sticking points	0.869	0.827	0.918	0.827	0.928	0.873	0.939	0.873	0.936	0.852	0.898	0.852

Note. $N = 6$, $aaabbb$, INCENT = 0. R indicates rubberstamping CEO. Performance for each level of K is an average across 10,000 landscapes. Sticking point results are an average across 500 landscapes. Differences denoted by * are significant with $p < 0.001$.

tual features may also provide balance. We start by analyzing the effects of “smart” subordinates who can consider a wide array of options within their departments. We show that the search undertaken by such subordinates can balance the stability of an active CEO, making capable subordinates and active CEOs complementary in our model. This runs contrary to conventional wisdom for reasons we explain below.

To analyze the interdependency between hierarchy and subordinate ability, we engage in simulations with the same parameter settings as in the previous subsection, but vary the number of alternatives that each manager is able to evaluate in each period (ALTSUB). The left, middle, and right panels of Table 3 show performance in period 100 for firms with subordinates who are able to evaluate one, four, and seven alternatives per period, respectively. As before, all subordinates pursue departmental performance (INCENT = 0).⁸

The first striking result shown in Table 3 is that a firm can *undermine* its long-run performance by

hiring smarter managers, especially when decisions richly interact and the CEO rubber-stamps decisions. A comparison of the three firms with rubberstamping CEOs—Firms 3A, 3E, and 3I—makes this clear. For $K = 5$, for instance, performance declines steadily as ALTSUB rises from 1 to 4 to 7. For all values of $K > 0$, an increase in ALTSUB from 4 to 7 leads to worse performance. Why can smarter managers undermine long-run performance? The answer lies in the problem that smart subordinates create for each other when their domains influence one another. A smart subordinate searches broadly and undertakes far-reaching changes to improve the performance of his department. In doing so, however, he undermines the improvement efforts that the other equally smart subordinate is making. Hence, a pair of smart subordinates can dance about—each making radical changes that seem like uphill movements from his perspective but that deform the landscape as the other sees it. The higher is ALTSUB, the graver is the danger of this ongoing dance. Reflecting this, the portion of firms that still wander after period 80 rises from 3.9% for Firm 3A to 60.8% for Firm 3I, and the number of sticking points falls from 3.9 to 1.0.

⁸ This subsection and the following two focus on interdependencies as exhibited in long-run performance. Short-run results, available from the authors, display interdependencies that are qualitatively similar.

The broad search pursued by smart subordinates may be harnessed to good effect if an active, stabilizing CEO is present. Accordingly, we see performance rise with ALT_{SUB} , or at least fail to fall dramatically as ALT_{SUB} increases, when a CEO reviews and ratifies proposals. For any given level of ALT_{SUB} , the portion of firms wandering after period 80 falls dramatically and the number of sticking points rises once an active CEO is in place. Moreover, the performance benefit of an active CEO relative to one who rubberstamps proposals, a figure reported in the bottom portion of Table 3 grows with ALT_{SUB} . For instance, for a CEO with $P = 1$ and $ALT_{CEO} = 1$, the benefit increases from -0.017 to $+0.086$ as ALT_{SUB} rises from 1 to 7.

An active hierarchy and capable subordinates are complements in our model; that is, the benefit of having an active CEO rises with ALT_{SUB} . In contrast, received wisdom holds that the two are substitutes: Senior management oversight is less necessary when subordinates are highly capable. The difference stems from distinct perspectives on the role of senior management. To simplify, the conventional view holds that senior managers are responsible for handling exceptions thrown up by the outside world: "The hierarchy of authority is employed on an exception basis. That is, the new situation, for which there is no preplanned response, is referred upward in the hierarchy to permit the creation of a new response" (Galbraith 1973, p. 11). Senior managers are presumed to have the expertise and experience to create new responses and to solve unusually difficult problems (Garicano 2000). The more capable are subordinates, the rarer are such exceptions and the less necessary is an active hierarchy. In contrast, we envision senior managers as the integrators and ratifiers of subordinates' proposals for change. An increase in the potential scope of subordinate proposals makes senior management all the more necessary, especially when decisions interact with one another. We interpret our findings as a boundary condition on conventional wisdom. Where senior managers are repositories of expertise that enable a firm to handle surprising circumstances, an active hierarchy and capable subordinates are likely to be substitutes. But where senior managers serve to check and integrate internal proposals, the two may be complements. Smart managers

may offset any excess stability of the active hierarchy while the hierarchy stabilizes the radical search undertaken by smart managers.⁹

