

Temporarily Divide to Conquer: Centralized, Decentralized, and Reintegrated Organizational Approaches to Exploration and Adaptation

Nicolaj Siggelkow • Daniel A. Levinthal

2017 Steinberg Hall-Dietrich Hall, Wharton School, University of Pennsylvania,
Philadelphia, Pennsylvania 19104

2028 Steinberg Hall-Dietrich Hall, Wharton School, University of Pennsylvania,
Philadelphia, Pennsylvania 19104

siggelkow@wharton.upenn.edu • levinthal@wharton.upenn.edu

Abstract

To create a competitive advantage, firms need to find activity configurations that are not only internally consistent, but also appropriate given the firm's current environment. This challenge is particularly acute after firms have experienced an environmental change that has shifted the existing competitive landscape and created new, high-performing sets of activity choices. How should firms organize to explore and search such an altered performance landscape? While it has been noted that adaptive entities need to maintain a balance of exploration and exploitation, little is known about how different organizational structures moderate this balance. With the help of an agent-based simulation model, we study the value of three different organizational structures: a centralized organization, in which decisions are made only at the level of the firm as a whole; a decentralized organization, in which decisions are made independently in two divisions; and a temporarily decentralized firm, which starts out with a decentralized structure and later reintegrates.

We find that if interactions among a firm's activities are pervasive, neither the centralized nor the permanently decentralized organizational structure leads to high performance. In this case, temporary decentralization—an organizational structure that has not found much attention in the literature—yields the highest long-term performance. This organizational structure allows the firm both to avoid low-performing activity configurations and to eventually coordinate across its divisions. Thus, even if the decision problem a firm faces is not fully decomposable, a temporary bifurcation can lead to a higher long-term performance outcome. Initial decentralized exploration is, however, costly in the short run, as compared to centralized exploration. As a result, a tradeoff exists between the short-term costs of decentralized exploration and the long-term benefits of reaching higher performance. As interactions across and within divisions increase, the optimal length of decentralized exploration tends to grow.

Paralleling our first result, we further show that even if a decision problem is decomposable, that is, can be perfectly

modularized, it can be beneficial to create a temporary decision allocation that creates “unnecessary” interdependencies across the subsystems. This benefit arises in particular when the modules are complex by themselves. In both cases, an initial phase of exploration, enabled by an appropriate organizational structure, followed by refinement and coordination, enabled by a different structure, leads to high performance. To illustrate our general model, we focus on incumbent firms' responses to the Internet and discuss implications for the product design process.

(Organizational Adaptation; Organizational Design; Activity Systems; E-Commerce; Agent-Based Simulations)

1. Introduction

Management scholars have exhorted firms to maintain a tight coherence among their multidimensional policy and activity choices to create systems with high internal fit (e.g., Miller and Friesen 1984, Porter 1996). To create a competitive advantage—and at times simply to survive—firms face the challenge of identifying activity configurations that are not only internally reinforcing, but also appropriate, given the firm's current environment. For incumbent firms, this challenge is particularly acute after environmental changes, such as technological shifts, that allow new ways of performing existing activities and that create new, high-performing, coherent sets of activity and policy choices, i.e., new business systems.

How should firms organize to explore and search such a new space of possibilities? Consider, for instance, the impact of the Internet and the World Wide Web on incumbent firms. To many firms, the Web constituted a significant environmental shock, creating an altered

performance landscape that enabled many new activity configurations and that required firms to readjust their existing activity systems. Which organizational structure to adopt in the search for an appropriate strategic response to the new competitive landscape was a nontrivial question for many firms, and a large variety of organizational responses ensued. Some firms pursued centralized change efforts in which the firm incorporated its e-commerce activities seamlessly into the existing organization. For instance, The Gap, a fashion retailer, considered its website simply as one more store and serviced it using its existing infrastructure. Similarly, The Vanguard Group and Dell considered the Internet as one more distribution channel to be exploited by the existing organization. In contrast, other firms pursued more decentralized search efforts. For instance, Bank One formed an independent Web subsidiary, Wingspan, to explore Internet-based banking. Likewise, Disney created Go.com, grouping all its Web-based activities into one organizational entity and creating a separate tracking stock for this business.

The academic literature has offered a paralleling variety in its prescriptions: Agreement exists that organizational structure plays a critical role in the search and exploration of new competitive landscapes (Schoonhoven and Jelinek 1990), but views differ on what organizational structures are particularly suitable. For instance, Christensen (1997) argues that the claims of established business lines on a firm's resources are such that the only hope for a rapid and effective response to a major shift in a firm's competitive landscape is to create an independent business to pursue the possible new business models suggested by these changes. Other scholars (Brown and Eisenhardt 1997, Burgelman 1991, Tushman and O'Reilly 1996) suggest that firms can engage in a mix of new, exploratory initiatives at the same time they devote substantial resources to continuing to pursue existing initiatives within the same structural context.

Common to these perspectives is the premise that adaptive entities are charged to maintain a balance of exploration and exploitation—a theme with deep roots in the literature on complex adaptive systems (Holland 1975) that has become central in the management discourse as well (March 1991). Perspectives differ, however, on the suggested organizational structure to manage this process. On one hand, firms need to search broadly for new activity configurations. On the other hand, firms need to coordinate across their interdependent activities to avoid misfits and instability. While the call to balance search and coordination has been well understood,

there is still little known about how different organizational structures moderate this balance. Using an agent-based simulation model, we start making an attempt to fill this gap in a structured manner. To focus our analysis, we concentrate on a particularly salient aspect of organizational structure, the degree of centralization of decision making. Thus, while our main motivating examples are concerned with organizational responses to the Internet, our model addresses a broader organizational phenomenon by examining the effects of centralized and decentralized decision making on the balance of exploration and coordination.

To be precise, we call an organizational structure “decentralized” when decision making has been disaggregated into a number of subunits, or divisions, each making its own decisions. In contrast, an organizational structure is called “centralized” when decisions are made only at the level of the firm as a whole. Since we conceptualize firms as systems of interdependent activities, the key contextual variable of our analysis is the degree and pattern of interactions among the activities in which a firm is engaged. If the decisions that a firm has to make concerning its activity choices can be grouped such that all interactions are captured within separate divisions, and no cross-division interdependencies exist, the overall decision problem the firm faces is called “decomposable” (Simon 1962). The decomposability of decision problems has recently attracted renewed attention in the literature on modularity (Schilling 2000). Firms have been exhorted to decompose decision problems, i.e., to “modularize” them, in particular in the context of product design (e.g., Sanchez 1995). Modularization of design problems has been seen, for instance, as a way to speed up product improvement (Baldwin and Clark 2000) because decision problems that have been divided into independent subproblems reduce the amount of required coordination.

Thus, at the heart of our analysis lie different combinations of organizational structures (centralized/decentralized) and decision problems (decomposable/nondecomposable). For which types of interaction patterns among decisions is which type of organizational structure suitable? Related to our study is prior work concerning the organizational implications of product modularity, which poses that modular product design will be paralleled by a “modular” or decentralized organizational structure (Sanchez and Mahoney 1996). Similarly, Galunic and Eisenhardt (2001) and Eisenhardt and Brown (1999) consider the benefits of modular organizational structures when “business charters” need to be reassigned, while Schilling and Steensma (2001) study the industry characteristics under which, in their terms,

modular organizational features such as contract manufacturing and alliance formation arise.

To summarize our main result, we find that if interactions among a firm's activities are nondecomposable, neither a centralized nor a permanently decentralized organizational structure leads to high performance. In this case, temporary decentralization with subsequent reintegration—an organizational structure that has not found much attention in the literature—yields the highest long-term performance. This organizational structure allows the firm to both initially explore possible solutions and eventually coordinate across its divisions. Similarly, we find that if a decision problem is decomposable, it can be beneficial not to decompose—i.e., modularize—the system from the beginning. A temporary introduction of “unnecessary” interdependencies can yield higher long-term performance, especially when the subunits of the decomposed system are complex. In both cases, an initial phase of exploration, enabled by an appropriate organizational structure, followed by refinement and coordination, enabled by a different structure, leads to high performance.

We apply the findings of our analysis to the recent challenge faced by incumbent firms to adapt to the changes and the economic opportunities brought about by the Internet. Interestingly, in this context, a number of temporary decentralized structures with subsequent reintegration did in fact arise. Since pervasive interdependencies between the Internet divisions and the mother firms are likely to exist, our model would suggest that such temporary decentralization followed by reintegration is indeed a helpful search strategy, whether the temporary aspect of the decentralization was originally intended or not.

The rest of the paper is organized as follows. The next section provides a detailed description of the notion of performance landscapes—a helpful framework for modeling firms as systems of interdependent choices. In §3, we analyze the implications of centralized, decentralized, and reintegrated organizational structures for organizational adaptation. Section 4 discusses the implications of the findings, while §5 concludes.

