MANAGEMENT SCIENCE
Vol. 53, No. 7, July 2007, pp. 1068–1085
ISSN 0025-1909 | EISSN 1526-5501 | 07 | 5307 | 1068

DOI 10.1287/mnsc.1060.0626

Patterned Interactions in Complex Systems: Implications for Exploration

Jan W. Rivkin
Harvard Business School, Harvard University, 239 Morgan Hall, Boston, Massachusetts 02163,
jrivkin@hbs.edu

Nicolaj Siggelkow
The Wharton School, University of Pennsylvania, 2211 Steinberg Hall-Dietrich Hall,
Philadelphia, Pennsylvania 19104, siggelkow@wharton.upenn.edu

Scholars who view organizational, social, and technological systems as sets of interdependent decisions have increasingly used simulation models from the biological and physical sciences to examine system behavior. These models shed light on an enduring managerial question: How much exploration is necessary to discover a good configuration of decisions? The models suggest that, as interactions across decisions intensify and local optima proliferate, broader exploration is required. The models typically assume, however, that the interactions among decisions are distributed randomly. Contrary to this assumption, recent empirical studies of real organizational, social, and technological systems show that interactions among decisions are highly patterned. Patterns such as centralization, small-world connections, power-law distributions, hierarchy, and preferential attachment are common. We embed such patterns into an NK simulation model and obtain dramatic results: *Holding fixed the total number of interactions among decisions, a shift in the pattern of interaction can alter the number of local optima by more than an order of magnitude*. Thus, the long-run value of broader exploration is significantly greater in the face of some interaction patterns than in the face of others. We develop simple, intuitive rules of thumb that allow a decision maker to examine two interaction patterns and determine which warrants greater investment in broad exploration. We also find that, holding fixed the interaction pattern, an increase in the number of interactions raises the number of local optima regardless of the pattern. This validates prior comparative static results with respect to the number of interactions, but highlights an important implicit assumption in earlier work—that the underlying interaction pattern remains constant as interactions become more numerous.

Key words: exploration; simulation model; pattern of interaction
History: Accepted by Brian Uzzi and Luis Amaral, special issue editors; received August 31, 2004. This paper was with the authors $1\frac{1}{2}$ months for 2 revisions.

1. Introduction

How much should an organization invest in the broad exploration of new possibilities? This enduring question arises in a wide array of contexts, including the management of production processes (Abernathy 1978), the search for new technologies (Wheelwright and Clark 1992, Fleming 2001), the structuring of organizations (Tushman and O'Reilly 1996), the design of products (Ulrich and Eppinger 2007, Baldwin and Clark 2000), and the design of individual and organizational learning processes (Ashby 1960, Argyris and Schön 1978, March 1991). The question poses a managerial dilemma. On one hand, managers of an organization must embrace the exploration of new possibilities. Otherwise, the organization fails to innovate. On the other hand, managers must contain exploration because it competes for resources with another crucial organizational process, the exploitation of known opportunities (March 1991). It is widely acknowledged that effective organizations strike a healthy balance between exploration and exploitation, even though it is organizationally difficult to accomplish both (Ghemawat and Ricart i Costa 1993, Tushman and O'Reilly 1996, Benner and Tushman 2003). How, however, can one know whether a particular balance is healthy? Under which conditions is it essential to rein in exploration, and when must one unleash it?

Studies of complex adaptive systems, set initially in the physical and biological sciences, have begun to shed light on this issue. Many of these studies seek systems that relax the exploration/exploitation trade-off—systems that are responsive and creative, yet stable and orderly, neither frozen nor chaotic (e.g., Langton 1990, Kauffman 1993). Among the complex adaptive systems frameworks that have made the transition to management science, the NK model from theoretical biology (Kauffman and Levin 1987, Kauffman and Weinberger 1989, Kauffman 1993) has become a particularly popular platform for studying organizations as complex adaptive systems (e.g., Levinthal 1997, McKelvey 1999, Gavetti and Levinthal 2000, Rivkin 2000, Sorenson 2002, Ethiraj and Levinthal 2004). The model grants a researcher control over the interactions among the elements that make up a system. Results of the model have shed light on the question of optimal exploration: As the degree of interaction among a firm's choices rises, the poor local optima that can disrupt a firm's search efforts proliferate and it becomes preferable, ceteris paribus, for a firm to undertake more exploration in order to escape those optima (Kauffman 1993, Levinthal 1997, Rivkin and Siggelkow 2003). Decision-making processes that focus on incremental change run out of improvement possibilities quickly when choices are intertwined, so broader search becomes vital.