So far, we have presented smarter managers as a source of broader search. There is, however, a way in which smarter managers can prevent wide search and exacerbate the potential for excessive stability. Note in Table 3 that performance sometimes declines with higher ALT_{SUB} even for a firm with an active CEO; for instance, the performance of Firm 3K with $ALT_{SUB} = 7$, is modestly lower than that of Firm 3C with $ALT_{SUB} = 1$ for all $K > 0$. Because active CEOs prevent excessive wandering, the reason for the performance difference in such cases cannot be a failure to achieve stability. Rather, the difference is driven by restricted search. In Firm 3C, each subordinate is essentially forced to send up two random proposals every period. In contrast, in Firm 3K, each subordinate ranks seven alternatives and sends up his two most preferred. As a result, the heterogeneity of proposals received by the CEO is larger in Firm 3C than in Firm 3K, leading Firm 3C to experience wider search. When a subordinate is so smart that the number of alternatives he considers is much greater than the number of proposals he must reveal, the pre-screening he performs can restrict search. In essence, the subordinate is able to hide options he dislikes behind the tall stack of alternatives he has considered. This dynamic plays a part in the next subsection.

4.3. Active Vertical Hierarchy and Firm-Wide Incentives

Conventional wisdom suggests an interdependency between incentive systems that stress firm-wide outcomes and the presence of an active CEO: Firm-wide incentives may reduce or even eliminate the need for an active CEO. Bounding this wisdom, we find situations in which firm-wide incentives can induce broad

⁹ Smart managers and an active hierarchy are complements in the sense of Milgrom and Roberts (1990): The adoption of one increases the benefit of adopting the other. To an actual manager in such a firm, however, the two countervailing items may not feel particularly consistent. A smart manager may be frustrated, for instance, by what he perceives as the meddling of an active CEO. This may even undermine the subordinate's motivation to search for better choices, a consideration we do not model here.

search, thereby making the stability of an active CEO more valuable.

In all previous simulations, each subordinate manager was under a parochial incentive system, evaluating alternatives only from the perspective of his individual department. Here, we conduct simulations that have the same parameter settings as those in the previous subsection except that the value of INCENT is set to 1. Thus, in evaluating alternatives, each subordinate manager takes into account the full effect of his actions on the firm as a whole. Table 4, which reports the results, has precisely the same structure as Table 3.

The key results are most easily seen by comparing the lines labeled "Benefit of active CEO" in Tables 3 and 4. Especially for firms with ALT_{SUB} = 4 or 7, the benefit of having an active CEO increases as INCENT rises from 0 to 1. For example, for Firm 3J with INCENT = 0, the benefit of having an active CEO is +0.086, while for Firm 4J with INCENT = 1, the benefit is +0.211. Accordingly, the active CEO and firm-wide incentives complement one another. This result arises because firm-wide incentives within an active vertical hierarchy promote search that balances the stability provided by the active CEO. Why do firm-

wide incentives promote search in the presence of an active CEO? Recall that the active CEO rejects alternatives that are detrimental to the firm as a whole. When incentives are parochial, the CEO often receives proposals that do not benefit the firm. She rejects these proposals out of hand so the proposals do not lead to effective exploration of the landscape. In contrast, when subordinates face firm-wide incentives, the CEO receives far more proposals that are acceptable to her, and much more movement ensues. In line with this intuition, we see that firms with an active CEO experience fewer sticking points when INCENT = 1 than when INCENT = 0.

In the absence of an active CEO, however, firm-wide incentives can lead to *less* search than parochial incentives, because subordinates have to find alternatives that are beneficial for the firm as a whole, rather than only for their departments. This effect is suggested by the larger number of the sticking points perceived by the rubberstamping firms in Table 4 versus those in Table 3.