2. Performance Landscapes


Firms engage in a wide array of choices, for instance, with respect to product line, target markets, marketing, sales, distribution, manufacturing, human resources, purchasing, research and development, and finance and control. The view that these choices have to fit together and form consistent configurations has a long tradition in the management literature (Chandler 1962, Drazin and

Van de Ven 1985, Learned et al. 1961, Meyer et al. 1993, Miller 1986, Miller and Friesen 1984, Mintzberg 1979). While previous work on fit focused on consistency among functional areas, the recent management literature has adopted a more fine-grained level of inquiry, focusing on fit at the level of individual activities (e.g., Porter 1996, Siggelkow 2002). In particular, the implications of interactions among choices made by a firm have found much recent attention (e.g., Milgrom and Roberts 1990, Siggelkow 2001, Whittington et al. 1999).

The analytical framework of a performance landscape (Kauffman 1993; Levinthal 1997; Rivkin 2000, 2001) allows us to highlight organizational consequences of interactions among choices. A performance landscape is a mapping of all possible sets of a firm's choices onto performance values (such as a profitability measure). If a firm's set of choices is described by a vector of N choices, then the performance landscape consists of N dimensions depicting the firm's alternatives along each dimension and an $(N + 1)$ th dimension depicting the performance associated with each vector of N choices. For a simple example with $N = 2$, see Figure 1. For each combination of X and Y , the performance landscape depicts the resulting performance V . For instance, one dimension might be product variety, while the other dimension characterizes the flexibility of the manufacturing setup.

The degree of interaction among a set of choices affects the ruggedness of the performance landscape

Figure 1 Performance Landscape



(Kauffman 1993). For instance, if the payoff to each individual choice is independent of all other choices, the landscape will be smooth and single peaked. A set of choices $P = (a_1, a_2, \dots, a_N)$ represents a peak in the landscape, if there is no other set of choices that differs from P in only one element and yields a higher performance than P . In other words, no incremental change of P exists that is performance enhancing.¹

If choices interact, the performance landscape will typically have more than one peak. Moreover, the greater the number of choices that interact with each other, the more rugged and multi-peaked a landscape becomes. One can think of interactions as imposing constraints on efforts to improve performance with respect to a single variable. Shifting a focal variable may yield some positive direct effects, but it may also generate a number of negative indirect effects on the performance contributions of other variables. In such an interconnected system, appropriate choices are highly contextualized, i.e., dependent on the values of other activity choices, leading to a number of distinct, consistent sets of activity configurations (Porter and Siggelkow 2002).

Aspects of search behavior that have been studied in prior work using the framework of performance landscapes include the efficacy of incremental versus more radical changes (Levinthal 1997) and the effect of simplified mental representations of the underlying performance landscape (Gavetti and Levinthal 2000). The effect of organizational structure on search performance has received only limited attention (e.g., Rivkin and Siggelkow 2003, Ethiraj and Levinthal forthcoming). In this paper, we focus on a particular feature of organizational structure: the degree of decentralization. As a firm searches for new activity configurations, does it search as one entity, or does it create divisions that search independently for division-level activities that yield high performance from each division's own perspective?

3. A Simulation Model of Organizational Exploration

To search a performance landscape, firms can organize in different ways. For instance, in the context of the Internet, some firms adopted decentralized structures. These firms created internal divisions focusing on Web activities or created free-standing subsidiaries. Thus, these firms created more or less independent, new organizational divisions. In contrast, other firms engaged in centralized and integrated responses, in which their Web activities were fully integrated into the firm's set of existing activities. Integrated responses were not neces-

sarily incremental. While some firms were able to incorporate Web activities through incremental adaptation of their existing business system, other firms engaged in higher-dimensional, systemic reconfigurations. The common feature of these responses was, however, the firms' unified organizational structure during their search process.


To study more formally the efficacy of different organizational structures for discovering good activity configurations, we use an agent-based simulation model. Firms with different organizational structures are "released" on performance landscapes. Over a number of periods, firms search this performance landscape for high-performing activity configurations, i.e., high points on the landscape. By comparing the performance of firms with different organizational structures over a large number of landscapes, we are able to identify the typical effects of different organizational structures on firm performance.

3.1. Model Specification

3.1.1. Modeling Performance Landscapes. Conceptualizing firms as systems of activities, we assume that a firm has to make decisions concerning N activities, $a_1, \dots, a_N$. For simplicity, it is assumed that each activity can take on two states. For instance, a_1 may represent the decision to sell products directly ($a_1 = 1$) or not ($a_1 = 0$), while a_2 may represent the decision to use resellers ($a_2 = 1$) or not ($a_2 = 0$). As described above, a performance landscape is a mapping of any possible set of firm activities $A = (a_1, a_2, \dots, a_N)$ to performance values $V(A)$. We create performance landscapes with a variant of the NK model (Kauffman 1993), which has been employed in a number of organizational studies (for a survey, see Sorenson 2002). The value of each individual activity a_i is affected by both the state of the activity itself and the states of a number of other activities $\mathbf{a}_{-i}$. For instance, the value of selling products directly is likely to be affected by the presence or absence of existing distribution through resellers. (Consider, for example, the problems encountered by Compaq to duplicate Dell's direct distribution model.) Denote the value of activity a_i by $c_i(a_i, \mathbf{a}_{-i})$. For each landscape, the particular values of all possible c_i 's are determined by drawing randomly from a uniform distribution over the unit interval, i.e., $c_i(a_i, \mathbf{a}_{-i}) \sim u[0, 1]$. The value of a given set of activities A is then given by

$$V(A) = [c_1(a_1, \mathbf{a}_{-1}) + c_2(a_2, \mathbf{a}_{-2}) + \dots + c_N(a_N, \mathbf{a}_{-N})] / N.$$

The identity of $\mathbf{a}_{-i}$, i.e., the set of activities that affect each activity a_i , is given by the interaction structure of

Figure 2 Different Interaction Structures


the firm's decision problem. The interaction structure is either prespecified directly or generated randomly in a controlled manner. The two interaction structures that we prespecify are depicted in panels A and B of Figure 2. An interaction structure is an $N \times N$ matrix, where an $\mathbf{x}$ in column j , row i , denotes that activity j affects activity i . In all simulations, we use $N = 6$. Panel A depicts a block-diagonal structure. In this case, activities a_1 – a_3 all affect each other and activities a_4 – a_6 all affect each other, but no interactions occur between any of the activities a_1 – a_3 and any of the activities a_4 – a_6 . Panel B depicts the case of full interdependence in which all activities affect each other. When interaction structures are generated randomly, as in the typical NK setup (Kauffman 1993), only the overall degree of interaction is specified. In particular, the parameter K denotes the number of off-diagonal entries in each row, i.e., the number of activities that each focal activity is affected by. Panel C depicts an example of a $K = 2$ random interaction structure. One should note that the block-diagonal interaction structure shown in panel A is a special case of a $K = 2$ interaction structure. Similarly, the case of full interdependence shown in panel B corresponds to the case of $K = 5$.

3.1.2. Alternative Organizational Structures. We model three different firms: a centralized firm, a decentralized firm, and a temporarily decentralized firm that later reintegrates. After a landscape is created, all three

firms are placed on the same, randomly chosen spot on the landscape. In each subsequent period, each firm attempts to find performance-improving changes in its activity configuration. For illustrative purpose, assume that a firm is located at $a_i = 0$ for all six activities, i.e., $A = 000000$. The three firms differ in the way they evaluate and choose performance-improving alternatives. A centralized firm will evaluate a randomly chosen local alternative and compare it to the current status quo. Local alternatives are activity configurations that differ in one activity from the status quo, in this case, e.g., 001000. If the centralized firm finds that the alternative has a higher performance value than the status quo—in this case, $V(001000) > V(000000)$ —the firm will implement the alternative. If the alternative is implemented, say 001000, the firm will search in the next period for further improvements, using 001000 as a starting point for local search. If the evaluated alternative has a lower performance than the status quo, $V(001000) < V(000000)$, the firm will continue, in the following period, to look for improvements around the status quo, i.e., around 000000.

The decentralized firm is assumed to consist of two divisions M and W (e.g., the mother firm and the Web division). Division M controls decisions a_1 – a_3 ; division W controls decisions a_4 – a_6 . In each period, each division is searching for local improvements for its own division. In evaluating alternatives, each division takes only the payoff to its own division into account. Thus, rather than using $V(\bullet)$ to evaluate alternatives and the status quo, division M computes

$$V_1 = [c_1(a_1, \mathbf{a}_{-1}) + c_2(a_2, \mathbf{a}_{-2}) + c_3(a_3, \mathbf{a}_{-3})]/3$$

while division W computes

$$V_2 = [c_4(a_4, \mathbf{a}_{-4}) + c_5(a_5, \mathbf{a}_{-5}) + c_6(a_6, \mathbf{a}_{-6})]/3.$$

In evaluating alternatives, each division takes the other division's status quo choices as given. For instance, if the firm as a whole is currently at 000000, division M might evaluate alternative 100 (given W 's current choices of 000) and division W might evaluate 010 (given M 's current choices of 000). If either division finds a profitable change as compared to the status quo, it is implemented. For instance, if M finds 100 to be more profitable than 000, and W finds 010 to be less profitable than 000, the firm's overall activity configuration would become 100000. In the next period, each division would again search for local, division profit-enhancing alternatives, starting its search at 100 and 000, respectively.