By embedding recent empirical results in a simulation model, this paper takes the NK model's insights on optimal exploration an important step further. Past modeling efforts have looked exclusively at how the *degree* of interaction among a firm's choices affects appropriate exploration. Much less attention has been placed on the *pattern* of interaction among these choices. Indeed, in most NK analyses it is assumed that interactions among choices have a random pattern. This made sense in the biological context, where the interactions were among genes and it was "useful to confess our total ignorance and admit that, for different genes and those which epistatically affect them, essentially arbitrary interactions are possible" (Kauffman 1993, p. 41). In the context of organizational, social, and technological systems, however, recent empirical work has shown that interactions are often very patterned. Our paper exploits this newly gained knowledge in three ways. First, we examine how commonly observed patterns of interactions affect the proliferation of local optima and, accordingly, the appropriate amount of long-run exploration. We find that systems of choices with the same number of total interactions but different patterns of interactions can display very different numbers of local peaks. Second, we study the relationship between the number of interactions and the number of local peaks, holding the pattern of interaction constant. This analysis sheds light on the question of whether prior comparative static results with respect to K, derived from random interaction patterns, are likely to hold for other patterns as well. Third, we identify easily observable characteristics of interaction patterns beyond the overall degree of interaction that in many cases allow one to look at two patterns of interaction and tell immediately which one generates more local optima and warrants greater investment in broad exploration. This can enable managers to convert their knowledge of the interactions among the choices they face into concrete guidance for optimal exploration.

For insight into real patterns of interactions, we rely on empirical work conducted in diverse domains. Detailed work at the level of individual firms (e.g., Porter 1996, Siggelkow 2002), and at the level of individual product systems (e.g., Eppinger et al. 1994, Ulrich and Eppinger 2007, Baldwin and Clark 2000), has yielded a number of explicit maps that show the interdependencies among the various system elements, allowing us to start seeing patterns. Likewise, recent network analyses, such as work on small-world networks (Watts and Strogatz 1998), has generated a great deal of research describing the patterns of real-world networks of interactions. As most of these studies show, networks tend not to be random, but are highly patterned.[1] Specifically, recent empirical work led us to study 10 different interaction patterns: a *small-world interaction* structure (Watts and Strogatz 1998), which includes as extreme cases the *random* structure and the *local* structure; the *preferential attachment* and the *scale-free* structures, two structures currently under intense investigation (e.g., Barabási 2002); and the *centralized*, *hierarchical*, *block-diagonal*, *diagonal*, and *dependent* structures, which capture various patterns observed in product design and studies of firms.

We emphasize the implications of interaction patterns for optimal exploration. Prior research has shown that interaction patterns affect other organizational phenomena as well, including the ability of a firm to adapt to environmental change, to imitate the effective configurations of other firms, and to replicate one's own effective configurations (e.g., Levinthal 1997; Rivkin 2000, 2001). We speculate below on how interaction patterns may influence these phenomena. Moreover, firms might be able to affect interaction patterns through system design decisions (Levinthal and Warglien 1999, MacCormack et al. 2006). Our findings suggest how firms might design systems to be more readily searchable.

This paper is structured as follows: Section 2 describes in detail the 10 interaction patterns we analyze. Section 3 outlines how we create decision problems with these different underlying interaction patterns. The results in §4 characterize the local

[1] It is interesting to note that, similar to the NK framework, network and graph theory, building on the seminal work by Erdős and Rényi (1959), traditionally relied on a randomness assumption as well. As Barabási (2002, p. 23) points out, "The random network theory of Erdős and Rényi has dominated scientific thinking about networks since its introduction in 1959. It created several paradigms that are consciously or unconsciously imprinted on the minds of everyone who deals with networks. It equated complexity with randomness. If a network was too complex to be captured in simple terms, it urged us to describe it as random."

optima that arise from the various interaction patterns. Section 5 explains in an intuitive way the link between different interaction patterns and the number of local optima they create. The different numbers of local optima, in turn, affect the benefit of broad organizational exploration, as §6 shows. Section 7 concludes.

2. Types of Influence Matrices

While the model we study is general enough to encompass a wide range of organizational, social, and technological systems, for expositional purposes we focus on firms as our system of interest. A long tradition in the organization literature (e.g., Learned et al. 1961), reinforced recently by empirical, prescriptive, and computational studies (e.g., Siggelkow 2002, Porter 1996, Levinthal 1997), leads us to conceptualize a firm's management team as facing a number of interdependent decisions. Each firm must choose, for instance, whether to distribute broadly or through narrow channels, whether to advertise in mass media, whether to invest in large-scale production facilities, etc. These decisions might interact with each other. For instance, broad distribution might make mass-market advertising more attractive.

In the context of modeling search behavior of firms, the NK framework assumes that a firm faces N decisions, each of which can be configured in a number of different ways (two, in our simulations). The contribution of an individual decision to a firm's overall performance depends on the resolution of that decision, and possibly other decisions. It is common to think of the space of decisions and the payoffs from combinations of choices as defining a "performance landscape:" each of the N decisions corresponds to a "horizontal" dimension, while the payoff is represented on the "vertical" axis.