Our finding that an active hierarchy is a complement to firm-wide incentives runs contrary to conventional wisdom. The intuition behind the conventional

Table 4 Interdependency Between Active CEO and Subordinate Capability with INCENT = 1

Firm	4A	4B	4C	4D	4E	4F	4G	4H	4I	4J	4K	4L
ALTCEO	R	1	1	3	R	1	1	3	R	1	1	3
P	—	1	2	1	—	1	2	1	—	1	2	1
ALT _{SUB}	1	1	1	1	4	4	4	4	7	7	7	7
Performance in period 100 for:												
K = 0	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
K = 1	0.971	0.970	0.987	0.972	0.973	0.982	0.987	0.986	0.941	0.981	0.987	0.987
K = 2	0.946	0.945	0.970	0.946	0.953	0.964	0.972	0.968	0.864	0.964	0.974	0.972
K = 3	0.922	0.921	0.956	0.923	0.926	0.946	0.957	0.952	0.796	0.948	0.960	0.959
K = 4	0.904	0.902	0.945	0.905	0.894	0.935	0.946	0.942	0.755	0.938	0.948	0.951
K = 5	0.884	0.883	0.937	0.885	0.845	0.925	0.937	0.934	0.718	0.929	0.937	0.945
For K = 5:												
Benefit of active CEO	—	-0.001	0.053*	0.001	—	0.080*	0.092*	0.089*	—	0.211*	0.219*	0.227*
Portion of landscapes on which a firm												
still wanders after period 80	0.0%	0.0%	2.3%	0.0%	30.3%	0.0%	0.1%	0.0%	71.6%	0.0%	0.0%	0.0%
Number of sticking points	9.2	9.2	4.0	9.2	4.9	4.9	2.9	4.9	4.3	4.3	4.0	4.3
Average height of sticking points	0.857	0.857	0.921	0.857	0.903	0.903	0.939	0.903	0.913	0.913	0.917	0.913

Note. $N = 6$, *aaabbb*, INCENT = 0. R indicates rubberstamping CEO. Performance for each level of K is an average across 10,000 landscapes. Sticking point results are an average across 500 landscapes. Differences denoted by * are significant with $p < 0.001$.

wisdom is straightforward: If subordinates have the interests of the overall firm at heart, why does one need oversight? Our answer is that firm-wide incentives can coordinate the *intentions* of subordinates, but they do not necessarily coordinate the *actions* of subordinates when decisions interact. Capable subordinates can engage in aggressive, well-intentioned search that results in mutually destructive “improvement.” This possibility is most vividly illustrated by Firms 3I and 4I, which have highly capable managers and rubberstamping CEOs. Firm-wide incentives cause a precipitous drop in performance for $K > 0$ as subordinates engage in search unchecked by an active CEO. In contrast, firm-wide incentives improve performance when an active CEO is present. The active hierarchy provides a device to coordinate actual moves, not just motives.

If firm-wide incentives encourage broad search in firms with active CEOs, one might expect the marginal benefit of such incentives to rise as interactions among decisions become pervasive, which introduces “natural” stability. A comparison of Tables 3 and 4 confirms this: For firms with active CEOs, the marginal benefit of firm-wide incentives increases as K rises from 0 to 2, then levels off. Comparing Firms 3J and 4J, for instance, we see that the marginal benefit of firm-wide incentives rises from 0.000 for $K = 0$ to +0.021 for $K = 1$ to +0.035 for $K = 2$, then stabilizes for higher K .

4.4. Active Vertical Hierarchy and Decision Decomposition

Lastly, we turn to the interdependency between an active CEO and the decomposition of decisions into departments. As discussed in §2, a consensus exists in the qualitative literature on organizational design that firms should, as much as possible, assign decisions that influence one another to the same manager. This manager is able to “internalize” the interactions among decisions and find the department’s choice configuration that is best for the firm. Moreover, conventional wisdom points out that when decisions are fully decomposed—grouped such that all interactions are internalized—an active vertical hierarchy might be unnecessary.