Three features of the decentralized firm should be noted. First, the decentralized firm is able to implement more changes in any given period than the centralized

firm. Given that decisions are divided between two divisions, the decentralized firm can implement up to two changes in any period, while the centralized firm can only implement one change. The rationale for this modeling assumption is to capture the effect of parallelism that is afforded by decentralized structures. In fact, a key reason for decentralization is the ability to overcome the slowness of sequential decision processes: If all decisions have to be made at one central location, the process is generally slower than if decisions can be made independently and simultaneously by more than one actor.²

Second, the assumption that each division is using its division profit (V_1 or V_2) to determine which alternative to implement can be interpreted in two ways, based either on incentives or on information availability. The incentive-based interpretation poses that divisions are treated as profit centers, i.e., that managerial incentives are linked entirely to division performance. In this case, divisions would implement those decisions that benefit them most, regardless of firmwide implications. Alternatively, divisions may not have sufficient information to understand the implications of their actions on other divisions. In this case as well, they may choose to ignore these externalities and base their decisions exclusively on the payoffs to their own divisions.


Third, by assigning the first three decisions to one division and the last three decisions to the other division, the decentralized firm's decision allocation perfectly decomposes the block-diagonal interaction structure (Figure 2A). In other words, no cross-divisional interdependencies exist for the decentralized firm when the interaction structure is block-diagonal. In contrast, when the interaction structure is generated randomly, the decentralized firm's decision allocation generally does not decompose the interaction structure; i.e., cross-divisional interdependencies typically exist. For example, for the $K = 2$ interaction structure depicted in Figure 2C, activity a_6 , which is under control of the second division, affects activity a_1 , which is under control of the first division. In the case of full interdependence ($K = 5$, Figure 2B), cross-divisional interdependencies exist in any decision problem that the decentralized firm faces.

Finally, the third type of firm is the reintegrator. Whereas centralized and decentralized firms never change their organizational structure, the reintegrator starts out with a decentralized structure, but switches to a centralized structure after period t .

3.2. Results

We study the effectiveness of the three different organizational forms for different types of interaction structures. The performance of each firm is expressed as the

Figure 3 Performance on Decomposable Landscape



percentage of the highest possible performance in any given landscape. Thus, a firm that is on the global peak has a performance of 1.00. Results are averaged over 10,000 landscapes. If interaction structures are not prespecified, both the interaction structure and the contributions of the various possible configurations of choices are generated anew in every landscape; with prespecified interaction structures, only the contributions are generated anew in every landscape.

3.2.1. Search in a Decomposable Environment. Figure 3 shows the results for the centralized and decentralized firms and for a firm that reintegrates after period 25, when firms face landscapes that have an underlying block-diagonal interaction structure.

In the short run, the decentralized firm has a performance advantage over the centralized firm, while in the long run both firms converge on the same performance. The short-run advantage of the decentralized firm shows the value of parallel search, which allows the decentralized firm to change its activity set more quickly than the centralized firm. Since the partitioning of the decision problem in the decentralized firm matches the underlying interaction structure, decisions of one division do not affect the decisions in the other division. As a result, any parochial decision made in one division is also beneficial for the firm as a whole. Put conversely, given a decomposed activity interaction structure, no benefit of coordination is provided by centralization. At the same time, in the long run, both firms reach similar average performance levels. Since both types of firms consider only local changes, they do not always reach the global peak, but can get stuck on local peaks. As a result, their performance is less than 1.00.

The reintegrator's performance is practically identical to that of the decentralized firm. Up to period 25, both firms share the same organizational structure and hence have the same performance (except for statistical fluctuation). By period 25, all firms that are decentralized are stuck on local peaks; i.e., neither division can find a local change that is performance enhancing. After reintegration, the firm is still looking only at local changes. Moreover, since decisions from one division do not affect the payoffs of the other division, the reintegrated firm ranks alternatives in the same manner as the decentralized firm. As a result, the firm remains stuck at the same point. Thus, reintegration does not provide any performance benefit over keeping a decentralized structure.

To summarize, if decisions can be grouped so that all interactions are contained within divisions, a decentralized firm that decomposes the decision problem has a short-term advantage, while yielding a similar long-term performance as the centralized or reintegrated firm.

3.2.2. Search in a Nondecomposable Environment.

The findings reported in the previous section change dramatically when activity interactions cannot be perfectly decomposed. Figure 4 shows the results for the case of interaction structures randomly specified with $K = 2$.³ In this case, even though the decentralized firm still has the ability to change faster than the centralized firm, the decentralized firm no longer has a short-term performance advantage. Given that activities interact across

Figure 4 Performance on Nondecomposable Landscapes ($K = 2$)





Figure 5 Individual Performance Histories



divisions, the benefit of coordination—i.e., of taking into account the externalities that decisions in one division have on the other division—is large. The decentralized firm, in which interdependencies across divisions are ignored, suffers from this lack of coordination in the short run and especially in the long run.⁴

The most intriguing result, however, is the performance of the reintegrated firm. Since it is decentralized until period 25, its early performance is again practically identical to the decentralized firm. After reintegration in period 25, its performance improves. Strikingly, its performance improves to become significantly higher than that of an organization that is centralized throughout. (All performance differences that are mentioned in this paper are significant with $p < 0.01$.)

The increase in performance subsequent to integration stems from the coordination that is provided by the centralized organizational structure. To gain more intuition for this effect, it is helpful to observe a number of individual performance histories. Figure 5 displays the performance of two firms that start decentralized and reintegrate after period 25 and one firm that is permanently centralized. In contrast to Figure 4, the performance given in any period in Figure 5 is the performance of a single firm on a single landscape, rather than the average performance of firms of a particular type over 10,000 landscapes.

While the performance values shown in Figure 5 are idiosyncratic, the shapes of the performance histories are fairly typical. Firms that are permanently centralized improve their performance in a number of discrete steps. Moreover, the overall performance of these firms never declines because only performance-improving

local adaptations are adopted. Firms that start out in a decentralized form follow two broad classes of developments, represented by Reintegrator 1 and Reintegrator 2. Common to both is an early phase of exploration in which both divisions pursue their own interests. Even though the firm may encounter good overall solutions, division managers might find alternatives for their departments that would improve their division's performance but reduce overall firm performance. Due to parochial interests, these alternatives will be implemented and overall performance of the firm can decline.

One basic pattern, represented by Reintegrator 1, is a cyclic pattern of performance, at times rising and then subsequently falling, with the organization never fixating on a particular set of choices. Underlying this pattern is the fact that each subunit, in its efforts to modify its choices to improve its own performance, perturbs the performance landscape of its sister division. Each division is essentially trying to climb uphill on a shifting performance landscape (Levinthal and Warglien 1999). This ongoing process of mutual self-perturbation results in a cyclic pattern of performance with no systemic sense of progress. This cycling ceases with reintegration, since an integrated, centralized firm will never adopt changes that are harmful to the firm as a whole. As a result, reintegration yields an immediate performance improvement. In addition, as indicated in the case of Reintegrator 1, subsequent local exploration can further improve performance.

The other basic pattern is represented by Reintegrator 2. In this case, the performance cycling ends after a number of periods. However, the cycling can end in a poor compromise, with both divisions at a solution that is optimal given the other division's current actions but with poor performance for the firm as a whole. Divisions have reached a low-performing Nash equilibrium. This stalemate can only be broken by reintegration. With reintegration it is possible to implement decisions in one division that might be harmful to this division, yet greatly benefit the firm as whole. As shown with Reintegrator 2, these improvements may not occur immediately after reintegration, but after a number of periods the newly integrated firm has identified changes in its activity system that enable it to achieve higher performance.

While Figure 5 provides an intuition for why performance increases after reintegration takes place, the interesting puzzle remains why the initial phase of decentralization allows the organization to achieve ultimately an even higher performance than the permanently centralized firm. We postpone the explanation of this effect until the end of the next section, which provides a number of insights into the sources of the beneficial


effect of an initial epoch of decentralization. In short, the initial decentralization allows the firm to escape from poor, yet consistent, sets of choices by moving the firm to a good starting point for the subsequent firmwide local search process.

3.2.3. Search from a Disrupted Global Peak. The analysis to this point has examined the effect of alternative organizational structures on search when this search process commences on an arbitrary starting point on the performance landscape. The typical problem of organizational adaptation is, however, generally couched as a problem of an organization that is initially well positioned in the performance landscape and then faced with the challenge of responding to a shift in the landscape.

To explore this situation, we analyze the following setting, in which each firm's initial activity system is optimized given the existing environment; i.e., firms are located on the global peak of the landscape. Firms then encounter an environmental change that alters the optimal activity configuration; i.e., a shift in the global peak. A minor environmental change might require each firm to change only one of its activities to reascend the global peak. Larger environmental changes may require firms to change more activities. In other words, we model the degree of environmental change by the number of activity changes, d , required by each firm to reascend the global peak. In this light, the random starting point used up to now can be interpreted as a mixture of environmental shocks that dislocated firms from the global peak by an average of $d = 3$ activity choices.⁵

Figure 6 displays the performance in period 100 of permanently centralized firms and firms that reintegrate


Figure 6 Performance in Period 100, Given Different Degrees of Initial Environmental Change on $K = 2$ Landscapes



after period 25. Reported results are averaged over 10,000 landscapes with $K = 2$. Two interesting observations arise from the figure. First, as the degree of environmental change increases, the performance advantage of the reintegrator over the centralized firm increases. Thus, temporary decentralization appears to be particularly important after large environmental shocks. Second, temporary decentralization can have a cost, even in the long run. When the firm is close to, or already on the global peak, temporary decentralization has the potential to lead the firm astray—in fact, so far astray that the firm is unable to reascend the global peak after it reintegrates. Thus, experimentation does not come for free; it is accompanied by the risk of undermining an existing high-performing set of choices.