An *influence matrix* records which decisions affect each decision. If a firm makes N decisions, then an influence matrix is an $N*N$ matrix whose entry (i, j) is set to an x if the resolution of column decision j affects the value of row decision i. Because each decision affects itself, all influence matrices have xs along their diagonal. Influence matrices can differ, however, in the total number of off-diagonal xs, i.e., in the number of interactions among the decisions, and in the patterns of these interactions. In the original NK setup (Kauffman 1993), it was assumed that each decision is affected by exactly K other decisions, i.e., each row contained K off-diagonal xs. Thus, in total, an NK influence matrix contained $N*(K+1)$ interactions. While a number of studies have investigated various consequences that arise when K increases in a random influence matrix (e.g., Kauffman 1993), we are interested in the effect of different patterns of interactions holding K fixed. Hence, to allow for comparisons of different types of interaction structures, we keep the total number of interactions fixed at $N*(K+1)$, but alter the pattern of interactions among the decisions.

Even for relatively small values of N and K, many possible interaction structures exist. In particular, $N*K$ (off-diagonal) interactions can be placed in N^2-N locations (the N diagonal elements are always filled), creating

$$\frac{(N^2-N)!}{(N*K)!(N^2-N-N*K)!}$$

possibilities. For $N=12$ and $K=2$, for instance, this yields $1.36*10^{26}$ possible influence matrices. For all our analyses, the labeling of individual decisions does not matter (i.e., columns and corresponding rows can be rearranged).[2] This reduces the number of patterns by a factor as large as $N!$, the number of ways that N decisions can be reordered. For $N=12$ and $K=2$, this reduction still leaves a lower bound of $2.84*10^{17}$ different patterns. Given this vast space of possibilities, it is helpful to consider different types of interaction patterns. In particular, we focus on 10 types that were culled from current work on networks, from studies that depict firms as deploying systems of interdependent activities, and from product design analyses.

Influence matrices arise frequently in these contexts even though the term "influence matrix" might not have been used there. The representation of a network as an influence matrix is straightforward (Wasserman and Faust 1994). Each row corresponds to a node of a network, as does each column, with an entry in row i, column j denoting that node j has a link to (and affects) node i. The work on firms as systems of interdependent activities generally has represented firms as consisting of a network of activities that are linked by interactions among them (Porter 1996, Siggelkow 2002). Again, these networks can easily be transformed into influence matrices. Most directly, the product design literature has developed the tool of a "design structure matrix" (DSM) (Steward 1981, Eppinger et al. 1994, Baldwin and Clark 2000), which corresponds to an influence matrix by our definition. A DSM contains all design decisions (e.g., concerning particular design parameters) that have to be resolved. The DSM has an entry in row i, column j if the design choice of element j has an impact on the optimal design choice of element i. For instance, the choice of engine power (element j) might have an impact on the optimal design of the brake system (element i). Table 1 examines all activity system maps that have been published in

[2] For the analysis it would not matter, for instance, whether we label the decision concerning training of the sales force as decision 1 or as decision 2. As long as we keep track of which decisions interact with one another, the labels of the decisions can be interchanged.

Table 1 Characteristics of Actual Design Structure Matrices and Activity Systems

Example	N	$K^{\dagger}$
Design structure matrices		
Automobile brake system (Black et al.1990)	13	3.8
Kodak cartridge development process (Ulrich and Eppinger 2007)	14	2.5
Automobile climate control system (Pimmler and Eppinger 1994)	16	1.4
Automobile door (Dong 1999)	32	3.4
Automobile digital mock-up process for the layout for components in the engine compartment (Ulrich and Eppinger 2007)	50	3.5
Semiconductor development process (Osborne 1993)	60	6.5
Power plant design	72	6.8
Jet engine design (Mascoli 1999)	111	5.8
Activity systems		
Vanguard–1974 (Siggelkow 2002)	18	2.2
Vanguard–1977 (Siggelkow 2002)	24	2.8
Vanguard–1978 (Siggelkow 2002)	29	2.8
Vanguard–1991 (Siggelkow 2002)	41	2.9
Vanguard–1997 (Siggelkow 2002)	48	3.0
Liz Claiborne–1990 (Siggelkow 2001)	36	3.2
Liz Claiborne–1997 (Siggelkow 2001)	34	3.5
IKEA–1996 (Porter 1996)	20	3.4
Southwest Airlines–1996 (Porter 1996)	18	3.4
Vanguard–1996 (Porter 1996)	25	3.4

†The value of K is computed by dividing the number of off-diagonal interaction effects by N.

the literature (Porter 1996; Siggelkow 2001, 2002) and all DSMs that were published on the DSM home page (www.dsmweb.org), which is hosted by Steven Eppinger, Daniel Whitney, and Ali Yassine. For the firm activity systems, N ranges from 18 to 48, and K, calculated as the number of off-diagonal interaction effects divided by N, from 2.2 to 3.5. For the DSMs, N varies from 13 to 111, with K ranging from 1.4 to 6.8.