To study the interdependency between decision decomposition and the active CEO, we examine a series of simulations in which the influence matrix is block-diagonal as shown in Figure 1C. That is, firms face decision problems in which $N = 6$, and all interactions are among decisions 1–3 and decisions 4–6. While the particular profit contributions change from run to run, this pattern of interaction stays the same. Incentives are parochial ($\text{INCENT} = 0$), and each manager considers one local alternative to his current configuration of choices ($\text{ALT}_{\text{SUB}} = 1$). Firms differ in the manner in which decisions are allocated to subordinates and whether their CEOs exercise active discretion (see Table 5). Firms 5A and 5B have the decision allocation *aaabbb*, which completely decomposes the firm into two independent parts, while Firms 5C and 5D have decision allocation *aabbba*, which leaves interactions between the departments. Firms 5A and 5C have rubberstamping CEOs, whereas Firms 5B and 5D have active CEOs.

The simulation results are in line with the second aspect of conventional wisdom noted above: when decisions are completely decomposed, the performance benefit of an active CEO is nil (performance of Firm 5A = performance of Firm 5B). Suggested improvements that are sent up benefit the proposing department and—because of the complete decomposition—have no effect on the other department. Moreover, the most preferred departmental improvements are also the most beneficial for the firm as a whole. As a result, the CEO always accepts

Table 5 Interdependency Between Active CEO and Decomposition

Firm	5A	5B	5C	5D
Decision allocation	<i>aaabbb</i>	<i>aaabbb</i>	<i>aabbba</i>	<i>aabbba</i>
ALTCEO	R	3	R	3
P	—	2	—	2
Performance in period 100	0.937	0.936	0.898	0.979
Benefit of active CEO	—	−0.001	—	0.081*
Portion of landscapes on which a firm still wanders after period 80	0.0%	0.0%	22.1%	0.0%
Number of sticking points	4.1	4.1	1.8	1.8
Height of average sticking point	0.916	0.916	0.897	0.968

Note. $N = 6$, block-diagonal I , $\text{INCENT} = 0$, $\text{ALT}_{\text{SUB}} = 1$. R indicates rubberstamping CEO. Performance is an average across 10,000 landscapes. Sticking point results are an average across 500 landscapes. Differences denoted by * are significant with $p < 0.001$.

preferred proposals of each subordinate and, consequently, acts much like a rubberstamper. Accordingly, the performance of Firm 5B is indistinguishable from that of Firm 5A. In contrast when interactions remain across the departments, the CEO provides valuable stability (performance of Firm 5D $\gg$ performance of Firm 5C). Firm 5C, with a rubberstamping CEO, fails to come to a sticking point on 22.1% of the landscapes. In contrast, Firm 5D with an active, stabilizing CEO, always comes to a steady configuration of decisions.

Intriguingly, the best performance in Table 5 is attained by Firm 5D, which has an active, well-informed CEO and “unnecessary” overlap between departments. This result contradicts one of the most common pieces of received wisdom, that decisions should be allocated to minimize cross-departmental interactions. The cause of Firm 5D’s superior performance is again a helpful balance of search and stability: The overlap across departments generates wide search, as each subordinate proposes—and is sometimes allowed to enact—options that change conditions in the other department, prompting new search in that department. The active CEO ensures that the firm eventually stabilizes around any great option that the wide search produces. Put differently, the incomplete decomposition generates search that the active CEO can take advantage of. Thus, we see imperfect decomposition and an active CEO as complements, the search generated by one working well in concert with the stability generated by the other.¹⁰ In sum, we place an important boundary condition on the conventional wisdom that firms should strive for complete decomposition. “Unnecessary” overlap between departments can induce subordinates to make creative proposals that pry firms off of low sticking points. The ensuing search, coupled

with an active, stabilizing CEO, can produce superior performance.¹¹

5. Discussion

The existence of interdependencies among elements of organizational design has become a bedrock proposition in the literature on organizations, yet relatively little is known about the underlying forces that create these interdependencies. Using a simulation model of organizational design and search, we identify one such force: Interdependencies arise because design elements influence *both* how broadly a firm searches its environment to discover good sets of coordinated choices *and* whether the firm is able to stabilize around those sets once they are discovered. The adoption of an element that encourages broad search typically raises the marginal benefit of other elements that provide offsetting stability. This duality between search and stability has played a central role in substantial prior research—on the productivity dilemma (Abernathy 1978), static and dynamic efficiency (Ghemawat and Ricart i Costa 1993), exploration and exploitation (March 1991), and the ambidextrous organization (Tushman and O’Reilly 1996), for instance. We add to these precedents in three ways.