While the exact level of environmental change at which the performance of the reintegrator exceeds that of the centralized firm depends on the degree of interaction, the overall pattern of performance is similar for different degrees of interactions. In Figure 7, we show the performance in period 100 of centralized firms and firms that reintegrate after period 25 for landscapes with $K = 5$. Again one can observe the potential liability of temporary decentralization for firms that start on the global peak and the increasing benefit of temporary decentralization as the degree of environmental change increases. In sum, for environmental changes that require only a low-dimensional change in the activity system to regain internal fit, a centralized response is an appropriate organizational adaptation strategy. For larger changes, however, a temporary decentralized response tends to yield superior performance.

Figure 7 Performance in Period 100, Given Different Degrees of Initial Environmental Change on $K = 5$ Landscapes



3.2.4. Value of Temporary Decentralization. What are the sources for the performance differences between centralized and reintegrated firms? As shown in Table 1, the performance differences that could be seen in Figures 6 and 7 are driven by two factors: the number of firms that are able to reach the global peak and the performance of firms that do not reach the global peak. The top panel of Table 1 reports these measures for the firms on the $K = 2$ landscapes that were plotted in Figure 6, while the bottom panel reports the measures for the firms on the $K = 5$ landscapes plotted in Figure 7. Consider, for instance, firms that encounter a very large shock ($d = 6$) on very rugged landscapes ($K = 5$). Only 0.91% of the permanently centralized firms are able to reach the global peak, whereas 15.85% of the reintegrators are able to do so. Thus, one source for the performance advantage for reintegrators is their greater ability to find the global peak.

Now consider those firms that did not reach the global peak. The average performance of the centralized firms that did not reach the global peak is, in this case, 0.862, while the average performance of the reintegrators that did not reach the global peak is (a statistically significantly higher) 0.880. Thus, a second source of the performance advantage of reintegrators is their ability to find higher local peaks than the centralized firms.⁶

The table also provides an insight into the potential liability of reintegrators. Centralized firms that are already on the global peak will never leave the global peak, as they cannot find any incremental change that improves performance. Decentralized firms, however, can move away from the global peak. Indeed, we see for the $K = 5$ landscape that, out of the 10,000 reintegrators that start on the global peak, 2,796 do not end up on the global peak. Their initial exploration leads them so far astray from the global peak that they are unable to find their way back onto the global peak once they reintegrate.⁷

To describe these effects more precisely, the concept of a “basin of attraction” is helpful. A peak’s basin of attraction is the set of locations on a landscape from which the particular peak can be reached via local exploration. In other words, a firm that searches locally and that finds itself in a particular basin of attraction tends to gravitate toward a particular activity configuration, i.e., a particular peak in the landscape. Two consequences arise from the firm’s local search. First, a firm’s starting location can have a profound effect on its eventual activity system. This path-dependence is related to Stinchcombe’s (1965) argument regarding imprinting. In contrast to Stinchcombe, however, the history dependence that results from local search through a rugged

Table 1 Sources of the Benefit of Reintegration

	Degree of Environmental Change						
	0	1	2	3	4	5	6
<i>K</i> = 2							
<i>% at global peak</i>							
Reintegrator	77.81	66.41	55.62	41.62	43.37	40.12	35.64
Centralized	100.00	78.26	58.64	48.27	29.46	20.55	13.32
<i>Performance if not at global peak</i>							
Reintegrator	0.927	0.919	0.913	0.909	0.907	0.904	0.902
Centralized	—	0.916	0.908	0.901	0.895	0.890	0.887
<i>K</i> = 5							
<i>% at global peak</i>							
Reintegrator	72.04	35.52	25.02	20.24	18.72	16.27	15.85
Centralized	100.00	48.10	21.50	9.98	4.53	1.75	0.91
<i>Performance if not at global peak</i>							
Reintegrator	0.896	0.884	0.880	0.882	0.882	0.880	0.880
Centralized	—	0.864	0.861	0.860	0.862	0.861	0.862

performance landscape does not require organizations to be inert; rather, the particular local peak the organization reaches is a function of the starting point of the search process (Levinthal 1997). Second, a local peak and its associated basin of attraction fence in firms that pursue local search. If we consider all configurations close to a particular peak as reachable with a firm's existing capability set, then peaks and their associated basins of attraction are the landscape representations of "capability traps" (Levitt and March 1988).

The two sources of the performance advantage of the reintegrator can be rephrased concisely as follows: Temporary decentralization allows the firm to avoid basins of attraction of low peaks. As a result, reintegrators are more likely to end up on the global peak, and, if they do not reach the global peak, they end up on higher local peaks than permanently centralized firms. This insight, to be explained in more detail, solves the puzzle of why reintegrators can achieve higher performance than permanently centralized firms. Reintegrators are able to avoid low local peaks, i.e., avoid competency traps that are associated with low performance.⁸

Why are reintegrators able to avoid low local peaks? Since reintegrators become centralized firms, reintegrators, as well as centralized firms, find themselves by the end of their search process on (local) peaks on the landscape. The last question thus can be rephrased as: Why are reintegrators able to avoid basins of attraction of low local peaks? A firm that starts its search process in a centralized manner is stuck with whatever basin of attraction is associated with its initial location, i.e., its

initial activity system. Similarly, at the time of reintegration, a reintegrator is stuck with whatever basin of attraction it has reached at that time. This starting point for the reintegrator, however, is influenced by the prior phase of local search, which is driven by the divisional performance perspective. This decentralized search can dislodge firms from basins of attraction. Moreover, as shown below, decentralized search is more likely to dislodge a firm from a basin of attraction of a low peak than from a basin of attraction of a high peak. As a result, at the time of reintegration, the reintegrator finds itself, on average, in a basin of attraction that is more attractive than the basin of attraction in which it and the centralized firm originally started their search.

Why is decentralized search more likely to dislodge a firm from a basin of attraction of a low peak than from a basin of attraction of a high peak? Consider a reintegrator that is still in its decentralized phase. As this firm moves through the landscape, it may reach a point that is a local peak, i.e., an activity combination from which a centralized firm would never move. Since in the decentralized firm each division is only concerned with its own profits, it is possible that the decentralized firm will move away from this local peak. In the terminology of Rivkin and Siggelkow (2002, 2003), this local peak may not be a "sticking point" for the decentralized firm. (A "sticking point" is a configuration of choices from which a firm does not move. While local peaks are sticking points for the centralized firm, they may not be sticking points for the decentralized firm.) If this local peak is a not sticking point for the decentralized firm, it is possible for the firm to escape this peak's basin of

attraction. If it were as likely that a decentralized firm would move away from a high local peak as from a low local peak, overall performance would not be affected substantially.⁹ Yet decentralized firms are actually more likely to move away from low local peaks than from high local peaks. In other words, the probability that a local peak is a sticking point is (weakly) increasing in the height of the local peak.

To gain an intuition for this result, recall that the performance of an activity configuration is the sum of the performance contributions of the two divisions, divided by two. For an activity configuration to be a local peak, the sum of the two division contributions has to be larger than the sum of the two division contributions for each of its local alternatives. For instance, on an $N = 2$ landscape, for configuration 00 to be a local peak, its performance has to be larger than the performance of 01 and 10. For an activity configuration to be a sticking point for a decentralized firm, however, each division's contribution has to be larger than the respective divisions' contributions of the local alternatives. Thus, for configuration 00 to be a sticking point for a firm that assigns the first activity to one division and the second activity to the other division, it has to be more profitable for the first division to choose 0 than to choose 1 (given that the other division chooses 0), and it has to be more profitable for the second division to choose 0 than to choose 1 (given that the other division chooses 0).