The 10 different types of influence matrices we explore can be divided into two groups. For five types, each decision is affected by exactly K other decisions. That is, each row of the influence matrix contains exactly K off-diagonal entries. The other five types allow for more heterogeneity among the decisions. For instance, some decisions are allowed to be affected by many other decisions, while other decisions might depend only on themselves.

Random. In a random influence matrix, exactly K xs are placed at randomly chosen off-diagonal positions in each row. For one example with $N = 12$ and $K = 2$, see Figure 1A. This specification is one of the two original specifications of the NK model (Kauffman 1993) and is the setup most commonly used in the organization literature (e.g., Westhoff et al. 1996, Rivkin 2000).

Local. In a local influence matrix, the other original specification, each decision i is assumed to be influenced by its $K/2$ neighbors on either side of it (Figure 1B). For instance, if $K = 2$, decision 3 is affected by decisions 2 and 4. Decisions are assumed to lie on a "ring," i.e., if $K = 2$, decision 1 is affected by decision 2 and decision N. This influence structure is related to Thompson's (1967) notion of "sequential interdependence" and has been employed previously in the organization literature (Levinthal 1997, Gavetti and Levinthal 2000). Moreover, it forms the starting point of the small-world influence structure.

Small world. Although not new, the notion of small-world networks (Milgram 1967) has attracted renewed attention due to recent theoretical advances (Watts and Strogatz 1998). A core feature of small-world networks is that most interactions are local, yet a few interactions exist between elements of the system that are distant from each other. Small-world interaction patterns have been documented in a variety of settings, including ownership patterns among German firms (Kogut and Walker 2001), board of directors' interlocks (Davis et al. 2003), memberships in underwriting syndicates (Baum et al. 2003), firm-alliance networks (Schilling and Phelps 2004), career networks of artists (Uzzi and Spiro 2005, Guimera et al. 2005), and collaboration networks of scientists (Newman 2001).

Following the algorithm by Watts and Strogatz (1998), we create small-world influence matrices in two steps. First, a matrix is initialized with a local influence structure. Second, each off-diagonal x is exchanged with a randomly chosen location in the same row with probability p. For one example, see Figure 1C. One should note that $p = 0$ yields an influence matrix with a local structure, while $p = 1$ creates a random influence structure, as every off-diagonal interaction in the matrix is randomly "rewired."

Block-diagonal. Interactions can be local in a different sense as well. In some systems, decisions can be grouped such that decisions within each group all affect each other, while no interactions across groups exist. This structure relates to the notion of decomposability (Simon 1962) and is the key characteristic of modularity (Eisenhardt and Brown 1999, Baldwin and Clark 2000, Schilling 2000). Block-diagonal structures have been used in a number of NK models (Marengo et al. 2000, Rivkin and Siggelkow 2003, Siggelkow and Levinthal 2003), yet their characteristics have not been compared to other structures. For an example of a symmetric block-diagonal influence matrix, see Figure 1D.

Preferential attachment. In all influence matrices discussed up to this point, each decision is affected by precisely K other decisions, while each decision itself affects K other decisions, on average. In some systems, however, certain decisions exist that are more central than others in the sense that they affect a larger number of other decisions than do most

Figure 1 Different Types of Influence Matrices, All with the Same Number of Total Interactions ($N = 12$, $K = 2$, $N * (K + 1) = 36$)