First, in the context of a formal model of search, we associate search and stability with specific combinations of design elements (see Figure 2). The designers of an organization can promote stability by employing an active CEO, who rejects proposals that make the firm as a whole worse off, or by decomposing decisions such that no cross-departmental interactions remain. Designers can broaden search by increasing the number of proposals sent to the CEO. The effects of firm-wide incentive systems and subordinates who are able to evaluate many alternatives are more subtle; they depend on the degree of discretion the CEO exercises. When the CEO is active and subordinates are allowed to send only a few proposals, smarter subordinates can curtail search because they are more

¹⁰ The bottom of Table 5 supports this interpretation. Compared to the completely decomposed Firm 5B, Firm 5D gets stuck on fewer and higher points; departmental overlap shakes Firm 5D off of low sticking points and encourages it to explore possibilities more widely. Compared to Firm 5C, which has a rubberstamping CEO, Firm 5D is unlikely to wander forever; it stabilizes around its high sticking points. While Firm 5B has only the advantage of stability and Firm 5C has only the advantage of search, Firm 5D enjoys both.

¹¹ As Siggelkow and Levinthal (2003) show, such unnecessary overlap can also be helpful for a rubberstamping firm if the firm subsequently centralizes decision making by assigning all decisions to one department.

Figure 2 Organizational Design Elements and Their Effects on Search and Stability

	Rubberstamping CEO	Active CEO
Active CEO		stability
Rich information flow		search
Smart managers	search	stability*
Firm-wide incentives	stability	search
Decomposition	stability	

*if information flow is limited

able to hide parochially distasteful alternatives from the CEO. In contrast, when the CEO rubberstamps proposals, firms with smarter subordinates engage in broader search because managers are given free rein. Firm-wide incentives can broaden search when the CEO is active by reducing the proportion of proposals that are rejected outright by the CEO, but they tend to restrict search when the CEO is passive because each subordinate then has to find solutions that benefit the firm as a whole, not his department alone.

Our second contribution is to emphasize the need for an organization to strike a balance between search and stability. While much of the prior literature highlights the tension between the two, we focus on ways in which they can work together. We find, for instance, that it can be useful to couple an active, stabilizing CEO with a rich vertical flow of information that promotes search (§4.1). Similarly, the broad search generated by smarter managers (§4.2), by firm-wide incentives (§4.3), or by an incomplete decision decomposition (§4.4) can be harnessed if it is balanced by the stability of an active CEO. Our results thus pinpoint the conditions that make an active vertical hierarchy especially valuable: very capable subordinates, incentives that stress firm-wide outcomes, and decision interactions across departments. Our findings that capable subordinates and firm-wide incentives complement an active hierarchy depart from conventional wisdom because we view top-level managers as integrators, not as exception handlers. Smart managers, even if they have the firm's interest at heart, may still require coordination. Similarly, firm-wide

incentives, while aligning intentions, may still fail to align actions.

Third, we show how the underlying pattern of interaction among decisions affects the appropriate balance between search and stability. The greater is the degree of interaction among decisions, the more rugged are the landscapes that firms face. This ruggedness provides built-in stability. A firm can productively counter this stability by shifting its organizational arrangements in favor of search. Hence, we see a need for a rubberstamping CEO or a rich vertical flow of information when interactions are pervasive and subordinates do not evaluate many alternatives (§4.1). Similarly, we find that the marginal benefits of search-promoting, firm-wide incentives increase with the density of interactions (§4.3).

The three contributions we just identified suggest a set of empirical propositions.