Now consider a peak that has maximal height, i.e., a peak for which each division's contribution is at its highest possible value. (If division contributions are restricted to the unit interval, this would correspond to the case in which each contribution equals 1 and the height of the peak is 1.)¹⁰ The probability that this highest of all possible peaks is a sticking point is 1, because no division could possibly find a profitable deviation. Now consider a peak that has less than maximal height, say 0.6, with underlying contributions of 0.5 and 0.7 for each division. A neighboring activity configuration may yield contributions of 0.6 and 0.4, for an overall performance value of 0.5. Even though overall performance is smaller, the first division will implement the change, since it improves its performance, and thereby move the firm away from this local peak. Thus, for peaks that have less than maximal height, the probability that the peak is a sticking point has to be less than 1. More generally, as shown in the Appendix, the lower a peak, the more likely it is that one of the neighboring activity configurations has an underlying divisional contribution that is larger, thereby causing the decentralized firm to move.¹¹ As a result, the probability that a local peak is a sticking point for the decentralized firm is increasing

in the height of the local peak. Consequently, decentralized firms are more likely to escape basins of attraction of low local peaks than of high local peaks. Thus, once a firm reintegrates, it finds itself in a basin of attraction that leads, on average, to a peak that is higher than the peak that would be reached via centralized search commencing on the same original starting point.¹²

Put differently, a consistent activity configuration (i.e., a peak) that generates very high overall firm performance is likely to be composed of activity configurations that yield high performance for both divisions (otherwise overall performance would not be very high). Since each division's performance is high, neither division is likely to find a new configuration that would yield higher divisional payoffs, and the firm retains this activity configuration. In contrast, a consistent configuration that creates intermediate performance can be generated by one division with high performance and one division with low performance. In this case, it is likely that the division with low performance can find a configuration for its division that will improve its performance. Since the firm is decentralized, the division will implement the change and the firm will move away from its original configuration. Thus, the firm has become unglued from its mediocre consistent activity configuration. In sum, the decentralized structure makes it more likely that the firm will escape a competency trap that is associated with lower overall performance than with higher overall performance. As a consequence, once the firm reintegrates, it has maneuvered itself into a good starting position for coordinated, centralized search.

3.2.5. Magnitude of the Reintegration Benefit. In §3.2.3 we investigated the value of reintegration as a function of the initial environmental shock that the firm experienced. In this section we illustrate further determinants of the magnitude of the reintegration effect.


The explanation for the performance advantage of reintegrators, as presented in the last section, stressed two factors: the possibility that centralized firms get stuck on low local peaks, and the ability of reintegrators to avoid getting stuck on these peaks. Reintegrators are able to avoid low local peaks because parochial decision making at the level of the divisions may dislodge a firm from a local peak. Note that cross-divisional interdependencies are necessary for this dislodging to take place. If each division's decisions are independent from the other division, each local peak is also a sticking point for the decentralized firm.¹³ As a result, in §3.2.1, we saw that reintegrators did not have any performance benefit on decomposable landscapes when decisions across divisions were independent of each other.

The magnitude of the reintegration benefit is thus likely to be driven by two factors: the probability of a centralized firm of getting stuck on a low local peak and cross-divisional interdependencies that dislodge decentralized firms from low local peaks. Consequently, the benefit of reintegration should be largest for firms that have many cross-departmental interdependencies and that operate on landscapes that have many local peaks. To investigate these effects, we analyze the long-run performance of reintegrators as compared to centralized firms on landscapes that differ in their overall degree of interaction. As the degree of interaction (K) increases, both the number of local peaks and the number of cross-divisional interdependencies increases. For instance, a $K = 2$ landscape has on average about 3.5 local peaks, while a $K = 5$ landscape has on average about 9 local peaks. Similarly, since interactions are determined randomly, some activities on low K landscapes generate only intradivisional interactions. Consider, for instance, the prior example of the $K = 2$ interaction pattern given in Figure 2C. Here, activity a_3 only affects activities a_1 and a_2 ; i.e., activity a_3 generates only intradivisional interactions. In contrast, for a $K = 5$ interaction structure, each activity generates not only intra- but also cross-divisional interactions.

The black line in Figure 8 displays the performance difference in period 100 between a permanently centralized firm and a decentralized firm that reintegrates after period 25 as a function of the overall degree of interaction of the landscapes on which the firms are operating. Consistent with our argument above, we find that the performance benefit of reintegration increases as the number of local peaks and the degree of cross-departmental interdependencies increase. One can further show that these two factors separately increase the benefit of reintegration, even if the overall degree of interaction stays constant (results available from the authors).

3.2.6. The Timing of Reintegration. The previous sections pointed to the long-term performance advantage that can be derived from an initial phase of decentralization, but we have yet to examine how long this initial phase should be. As we observed in Figure 4, performance rises precipitously with reintegration; as a result, there is some cost to persisting in a decentralized structure. Thus, we are faced with the question of how long firms must stay decentralized to realize the benefits of reintegration. Figure 9 sheds light on this question by showing the long-term performance of firms that reintegrate at different times on landscapes that are characterized by full interdependence. Even five periods


Figure 8 Magnitude of the Reintegration Benefit and Optimal Reintegration Period



of decentralized exploration with subsequent integration generates a significantly higher performance than no decentralized exploration at all. A longer phase of decentralized exploration appears to be beneficial, though. Firms that reintegrate after 10 periods outperform firms that reintegrate after five periods. Similar relationships are found for firms reintegrating after 10, 30, and 40 periods. Substantially more than 40 periods of decentralized exploration, however, does not appear to yield any performance advantages. The long-term performance of firms that reintegrate after 50 or 60 periods is not significantly different from firms that reintegrate after 40 periods.

To repeat: Decentralization alone does not yield high performance. Firms that never reintegrate perform sig-

Figure 9 Performance in Period 100 on Fully Interdependent Landscapes



nificantly worse than reintegrators, as shown by the bar on the far right of Figure 9.

As these results show, to achieve high long-term performance, reintegration should not occur too early. At the same time, keeping the decentralized structure for a very long time does not yield any long-term benefits, while it does create costs of underperformance. What affects the optimal timing of reintegration? Similar to our analysis of the magnitude of the reintegration benefit, we analyze the optimal timing of reintegration as a function of the overall degree of interaction, which has an effect on both the ruggedness of the landscape and the degree of cross-divisional interdependence. The white line in Figure 8 displays the optimal period of reintegration as a function of K .¹⁴ As the ruggedness of the landscape and the degree of interdependence between the divisions increase, the optimal length of the initial phase of decentralization increases as well. For instance, while for $K = 2$ landscapes 12 periods of decentralization are sufficient, for $K = 5$ landscapes decentralization up to 40 periods yields long-term performance benefits.

In sum, firms need to trade off the short-term costs of decentralized exploration and the long-term benefits that can be achieved from coupling initial decentralized exploration with subsequent integration. Since the costs of decentralized exploration are caused by divisions that only care about their own payoffs, one may wonder whether a different incentive structure might help alleviate this problem. For instance, rather than making the divisions profit centers, a firm might employ firm-based compensation schemes that reward only decisions that are beneficial for the firm as a whole. In managerial terms, this is a central question: In making decisions, how independent should each division be? Organizationally, firms have two levers: a structural choice of how to allocate decisions (in this case, “Should we allocate decision rights to two separate divisions?”) and a choice of how to guide decision making within each division (in this case, “When deciding what choices to make for each activity, should a division care about the firm as a whole or only about the division itself?”). Can the firm still reap the benefits of its structural choice of allowing decentralized exploration, while reducing the short-term performance costs by requiring divisions to take only decisions that are beneficial for the firm as a whole? In short, the answer appears to be “no.”

In Figure 4, we also report the performance of a second firm that is decentralized until period 25 and integrated thereafter. In this firm, however, divisions use V , rather than V_1 and V_2 , to evaluate alternatives; i.e., divisions take the firm’s overall performance into

account when making decisions concerning new activity configurations. As expected, the short-term costs of the decentralized structure are reduced. Indeed, this firm outperforms the centralized firm in the short-run because of the virtues of parallelism in the search process that decentralization permits. At the same time, however, this firm does not experience any performance boost from later reintegration, in stark contrast to the firm that started out using parochial decentralized exploration. Early decentralization with firmwide incentives simply channels the firm more quickly onto a similar developmental path as the one taken by a firm that is centralized from the beginning—it does not allow the firm to escape from its existing basin of attraction. Formally, with firmwide incentives, each local peak is a sticking point for the decentralized firm (Rivkin and Siggelkow 2002). As a result, no low local peaks are avoided and long-term performance between the centralized and the reintegrated firm are not different from each other. Since no “risks are taken,” i.e., no activity configurations that lead to performance declines in the present period are adopted, organizational decentralization with firmwide incentives ultimately does not generate further exploration than a permanently centralized firm. Only parochial decentralization creates long-term benefits.

3.2.7. Scrambling the Decomposable. As our last consideration, we return to the first simulation setup we analyzed: firms that search landscapes with underlying interaction patterns that are block-diagonal (Figure 2A). As noted before, firms that assign decisions concerning activities a_1 – a_3 and a_4 – a_6 to the two divisions, perfectly decompose this decision problem; i.e., no cross-divisional interdependencies remain. Since the decision problem is perfectly decomposed in these firms, each division faces an independent decision problem. However, within each decision problem, the three decisions still interact richly ($N = 3$, $K = 2$). Since each division’s “subscape” (the partial landscape comprised of the N_i elements under each division’s control) is itself rugged, divisions may not find the highest peak in their subsapes. As a result, as was shown in Figure 3, even firms that perfectly decompose the decision problem may not reach the overall global peak on the landscape.

The main result of this paper states that temporary, cross-divisional interdependencies can be beneficial because they help a firm to avoid low local peaks. The question arises whether it could be beneficial to introduce temporary cross-divisional interdependencies into situations in which such interdependencies could be avoided. In other words, even though a decision


problem might be decomposable given a “proper” decision allocation, might there be a benefit to using a “scrambled” decision allocation that creates cross-divisional interdependencies?