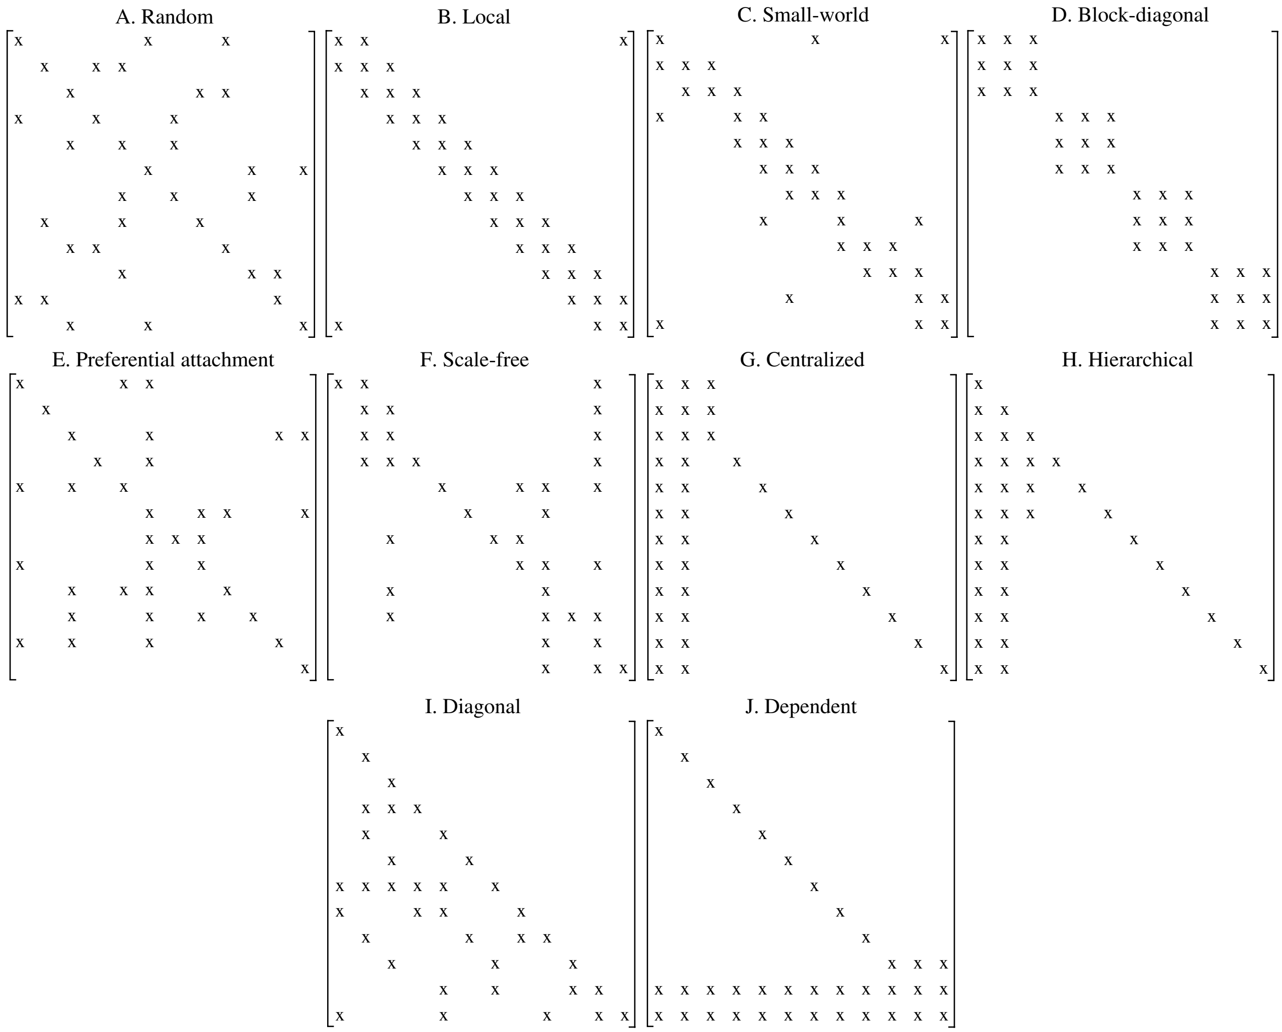

choices. For instance, in the analysis of the mutual fund company Vanguard, Siggelkow (2002) reports that certain of Vanguard's choices were much more central than other choices. Similarly, DSMs often show that certain design elements are much more central. For example, Figure 2 displays the DSM of an automobile brake system as reported by Black et al. (1990). In this DSM, element 4 (corresponding to "piston front size") affects seven out of the other 12 elements of the system, while element 11 ("booster—maximum stroke") influences only itself. Such imbalance in the influence exerted by various elements is sometimes reflected in a distinction between core and peripheral elements (e.g., Hannan and Freeman 1984).

One method of creating networks with elements that are more central than others has been provided by Barabási and Albert (1999). Their algorithm captures a "rich-get-richer" dynamic, by which nodes that already have many interactions are more likely to add a further interaction than are nodes that have few interactions. Thus, interactions are preferentially attached to nodes that already affect many other nodes. We create preferential attachment influence matrices in four steps. First, we initialize a matrix with xs along the main diagonal. Second, we pick one row randomly with equal probability. Call this row i. Third, we pick one column randomly with a probability that is proportional to the number of xs that are already in that column. In particular, if D_j is the number of xs in column j and S is the total number of xs in the matrix at the current point, then the probability that column j is picked is D_j/S. Fourth, if column j was picked, we replace the entry in row i, column j with an x (if there is already an x in (i, j), the x is not changed) and S is updated. We repeat Steps 2–4 until $S = N * (K + 1)$. For one resulting example, see Figure 1E.