HYPOTHESIS 1. *Both an active CEO and decision decomposition encourage stability in firm choices, ceteris paribus. A rich vertical flow of information, firm-wide incentives in the presence of an active CEO, and more capable subordinates in the absence of an active CEO promote wide search.*

HYPOTHESIS 2. *Organizations that couple design elements that foster search with elements that promote stability will be more successful than those that rely exclusively on one set of elements or the other.*

HYPOTHESIS 3. *Successful organizations in environments with pervasive interactions among decisions will rely more heavily on design elements that promote search than do successful organizations in environments with less pervasive interactions.*

Our findings both support and extend the extensive literature on organizational configurations (e.g., Miles and Snow 1978, Mintzberg 1979, Doty et al. 1993). This line of work contends that only a few internally consistent configurations of organizational design elements exist. This may be true in part because of the interdependencies generated by the need to balance search and stability. In this paper, we have envisioned firms with fixed organizational designs that struggle to find effective combinations of operational choices. One can imagine, however, that in the long

term, firms are also engaged in a search for good organizational designs. That is, they struggle on a higher order landscape whose horizontal dimensions are elements of organizational design. Interdependencies among organizational design elements may give rise to multiple local peaks on this landscape, with each peak corresponding to an organizational configuration. A natural extension of this paper is to allow modeled firms to tweak their internal structures—that is, explore the landscape of organizational elements—and examine whether firms gravitate toward effective organizational designs.

Though the interdependencies we identify may underpin configurations, our results call into question the recommendation that firms pursue pure “consistent” configurations. Some of the configuration literature suggests that a firm should be fully geared toward search or toward stability (e.g., Miles and Snow’s (1978) prospectors and defenders). In contrast, our results highlight the need to balance both attributes in each organization. The difference between our findings and the prior emphasis on pure configurations may arise for at least two reasons. First, scholars of configuration typically acknowledge that pure configurations are simply ideal types, useful for exposition, yet hybrids may arise in reality and have higher performance because of “the need to respond to more than one valid force at the same time” (Mintzberg 1979, p. 474).

Second, while we emphasize the need to balance search and stability, we also stress that contextual variables affect the nature of that balance. Hypothesis 3, for instance, contends that an increase in underlying interactions should tilt the balance toward search. Another critical contextual feature that we have purposely suppressed in this paper is environmental change. A distinguished line of research emphasizes that organizations design themselves in part to cope with environmental change (Burns and Stalker 1961, Chandler 1962, Lawrence and Lorsch 1967). A valuable extension of this paper would place our simulated firms on fluctuating landscapes and examine what kinds of organizational designs deal well with external turbulence. Different types of environmental change might necessitate different combinations of search and stability. Looking at a cross-section

of contexts, one might very well see a pattern that configuration scholars would predict—that numerous organizational elements associated with search are adopted in some settings and many elements associated with stability are adopted in others—even though any single organization requires a balance between search and stability, as we contend.

More broadly, this paper illustrates a general theme that has emerged from agent-based simulations of organizations: connections at one level of analysis drive connections at other levels. In this paper, the underlying interactions among a firm’s decisions shape interdependencies among organizational design elements. Interdependencies among design elements set the stage for organizational configurations, which, in turn, might mold the interplay among competing firms. Such layered connectivity is one of the features that make organizations fascinating yet challenging to study.

Acknowledgments

An earlier version of this paper circulated under the title “Choice Interaction and Organizational Structure.” Special thank to Howard Brenner for computer programming heroics and to Atul Tandia and Kevin Wang for excellent research assistance. For helpful comments, the authors thanks George Baker, Carliss Baldwin, Kathleen Eisenhardt, Lee Fleming, Daniel Levinthal, Will Mitchell, Olav Sorenson, Michael Tushman, Sid Winter, two anonymous referees, an anonymous associate editor, and seminar participants at the University of Michigan, Harvard Business School, University of Pennsylvania, the 2000 Academy of Management meetings, the November, 2000, INFORMS conference, and the 2001 EGOS conference. We are grateful to the Mack Center for Technological Innovation and the Division of Research of Harvard Business School for generous funding. Errors remain our own.