The answer to this question is “yes” and can be derived formally from our main result. In short, one can show that a subset of $K = 2$ interaction structures can be transformed into a block-diagonal interaction structure by reassigning activities between the two managers. Thus, a scrambled decision allocation of a block-diagonal decision problem (e.g., one division has control over activities a_1 , a_4 , and a_5 , while the other division has control over activities a_2 , a_3 , and a_6) has an equivalent counterpart in a decision allocation (a_1 – a_3 ; a_4 – a_6) with an appropriately scrambled $K = 2$ interaction structure. Since our main result holds for all scrambled $K = 2$ interaction structures, the result transfers over to the case of a scrambled decision allocation over a block-diagonal interaction structure. (A more formal exposition of this point is available from the authors.)

For illustration, consider a firm that operates on landscapes with block-diagonal interactions and that scrambles its decision allocation (one division has control over activities a_1 , a_4 , and a_5 , while the other division has control over activities a_2 , a_3 , and a_6) until period 25, at which point it switches to a decision allocation that decomposes the decision problem (one division controls activities a_1 – a_3 , the other a_4 – a_6). As a benchmark, Figure 10 contains the performance of the permanently decomposed (i.e., in this case, decentralized) firm, the firm with the highest performance in this landscape, as previously reported in Figure 3. As seen in Figure 10, in the long run, the temporarily scrambled firm is able to outperform the perfectly decomposed firm. Similar to the prior results concerning decentralization, we can also observe that the benefits of cross-divisional interdependencies are reaped only if the firm eventually adopts the “proper” organizational structure. Firms that remain scrambled on decomposable landscapes suffer the same problems (of cycling or low sticking points) as firms that stay permanently decentralized on nondecomposable landscapes.

The intuition for this result is similar as before: Without cross-interdependencies, divisions will settle on the first consistent set of choices they find. However, if the problem each division is solving is complex, it is likely that the first solution (the first local peak) each division finds may not be a very good one. If search is local, a division will find no improvement to this first solution and will be stuck with it. Cross-divisional interdependencies can dislodge a division from such a solution,

Figure 10 Performance on Decomposable Landscapes



and in particular from solutions that yield overall low performance.

3.2.8. Overview of Simulation Results. Overall, the simulation results point to a number of interesting findings. First, even if decisions between divisions interact, a temporary divisionalization can be beneficial in the long run. In other words, even if the decision problem the firm faces is not fully decomposable, a temporary bifurcation can ultimately lead to a higher performance outcome. In the short run, however, if interactions are dense, decentralization leads to lower performance than an integrated form. As a result, a tradeoff exists between the short-term costs of parochial, decentralized exploration and the long-term benefits of reaching higher performance. As interactions across and within divisions increase, the optimal length of decentralized exploration tends to grow.

Second, the benefit of temporary decentralization is greater the more intra- and cross-divisional interdependencies exist. Temporary cross-divisional interdependencies are important, since they are able to divert firms from low local peaks that are created by both types of interdependencies and that can create capability traps for centralized firms. Moreover, to be able to escape the basins of attraction of low local peaks, divisions must be allowed to pursue parochial interests. The benefits of reintegration also tend to increase with the degree of environmental change that a well-positioned firm encounters. For small environmental changes, temporary decentralization can lead to excess exploration and performance declines.

Third, even if a decision problem is decomposable, it can be beneficial to create a temporary decision allocation that creates “unnecessary” interdependencies across divisions. This is especially the case when the decomposed subproblems are themselves complex. Similar to the decentralization results, temporary interdependencies can prevent divisions from getting stuck on low-performing solutions.

4. Discussion

The findings of the last section have a range of implications for the appropriate structuring of search processes. In short, if one can afford it, ignoring interdependencies for a while can be beneficial for long-term performance. Likewise, creating temporary interdependencies, even though they could be avoided, can lead to more valuable long-term outcomes. Given the general nature of our model, these findings are applicable to various organizational levels. So far, we considered the allocation of decision rights over activities to different divisions of a firm, e.g., the Web division and the mother firm. An equivalent problem at a lower level of analysis would be posed by the design of a complex product. In this case, each choice element a_i would correspond to a design parameter, where, for example, the first three parameters describe the first component and the second three parameters describe the second component. Given a particular interaction structure among these parameters, what is an appropriate search strategy? Should each component be developed by an independent design team? How often and when should design teams meet to take into account their interdependencies? At a higher level of analysis, each division may correspond to a separate subsidiary of a multinational firm. Here, the issue can be framed as to how much autonomy to grant each subsidiary versus how much control to exert centrally. Last, each division may correspond to individual firms that may pose the question of how much value could be won by “centralizing” their decision making through joint ventures or mergers.

In this section, we will touch on a number—yet not all—of these organizational issues. Continuing the example from the Introduction, we start by discussing the implications with respect to firms’ organizational responses to the Internet. We then broaden the discussion to address the general issue of decentralization within organizational structures, in particular in environments that may experience more than one shock. Last, we apply our findings to a lower level of analysis and discuss potential implications for product design processes.

The advent of the Internet, which disrupted the existing performance landscape for many firms, provides a powerful natural experiment with which to consider the findings of our analysis. It appears that the Internet is likely to have a pervasive effect on a firm’s business system, leading to a nondecomposable decision problem. First, the Internet has many aspects of a general purpose technology (Bresnahan and Trajtenberg 1995, David 1990). As a result, it is unlikely that all necessary activity reconfigurations will be contained (or be containable) in a separate organizational structure. Consequently, changes in the Internet division often have reverberations throughout the firm. Second, external stakeholders, such as consumers or distributors, frequently induce interdependencies between activity choices in the Web division and mother firm, because both often operate in the same competitive context. Common examples of such interdependencies arise, for example, with respect to pricing, distribution, and human resource practices.

For instance, with respect to pricing, the Web division of a firm will have to face the question of whether to match the prices of other, potentially pure-Internet, competitors, or whether to keep price in parity with the mother firm. Should the prices offered, for example, by *barnesandnoble.com* be competitive with *Amazon.com*, or should they be the same as those offered in Barnes&Noble stores? This decision of the Web division affects the payoffs of other decisions taken by Barnes&Noble stores, e.g., the effectiveness of price discounts, as Barnes&Noble competes with Borders stores. In short, decision a_i of the Web division affects payoffs of decision a_j in the mother firm.

The degree of interdependencies between a Web division and its mother firm are thus likely to be high, suggesting that for firms that encountered a large change in their performance landscape, a temporary decentralized response might be more appropriate than either a fully centralized response or a long-term decentralized structure. Interestingly, this pattern of reintegration has been observed at a number of firms that first created an independent Internet division, such as Disney and Bank One (Kherrmouch and Byrnes 2001). Similarly, Gilbert (2003) finds that even though there are important operating connections between newspapers’ online units and their traditional press units, there was a need to allow each unit autonomy to first discover the broad outline of the appropriate business models for each, before the units were reintegrated.

It does not appear that many of these reintegrating firms, however, chose a decentralized structure with the intent of reintegration. Rather, many firms appear to have chosen an initial decentralized structure in order

to respond rapidly to the perceived dramatic changes in their performance landscape. Reintegration appears to stem from an effort to resolve many of the coordination challenges that this initial decentralized structure created (recall the cycling behavior of decentralized firms in Figure 5). Whatever the organization's intent or vision at the time of its initial response to the Internet, our analysis suggests that this dynamic of an initial epoch of decentralization followed by reintegration is in fact a powerful structuring of the search process when interdependencies across divisions exist.

More generally, our findings indicate that firms may change their organizational structure, or the balance between centralization and decentralization, both immediately after environmental changes (towards decentralization) and after a period of initial exploration, back to a more centralized structure. One should note that the second organizational change (reintegration) takes place even though the environment is not changing. Thus, environmental shocks and reorganizations may not always go hand in hand (Siggelkow and Levinthal 2003).

Since our model subjected firms to a one-time shock, firms engaged only in one cycle of decentralization and centralization. In environments that experience a series of shocks over time, however, firms may have to repeat this process. As a result, firms might cycle through different organizational structures, pulsating back and forth between decentralization, to ignite new search, and centralization, to increase coordination. Indeed, as Schoonhoven and Jelinek report (1990, p. 99), firms that operate in dynamic environments "adapt to change by reorganizing their formal structures." Moreover, these firms switch their organizational structures to both maintain their flexibility and their ability to remain efficient producers (*ibid.*). A similar observation has been offered by Bartlett and Ghoshal (1990) in the context of managing innovation in multinational companies. To manage the cross-divisional interdependencies within a multinational company, Bartlett and Ghoshal (1990, p. 237) noted that Ericsson had created a "constant ebb and flow in the centralization and decentralization of various responsibilities."