Scale-free. A different implementation of the notion that some elements are more central than others assumes that the degree distribution of nodes follows a power law. (Here, the degree of a node equals the number of other nodes it affects.) Networks with

Figure 2 **Design Structure Matrix of an Automobile Brake System Design**

$$\begin{bmatrix}
x & & & & & & & & & & & & \\
x & x & & & & & & & & & & & \\
 & x & x & & & & & & & & & & \\
 & x & x & x & x & & & & & & & & \\
 & x & x & x & x & & & & & & & & \\
x & x & & x & x & x & & & x & x & & & \\
 & x & & x & x & x & x & & x & x & & & \\
 & x & & x & x & x & & x & x & x & & & \\
 & x & x & x & x & x & x & x & x & x & & & \\
x & x & x & x & x & x & x & x & x & x & & & \\
 & & & & & & & & x & & x & & \\
 & x & & x & & x & & & x & & & x & \\
 & & & & & & & & & & & x & x
\end{bmatrix}$$

Source. Black et al. (1990) as in http://www.dsmweb.org/index.php?option=com_content&task=view&id=52.

power-law degree distributions are called scale-free (Barabási and Albert 1999). A number of networks have been shown to be scale-free (Albert et al. 1999, Strogatz 2001, Albert and Barabási 2002).[3] In the context of firm activity systems, the degree distribution in the influence matrix that Siggelkow (2002) reports for the mutual fund provider Vanguard also closely follows a power law.

We create a scale-free influence matrix in two steps. First, we initialize the matrix with xs along the main diagonal. Second, in each column, M off-diagonal xs are added, where M lies between 0 and $N-1$, such that $\text{Prob}(M) = (M+1)^{-\gamma}$. Thus, plotting the number of decisions (M) that a given decision affects against the probability of this occurrence yields a straight line on a log-log scale: $(\ln(\text{Prob}(M)) = -\gamma * \ln(M+1))$. The parameter γ is chosen such that on average the total number of xs in each influence matrix equals $N*(K+1)$. For instance, for $N=12$, setting γ to 1.37 produces the same total number of interactions on average as a random influence matrix with $K=2$. For an example of a resulting influence matrix, see Figure 1F.[4]

Centralized. The centralized influence matrix takes the notion of highly influential decisions to the extreme. It assumes that some decisions affect all other decisions, while other decisions only affect themselves. See, e.g., Barabási (2002, p. 103) for a mechanism that can lead to a "winner-take-all" interaction structure, and Ghemawat and Levinthal (2000) for an application of this influence matrix to organizational search. Starting with xs along the main diagonal, this matrix is created by adding xs into the first column, then into the second column, etc., until the matrix contains a total of $N*(K+1)$ interactions. See Figure 1G.

Hierarchical. The hierarchical influence matrix assumes that decisions are ordered in some fashion, with high-ranked decisions influencing all the decisions below them, but not the decisions above them (Ghemawat and Levinthal 2000). Starting with xs along the main diagonal, we create a hierarchical influence matrix by adding xs below the diagonal, starting with the first column, continuing with the second column, etc., until the matrix contains a total of $N*(K+1)$ interactions. See Figure 1H.

Diagonal. The diagonal influence matrix reflects a situation (as in the hierarchical structure) in which decisions can be ordered such that low-ranked decisions never affect high-ranked decisions, yet decision 1 is not necessarily the most central decision (as it was in the hierarchical structure). Starting with xs along the main diagonal, this matrix is created by randomly adding xs below the diagonal until the matrix contains a total of $N*(K+1)$ interactions. For an example, see Figure 1I. A number of DSMs have diagonal, or close to diagonal, influence matrices. See, for instance, Figure 3, which shows the DSM for the major tasks of a cartridge development project at Kodak, as reported by Ulrich and Eppinger (2007).

Dependent. The dependent influence matrix captures an instance in which a handful of decisions are affected by virtually every other decision the

Figure 3 **Design Structure Matrix for the 14 Major Tasks of Kodak's Cheetah Project (Cartridge Development)**

$$\begin{bmatrix}
x & & & & & & & & & & & & & \\
x & x & & & & & & & & & & & & \\
x & x & x & & & & & & & & & & & \\
 & & x & x & & & & & & & & & & \\
x & x & x & & x & & & & & & & & & \\
 & & x & x & x & x & & & & & & & & \\
x & x & x & & & x & x & x & x & & & & & \\
x & x & & & & x & x & x & x & & & & & \\
 & & & & & & x & x & x & & & & & \\
 & & & & x & & x & & x & x & & & & \\
 & & & & & & & x & & & x & & & \\
 & & & & & & x & x & & & x & x & & \\
 & & & & x & & & & & & x & & x & \\
 & & & & & & & & & x & & x & x & x
\end{bmatrix}$$

Source. Ulrich and Eppinger (2007) as in http://www.dsmweb.org/index.php?option=com_content&task=view&id=53.