References

- Abernathy, W. J. 1978. *The Productivity Dilemma: Roadblock to Innovation in the Automobile Industry*. John Hopkins University Press, Baltimore, MD.
- Aghion, P., J. Tirole. 1997. Formal and real authority in organizations. *J. Political Econom.* **105** 1–29.
- Alchian, A. A. 1950. Uncertainty, evolution, and economic theory. *J. Political Econom.* **58** 211–221.
- Anderson, P., A. Meyer, K. M. Eisenhardt, K. Carley, A. Pettigrew. 1999. Applications of complexity theory to organization science. *Organ. Sci. (Special Issue)* **10** 233–236.
- Axelrod, R. M., M. D. Cohen. 1999. *Harnessing Complexity: Organizational Implications of a Scientific Frontier*. Free Press, New York.

- Baldwin, C. Y., K. B. Clark. 2000. *Design Rules: The Power of Modularity*. MIT Press, Cambridge, MA.
- Bower, J. L. 1970. *Managing the Resource Allocation Process*. Harvard Business School Press, Boston, MA.
- Brown, S. L., K. M. Eisenhardt. 1997. The art of continuous change: Linking complexity theory and time-paced evolution in relentlessly shifting organizations. *Admin. Sci. Quart.* **42** 1–34.
- Burns, T., G. M. Stalker. 1961. *The Management of Innovation*. Tavistock, London, U.K.
- Camerer, C., A. Vepsäläinen. 1988. The economic efficiency of corporate culture. *Strategic Management J.* **9** 115–126.
- Carley, K. M., Z. Lin. 1997. A theoretical study of organizational performance under information distortion. *Management Sci.* **43** 976–997.
- , D. M. Svoboda. 1996. Modeling organizational adaptation as a simulated annealing process. *Sociological Methods and Res.* **25** 138–168.
- Chandler, A. D., Jr. 1962. *Strategy and Structure: Chapters in the History of Industrial Enterprise*. MIT Press, Cambridge, MA.
- Chang, M.-H., J. E. Harrington, Jr. 2000. Centralization vs. decentralization in a multi-unit organization: A computational model of a retail chain as a multi-agent adaptive system. *Management Sci.* **46** 1427–1440.
- Child, J. 1984. *Organization: A Guide to Problems and Practice*, 2nd ed. Harper & Row, New York.
- Doty, D. H., W. H. Glick, G. P. Huber. 1993. Fit, equifinality, and organizational effectiveness: A test of two configurational theories. *Acad. Management J.* **36** 1196–1250.
- Eppinger, S. D., D. E. Whitney, R. P. Smith, D. A. Gebala. 1994. A model-based method for organizing tasks in product development. *Res. Engrg. Design* **6** 1–13.
- Fleming, L., O. Sorenson. 2001. Technology as a complex adaptive system: Evidence from patent data. *Res. Policy* **30** 1019–1039.
- Galbraith, J. R. 1973. *Designing Complex Organizations*. Addison-Wesley, Reading, MA.
- Garicano, L. 2000. Hierarchies and the organization of knowledge in production. *J. Political Econom.* **108** 874–904.
- Gavetti, G., D. Levinthal. 2000. Looking forward and looking backward: Cognitive and experiential search. *Admin. Sci. Quart.* **45** 113–137.
- Ghemawat, P., D. Levinthal. 2000. Choice structures and business strategy. Working paper 01–012, Harvard Business School, Boston, MA.
- , J. E. Ricart i Costa. 1993. The organizational tension between static and dynamic efficiency. *Strategic Management J.* **14** 59–73.
- Gibson, J. L., J. M. Ivancevich, J. H. Donnelly, Jr. 2000. *Organizations: Behavior, Structure, Processes*. Irwin McGraw-Hill, Boston, MA.
- Harris, M., A. Raviv. 2002. Organization design. *Management Sci.* **48** 852–865.
- Holmström, B., P. R. Milgrom. 1994. The firm as an incentive system. *Amer. Econom. Rev.* **84** 972–991.
- Jensen, M. C., W. H. Meckling, G. P. Baker, K. H. Wruck. 1999. Coordination, control, and the management of organizations: Course notes. Working paper 98–098, Harvard Business School, Boston, MA.
- Kaplan, R. S., A. A. Atkinson. 1998. *Advanced Management Accounting*. Prentice-Hall, Englewood Cliffs, NJ.
- Kauffman, S. A. 1993. *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press, New York.
- . 1995. *At Home in the Universe: The Search for Laws of Self-Organization and Complexity*. Oxford University Press, New York.
- , J. Lobo, W. G. Macready. 2000. Optimal search on a technology landscape. *J. Econom. Behavior and Organ.* **43** 141–166.
- Khandwalla, P. N. 1973. Viable and effective organizational designs of firms. *Acad. Management J.* **16** 481–495.
- Lawrence, P. R., J. W. Lorsch. 1967. *Organization and Environment*. Harvard Business School Press, Boston, MA.
- Levinthal, D. A. 1997. Adaptation on rugged landscapes. *Management Sci.* **43** 934–950.
- , M. Warglien. 1999. Landscape design: Designing for local action in complex worlds. *Organ. Sci.* **10** 342–357.
- March, J. G. 1991. Exploration and exploitation in organizational learning. *Organ. Sci.* **2** 71–87.
- Marengo, L., G. Dosi, P. Legrenzi, C. Pasquali. 2000. The structure of problem-solving knowledge and the structure of organizations. *Indust. and Corporate Change* **9** 757–788.
- McKelvey, B. 1999. Avoiding complexity catastrophe in coevolutionary pockets: Strategies for rugged landscapes. *Organ. Sci.* **10** 294–321.
- Miles, R. E., C. C. Snow. 1978. *Organizational Strategy, Structure and Process*. McGraw-Hill, New York.
- Milgrom, P. R., J. Roberts. 1990. The economics of modern manufacturing: Technology, strategy, and organization. *Amer. Econom. Rev.* **80** 511–528.
- , ———. 1995. Complementarities and fit: Strategy, structure, and organizational change in manufacturing. *J. Accounting and Econom.* **19** 179–208.
- Miller, D., P. H. Friesen. 1984. *Organizations: A Quantum View*. Prentice-Hall, Englewood Cliffs, NJ.
- Mintzberg, H. 1979. *The Structuring of Organizations*. Prentice-Hall, Englewood Cliffs, NJ.
- . 1981. Organization design: Fashion or fit? *Harvard Bus. Rev.* **81**(1) 103–116.
- Nadler, D. A., M. L. Tushman. 1997. *Competing by Design*. Oxford University Press, New York.
- Porter, M. E., N. Siggelkow. 2002. Contextual interactions within activity systems and sustainable competitive advantage. Working paper, Harvard Business School, Boston, MA.
- Rivkin, J. W. 2000. Imitation of complex strategies. *Management Sci.* **46** 824–844.
- . 2001. Reproducing knowledge: Replication without imitation at moderate complexity. *Organ. Sci.* **12** 274–293.
- , N. Siggelkow. 2002. Organizational sticking points on NK-landscapes. *Complexity* **7**(5) 31–43.
- Siggelkow, N. 2001. Change in the presence of fit: The rise, the fall, and the renaissance of Liz Claiborne. *Acad. Management J.* **44** 838–857.