At a lower level of analysis, our model has a number of interesting implications for product design. For design problems with many interdependencies, our findings would suggest breaking the problem into parts and ignoring interdependencies for an initial phase. In this phase, each team should be given free reign and exhorted to find the best solution for its own subproblem. By allowing this (costly) initial phase of unfettered experimentation, premature lock-in into a suboptimal design

can be avoided. Our model also speaks to the appropriate timing of the modularization of a system. Most of the literature on product design suggests to modularize, i.e., to decompose a system as much as possible (e.g., Baldwin and Clark 2000, Eppinger et al. 1994). The advantage of decomposition or modularization is clear: Each subsystem can be developed independently without any need for coordination. Yet our analysis points out that if each subsystem is by itself complex, early modularization may yield quite suboptimal solutions within each module. As a result, an early phase in which the product design is split into interdependent subunits (even though a perfect decomposition is possible) can yield eventually higher-performing modules. As in the case of organizational structure, the benefits of the initial exploratory phase will only be reaped, however, after the correct decomposition (in the post-exploratory phase) takes place. Keeping a decomposable system partitioned in a manner that creates interactions across subsystems would create a very unstable system with low performance in the long run.

5. Conclusion

Prior research has often implied that if a system is not decomposable, it should not be split apart, since interaction effects would be ignored. Indeed, we find that if attention is restricted to organizational structures conforming to the pure forms of permanently centralized or decentralized structures, then for intermediate degrees of cross-divisional interdependencies the integrated form is superior (Figure 4). However, a third solution exists: a temporary decentralization followed by reintegration. As our results show, this temporal sequencing of different organizational structures can lead to higher performance than either pure form. Thus, even if a system is not decomposable, one still may usefully break it apart—temporarily.

A second central doctrine of organizational design states that tasks should be grouped so that the most intensive interactions are internalized (Thompson 1967). As Simon put it (1973, p. 270): "Any division of labor among decisional subsystems creates externalities, which arise out of the interdependence among the subsystems that are ignored. What is wanted is a factorization that minimizes these externalities." Our results show a limitation of this dictum as well. If the subsystems, or modules, that decomposition creates are complex by themselves, temporary cross-interdependence among modules can yield higher final performance than perfect decomposition from the start.

The underlying driver for these results is similar in both cases: Cross-interdependencies between divisions

lead to increased exploration because divisions will not get stuck with the first set of consistent choices they chance upon. Subsequent centralization (in the case of a nondecomposable system) or decomposition (in the case of a fully decomposable system) allows the firm to further refine the solutions and to coordinate across the divisions. In sum, the benefit of temporary decentralization arises from its ability to sufficiently dislodge a firm from its current set of practices so that the firm escapes its existing developmental trajectory.

While for minor environmental changes such radical exploration can be detrimental, for larger environmental shocks that require higher-dimensional changes in a firm's activity system, such initial exploration can have long-term benefits. To reap these benefits, however, a mechanism for stability and coordination is required. As Rivkin and Siggelkow (2003) point out, only firms that employ organizational features that both push the firm toward exploration and pull it toward stability tend to have high performance. In our case, exploration and stability are not achieved simultaneously through distinct organizational features as in their model, but sequentially by adopting different organizational structures.

Reconfiguring an organization to correspond more closely to a changed environment is a longstanding practical challenge and one that has been salient in the management literature for some time. However, this work has tended to have a comparative static quality—the environment has shifted from E to E' and therefore the organizational structure should shift from S to S' . Organizational structure is, however, not simply an attempt to capture an isomorphism between an interaction structure in the task environment and the organization. Organizational structure is also an important determinant of the search process—a process that is usefully guided by an integrated, holistic view of the organization, but a process that at times is usefully driven by narrower and conflicting perspectives on what constitutes progress.

Acknowledgments

The authors would like to thank seminar participants at the 2002 Innovation and Inertia Conference, Duke University, and three anonymous reviewers for helpful comments. The authors gratefully acknowledge financial support from the Reginald H. Jones Center for Management Strategy, Policy and Organization and the Sol C. Snider Entrepreneurial Research Center at the University of Pennsylvania.

Appendix

In this Appendix, we formally show for the $N = 2$, $K = 1$ case that the probability that a local peak is a sticking point for a decentralized firm is (weakly) increasing in the height of the local peak. (For higher levels of N and K , we have confirmed this relationship numerically.)

Without loss of generality, assume 00 is a local peak. Let us denote the underlying contributions of 00 with a and b , the underlying contribu-

tions of 10 with c and d , and the underlying contributions of 01 with e and f . Since point 00 is a local peak, it follows that $(a + b) > (c + d)$ and $(a + b) > (e + f)$. The point 00 is a sticking point if $a > c$ and $b > f$. The claim is that $\text{Prob}(a > c, b > f \mid (a + b) > (c + d), (a + b) > (e + f))$ is (weakly) increasing in $(a + b)$. Recall, all contributions are independent draws from a uniform distribution over the unit interval.

Let $s = a + b$. It turns out to be helpful to distinguish between the cases $0 \leq s \leq 1$ and $1 \leq s \leq 2$.

Case 1. $0 \leq s \leq 1$.

In this case, the probability density function (pdf) for x given that $x + y \leq s$, where x and y are independent draws from a uniform distribution over $[0, 1]$ is given by

$$\frac{2}{s} - \frac{2x}{s^2} \quad \text{if } x \leq s, \quad 0 \quad \text{otherwise.}$$

Since $a + b = s$, contribution a is uniformly distributed between 0 and s , with pdf of $1/s$. Hence,

$$\begin{aligned} \text{Prob}(a > c, b > f \mid (a + b) > (c + d), (a + b) > (e + f)) \\ &= \text{Prob}(a > c, (s - a) > f \mid s > (c + d), s > (e + f)) \\ &= \int_0^s \left(\int_0^s \left(\frac{2}{s} - \frac{2c}{s^2} \right) dc \int_0^{s-a} \left(\frac{2}{s} - \frac{2f}{s^2} \right) df \right) \frac{1}{s} da \\ &= \int_0^s \left(\frac{2a}{s} - \frac{a^2}{s^2} \right) \left(1 - \frac{a^2}{s^2} \right) \frac{1}{s} da = \frac{11}{30}. \end{aligned}$$

This result implies that for low local peaks, the probability that a local peak is a sticking point is actually independent of its height.

Case 2. $1 \leq s \leq 2$.


In this case, the probability density function for x given that $x + y \leq s$ is given by

$$\begin{aligned} \frac{2}{4s - s^2 - 2} \quad &\text{if } x \leq s - 1 \\ \frac{2s - 2x}{4s - s^2 - 2} \quad &\text{if } s - 1 \leq x \leq 1. \end{aligned}$$

Since $a + b = s > 1$, contribution a is uniformly distributed between $(s - 1)$ and 1 with pdf $1/(2 - s)$. With the above pdfs, we can write

$$\begin{aligned} \text{Prob}(a > c, (s - a) > f \mid s > (c + d), s > (e + f)) \\ &= \int_{s-1}^1 \left[\left(\int_0^{s-1} \frac{2}{4s - s^2 - 2} dc + \int_{s-1}^a \frac{2s - 2c}{4s - s^2 - 2} dc \right) \right. \\ &\quad \left. \cdot \left(\int_0^{s-1} \frac{2}{4s - s^2 - 2} df + \int_{s-1}^{s-a} \frac{2s - 2f}{4s - s^2 - 2} df \right) \right] \frac{1}{2 - s} da \\ &= \frac{56 - 192s + 184s^2 - 38s^3 + s^4}{30(4s - s^2 - 2)^2} = \frac{56 - 384h + 736h^2 - 304h^3 + 16h^4}{30(8h - 4h^2 - 2)^2}. \end{aligned}$$

We denote the height of a peak with h . Because the height of a peak is the average of the contributions, we substitute $s = 2h$ in the last equality to get a description of the relationship between the height of a local peak and its likelihood of being a sticking point. Figure A1 displays the relationship for the full range of h . As the graph indicates, the higher the local peak, the more likely that it is also a sticking point for the decentralized firm. Put conversely, decentralized firms are more likely to move away from a low local peak than from a high local peak.

Figure A1 The Relationship Between the Height of a Local Peak and Its Probability of Being a Sticking Point

Endnotes

¹Denote with $P^* = (a_1^*, a_2^*, \dots, a_N^*)$ the optimal set of choices, i.e., the global maximum or the highest peak in the landscape. Now consider any set of choices other than the global peak, e.g., $P = (a_1, a_2, \dots, a_N)$. For independent choices, this set cannot be a peak because in each dimension i , choice a_i^* is optimal regardless of all other choices. Hence, for instance, $Q = (a_1^*, a_2, \dots, a_N)$ is a profitable incremental improvement over P . Since this argument applies to all $P \neq P^*$, it follows that the landscape has only one peak (i.e., P^*). Also note that given any starting set of choices, the firm can always find incremental improvements (changing successively each choice to the optimal a_i^*) until it reaches the global peak. In this sense, the landscape is smooth, because from each starting point a firm can find the global peak through a series of incremental improvements.

²The speed advantage of the decentralized firm, and hence the exact degree of this speed advantage, does not affect, however, the main results of this paper, which are concerned with long run performance.

³Note this figure contains one additional firm, the reintegrator with firmwide incentives, which will be discussed later.