[3] The previously described preferential attachment algorithm can yield a power-law distribution if the matrix is allowed to grow, i.e., if nodes are added to the system every time the algorithm cycles through Steps 2–4 (Barabási and Albert 1999). Given the fixed value of N and differing values of K, this approach is not suitable here. As a result, we create a power-law distribution directly.

[4] Despite their difference with respect to containing asymmetrically influential elements, scale-free networks share a number of properties with small-world networks. For instance, their clustering coefficients are much larger than that of random networks of the same size and average degree.

firm makes, yet those decisions exert little influence themselves. We construct such a matrix by transposing the centralized influence matrix. See Figure 1J.

A conceptual distinction arises among our 10 types of influence matrices. For some types, the influence matrix is fully determined once one sets N and K. This is true of the local, block-diagonal, centralized, hierarchical, and dependent types. For the others—the random, small-world, preferential attachment, power-law, and diagonal types—N and K guide the number of entries in the matrix, but chance influences precisely where the xs lie.

3. Creation of Performance Landscapes

To investigate the properties associated with each type of influence matrix, we create a large number of performance landscapes based on each type. First, we specify the type, N, and K. N is the number of binary decisions a firm is assumed to make about how to configure its activities. Hence, an N-digit string of zeroes and ones summarizes all the decisions a firm makes that affect its performance. We represent this "choice configuration" as $\mathbf{d} = d_1 d_2 \cdots d_N$, with each d_i either zero or one. Next, the computer generates a specific influence matrix of the desired type, consistent with N and K.

The final and most intricate step is to assign a performance level to each of the 2^N possible configurations of choices, as follows. Each decision i is assumed to make a contribution C_i to overall firm value, and this contribution is affected by the resolution of decision i and the resolutions of other decisions: $C_i = C_i(d_i;\text{ other } d_j\text{s})$, where the identity of the "js" (i.e., those decisions that influence the contribution of decision i) is specified by the influence matrix. For each possible realization of d_i and the other relevant d_js, a contribution is drawn at random from a uniform $U[0, 1]$ distribution. (Suppose, for instance, that row i of an influence matrix shows decision i to be influenced by three decisions—decision i itself and two other decisions, as for the first decision in Figure 1A. Then, the computer constructs a table with 2^3 contributions for decision i: one random draw for each possible combination of decision i and the two others that affect it.) The overall performance associated with a specific configuration $\mathbf{d}$ is the average over the N contributions:

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

where the C_is are the relevant entries in the tables of contributions mentioned above. One should note that while the choice of influence matrix puts structure on the interaction pattern among decisions, the model remains agnostic as to the type of interaction (e.g., complementarity or substitution) that arises among specific decisions. An intriguing avenue for future research would be to consider the impact of landscape-generating mechanisms that prescribe particular relations among the contribution values. (See also the conclusion section of this paper.)

4. Landscape Characterization

We use each of the 10 different influence matrices to generate performance landscapes and determine a number of topographical characteristics of the resulting landscapes. For all simulations, we consider the case of $N = 12$. (Our results are insensitive to the precise choice of N. Results for $N = 8$ and $N = 16$ are available upon request from the authors.) For each set of landscapes with different interaction patterns, we hold the total number of interactions constant. In particular, we consider influence matrices with 24, 36, 48, 60, 72, and 84 interactions, corresponding to values of K in the traditional random setup of one through six.[5] Below, we occasionally report that the characteristics of the landscapes produced with different influence matrices are statistically significantly different from one another; in each case, we conducted a t-test on differences in the simulation results and found a p-value less than 0.001, making it highly unlikely that reported differences are simply chance occurrences.

A key characteristic of a landscape—indeed, the one we examine most closely—is the number of local peaks it contains. A local peak is a configuration $\mathbf{d}$ such that no configuration $\mathbf{d}'$ exists that differs from $\mathbf{d}$ in only one decision and has higher performance than $\mathbf{d}$. Prior work on the random NK model has documented that increases in K lead to an increase in the number of local peaks (Kauffman 1993). The organizational implications of this feature have been discussed by Levinthal (1997), Rivkin (2000), and others. In contrast, this study is concerned with the number of local peaks that are to be found in landscapes with different underlying patterns of interactions given a fixed value of K, i.e., holding the total number of interactions constant.

The top panel of Table 2 reports the number of local peaks for random, local, and small-world matrices. Recall that the small-world setup involves the parameter, p, the probability of nonlocal interactions, and it includes as special cases the local influence matrix ($p = 0$) and the random influence matrix ($p = 1$). Two patterns in the panel are noteworthy. First, as the interaction structure becomes increasingly

[5] For $K > 5$ (given $N = 12$), it is not possible to construct diagonal, hierarchical, or scale-free influence matrices. As a result, because we are interested in comparisons across influence matrices, we do not investigate values larger than $K = 6$.