- . 2002. Misperceiving interactions among complements and substitutes: Organizational consequences. *Management Sci.* **48** 900–916.
- Siggelkow, N., D. Levinthal. 2003. Temporarily divide to conquer: Centralized, decentralized, and reintegrated organizational approaches to exploration and adaptation. Working paper. Reginald H. Jones Center, Wharton School, University of Pennsylvania, Philadelphia, PA.
- Simon, H. A. 1957. *Models of Man: Social and Rational; Mathematical Essays on Rational Human Behavior in a Social Setting*. John Wiley and Sons, New York.
- . 1962. The architecture of complexity. *Proc. Amer. Philos. Soc.* **106** 467–482.
- . 1973. Applying information technology to organization design. *Public Admin. Rev.* **33** 268–278.
- Steward, D. V. 1981. The design structure system: A method for managing the design of complex systems. *IEEE Trans. Engrg. Management* **28**(3) 71–74.
- Thompson, J. D. 1967. *Organizations in Action*. McGraw-Hill, New York.
- Tushman, M. L., C. A. O'Reilly. 1996. Ambidextrous organizations: Managing evolutionary and revolutionary change. *California Management Rev.* **38**(4) 8–30.

Accepted by Daniel Levinthal; received July 2002. This paper was with the authors 5 months for 2 revisions.