⁴The long-run performance advantage of the centralized firm can diminish, however, when the degree of interaction is high. In the case of $K = 5$, the long-run performance of the permanently centralized and permanently decentralized firms are very similar (see Figure 9). In short, the centralized firm, while benefiting from coordination, suffers in the long run from a lack of search. Compared to the decentralized firm, it locks in relatively quickly on a set of activities before having explored much of the performance landscape. In contrast, the decentralized firm explores more of the landscape, yet may have difficulty in stabilizing around a good set of decisions, once found. When $K = 5$, the costs and benefits of stability for the centralized firm are equal and both firms have similar long-run average performance, though their behavior is very different. For a fuller discussion of the tension between search and stability, see Rivkin and Siggelkow (2003).

⁵With a random starting point, in expectation 1 out of 64 firms will start on the global peak, 6 firms will have one different activity than the global peak, 15 firms will have two different activities, 20 three different activities, 15 four different activities, 6 five different activities, and 1 firm six different activities.

⁶As an aside, one may note that in the case of $K = 5$ the performance of centralized firms that do not reach the global peak is (statistically) the same regardless of the distance between the firms' starting point and the global peak. Since performance values are completely uncorrelated in a $K = 5$ landscape, the distribution of local peaks with different heights is unrelated to the starting point's distance to the global peak. As a result, for centralized local search on completely random landscapes, the only driver for performance differences is the proportion of firms that reach the global peak. In contrast, for correlated landscapes, e.g., $K = 2$, the performance of firms that do not reach the global peak is negatively correlated with the distance of the starting point from the global peak.

⁷For firms that start one step away from the global peak, even centralized firms may not end up on the global peak. This effect arises because centralized firms only evaluate one randomly chosen local alternative (rather than all local alternatives).

⁸If the firm is "trapped" on the global peak, a capability trap can be helpful. The liability of decentralization thus can be described as the possibility that decentralization leads the firm to drift out of the basin of attraction of the global peak, even if the firm starts on the global peak.

⁹Even with equal probabilities of moving away from a low or a high local peak, a small performance effect is still possible since basins of attraction of low local peaks tend to be smaller than those of high local peaks. Consequently, it would be more likely that a single step away from a local peak would allow a firm to escape the basin of attraction of a low local peak than of a high local peak.

¹⁰Please note that in this paragraph we are referring to true contribution values, not contribution values that are scaled by the height of the global peak. Thus, not every global peak has the absolute height of 1.

¹¹As shown in the Appendix, in the case of $N = 2$, $K = 1$, the positive relationship between the height of local peaks and their probability of being a sticking point actually holds only for peaks with a height above 0.5. For lower peaks, the probability of being a sticking point is independent of the height of the peak.

¹²Decentralized search followed by centralized search thus produces an effect similar to simulated annealing (Carley and Svoboda 1996). In simulated annealing, organizations are exogenously perturbed with a probability that is decreasing in the performance of the firm. As a result, simulated annealing can also take an organization from one basin of attraction to another.

¹³The proof follows by contradiction. Let activity configuration L , composed of division choices A and B , be a local peak. Since activities across divisions are independent, the value of A is independent of the particular choices B , and vice versa. As a result, $V(L) = V_1(A) + V_2(B)$. Assume L is not a sticking point, i.e., that there exists a local alternative A' with $V_1(A') > V_1(A)$. In this case, however, $V(L') = V_1(A') + V_2(B) > V(L)$, contradicting the assumption that L is a local peak, since L' is a local alternative to L with higher performance. Hence, every local peak is a sticking point when activities are independent.

¹⁴Note that in showing the "optimal" reintegration period, we are only addressing long-term performance. Thus, the "optimal" reintegration point is taken to be the time period after which there is no

further increase in the long-term performance. The truly optimal reintegration period would have to be found by balancing the discounted short-term costs with the discounted long-term benefits.

References

- Baldwin, C. Y., K. B. Clark. 2000. *Design Rules: The Power of Modularity*. MIT Press, Cambridge, MA.
- Bartlett, C. A., S. Ghoshal. 1990. Managing innovation in the transnational corporation. C. A. Bartlett, Y. Doz, G. Hedlund, eds. *Managing the Global Firm*. Routledge, London, U.K.
- Bresnahan, T. F., M. Trajtenberg. 1995. General purpose technologies: "Engines of growth?" *J. Econometrics* **65** 83–108.
- 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.
- Burgelman, R. A. 1991. Intraorganizational ecology of strategy making and organizational adaptation: Theory and field research. *Organ. Sci.* **2** 239–262.
- Carley, K. M., D. M. Svoboda. 1996. Modeling organizational adaptation as a simulated annealing process. *Sociological Methods Res.* **25** 138–168.
- Chandler Jr., A. D., 1962. *Strategy and Structure: Chapters in the History of Industrial Enterprise*. MIT Press, Cambridge, MA.
- Christensen, C. M. 1997. *The Innovator's Dilemma*. Harvard Business School Press, Boston, MA.
- David, P. A. 1990. The dynamo and the computer: An historical perspective on the modern productivity paradox. *Amer. Econom. Rev.* **80** 355–361.
- Drazin, R., A. H. Van de Ven. 1985. Alternative forms of fit in contingency theory. *Admin. Sci. Quart.* **30** 514–539.
- Eisenhardt, K. M., S. L. Brown. 1999. Patching: Restitching business portfolios in dynamic markets. *Harvard Bus. Rev.* **77**(May–June) 72–82.
- 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.
- Ethiraj, S., D. Levinthal. Modularity and innovation in complex systems. *Management Sci.* Forthcoming.
- Galunic, D. C., K. M. Eisenhardt. 2001. Architectural innovation and modular corporate forms. *Acad. Management J.* **44** 1229–1250.
- Gavetti, G., D. Levinthal. 2000. Looking forward and looking backward: Cognitive and experiential search. *Admin. Sci. Quart.* **45** 113–137.
- Gilbert, C. G. 2003. Can competing frames co-exist? The paradox of threatened response. Working paper, Harvard Business School, Boston, MA.
- Holland, J. H. 1975. *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence*. University of Michigan Press, Ann Arbor, MI.
- Kauffman, S. A. 1993. *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press, New York.
- Khermouch, G., N. Byrnes. 2001. Come back to papa. *Bus. Week* **73**(February 19) 42.
- Learned, E. P., C. R. Christensen, K. R. Andrews, W. D. Guth. 1961. *Business Policy: Text and Cases*. Irwin, Homewood, IL.
- 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.
- Levitt, B., J. G. March. 1988. Organizational learning. *Annual Rev. Soc.* **14** 319–340.
- March, J. G. 1991. Exploration and exploitation in organizational learning. *Organ. Sci.* **2** 71–87.
- Meyer, A. D., A. S. Tsui, C. R. Hinnings. 1993. Configurational approaches to organizational analysis. *Acad. Management J.* **36** 1175–1195.
- Milgrom, P. R., J. Roberts. 1990. The economics of modern manufacturing: Technology, strategy, and organization. *Amer. Econom. Rev.* **80** 511–528.
- Miller, D. 1986. Configurations of strategy and structure: Towards a synthesis. *Strategic Management J.* **7** 233–249.
- , 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.
- Porter, M. E. 1996. What is strategy? *Harvard Bus. Rev.* **74**(6) 61–78.
- , N. Siggelkow. 2002. Contextual interactions within activity systems and sustainable competitive advantage. Working paper, Wharton School, University of Pennsylvania, Philadelphia, PA.
- 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.
- , ———. 2003. Balancing search and stability: Interdependencies among elements of organizational design. *Management Sci.* **49** 290–311.
- Sanchez, R. 1995. Strategic flexibility in product competition. *Strategic Management J.* **16**(Summer) 135–159.
- , J. T. Mahoney. 1996. Modularity, flexibility, and knowledge management in product and organization design. *Strategic Management J.* **17**(Winter Special Issue) 63–76.
- Schilling, M. A. 2000. Towards a general modular systems theory and its application to interfirm product modularity. *Acad. Management Rev.* **25** 312–334.
- , H. K. Steensma. 2001. The use of modular organizational forms: An industry-level analysis. *Acad. Management J.* **44** 1149–1168.
- Schoonhoven, C. B., M. Jelinek. 1990. Dynamic tension in innovative, high technology firms: Managing rapid technological change through organizational structure. M. A. von Glinow, S. A. Mohrman, eds. *Managing Complexity in High Technology Organizations*. Oxford University Press, New York.
- 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. Evolution toward fit. *Admin. Sci. Quart.* **47** 125–159.
- , D. A. Levinthal. 2003. Adapting to adapt: Sequencing organizational structures. Working paper, Wharton School, University of Pennsylvania, Philadelphia, PA.
- Simon, H. A. 1962. The architecture of complexity. *Proc. Amer. Philos. Soc.* **106** 467–482.
- . 1973. Applying information technology to organization design. *Public Administration Rev.* **33** 268–278.
- Sorenson, O. 2002. Interorganizational complexity and computation. J. A. C. Baum, ed. *Companion to Organizations*. Blackwell, Oxford, U.K.
- Stinchcombe, A. L. 1965. Social structure and organization. J. G. March, ed. *Handbook of Organizations*. Rand McNally, Chicago, IL.
- 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.
- Whittington, R., A. Pettigrew, S. Peck, E. Fenton, M. Conyon. 1999. Change and complementarities in the new competitive landscape: A European panel study, 1992–1996. *Organ. Sci.* **10** 583–600.