Table 2 Characteristics of Local Peaks for Small-World Influence Matrices

	(local) $p = 0.0$	0.05	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	(random) 1.0
Panel 1. Average number of local peaks												
$K = 1$	5.0	5.1	4.8	4.6	4.7	4.2	4.5	4.5	4.6	4.7	4.3	4.8
$K = 2$	14.2	13.2	13.8	13.4	12.2	12.5	12.0	12.0	11.4	11.9	11.3	11.7
$K = 3$	27.5	28.8	26.2	25.2	24.6	24.4	24.3	23.4	23.2	23.5	23.0	23.1
$K = 4$	48.5	46.9	44.5	44.5	42.5	41.7	40.3	39.9	40.0	39.4	39.3	39.3
$K = 5$	71.7	69.2	67.3	66.1	64.8	63.0	61.4	61.9	62.2	61.8	60.7	60.8
$K = 6$	98.2	97.2	95.0	93.3	92.1	90.6	90.1	89.4	88.9	88.3	88.5	88.7
Panel 2. Standard deviation of number of local peaks												
$K = 1$	3.5	3.5	3.3	3.4	3.1	2.9	3.0	3.2	3.2	3.4	2.7	3.0
$K = 2$	6.6	5.5	5.7	6.2	5.5	5.1	5.1	5.3	4.7	5.1	4.9	4.9
$K = 3$	7.4	7.5	6.7	6.4	6.6	6.0	6.3	6.5	5.5	5.8	5.8	6.1
$K = 4$	8.7	8.7	7.4	6.9	7.6	7.4	6.6	6.3	6.5	6.1	6.9	6.9
$K = 5$	7.4	8.2	8.2	7.8	7.5	7.4	7.5	7.4	7.3	7.7	7.1	8.1
$K = 6$	7.9	8.6	7.8	8.7	7.7	8.4	8.2	8.1	8.6	7.9	7.8	8.5
Panel 3. Average local-peak performance (as a portion of maximum performance attainable on landscape)												
$K = 1$	0.947	0.955	0.954	0.951	0.956	0.952	0.951	0.955	0.949	0.947	0.954	0.951
$K = 2$	0.907	0.912	0.912	0.908	0.909	0.910	0.910	0.914	0.911	0.913	0.910	0.909
$K = 3$	0.882	0.882	0.881	0.883	0.881	0.883	0.887	0.883	0.884	0.883	0.883	0.879
$K = 4$	0.860	0.864	0.866	0.865	0.865	0.866	0.866	0.867	0.865	0.866	0.863	0.866
$K = 5$	0.849	0.849	0.852	0.850	0.851	0.851	0.852	0.851	0.851	0.856	0.852	0.853
$K = 6$	0.837	0.843	0.844	0.843	0.842	0.840	0.841	0.843	0.843	0.843	0.842	0.844
Panel 4. Standard deviation of local-peak performance												
$K = 1$	0.053	0.042	0.047	0.050	0.045	0.050	0.049	0.045	0.054	0.052	0.046	0.050
$K = 2$	0.063	0.062	0.061	0.063	0.065	0.064	0.063	0.061	0.062	0.061	0.063	0.064
$K = 3$	0.066	0.066	0.067	0.068	0.069	0.068	0.066	0.066	0.067	0.069	0.067	0.067
$K = 4$	0.068	0.066	0.067	0.066	0.066	0.068	0.067	0.067	0.066	0.068	0.067	0.067
$K = 5$	0.066	0.066	0.066	0.067	0.065	0.066	0.067	0.067	0.067	0.066	0.066	0.066
$K = 6$	0.065	0.065	0.064	0.065	0.065	0.066	0.065	0.066	0.065	0.065	0.065	0.065

Note. Each result is an average over 200 landscapes of each type.

random (i.e., as p increases), the number of local peaks declines. The change, however, is rather modest—a decrease of 10%–20% as one moves from local to random influence. Second, the decline in the number of local peaks is fairly linear with respect to p. The correlation between the number of peaks and p ranges from -0.73 to -0.92 for different values of K. This near-linearity stands in stark contrast to the results of Watts and Strogatz (1998), who identify a number of highly nonlinear relationships in small-world networks, e.g., between the clustering coefficient and p, and between the characteristic path length and p. Thus, while certain aspects of small-world networks respond nonlinearly to p, the number of local peaks in performance landscapes based on small-world influence matrices behaves rather smoothly as p is changed.

The other panels of Table 2 examine, as a function of p and K, the standard deviation of the number of local peaks across landscapes, the average performance associated with a local peak (measured as a fraction of the maximum attainable performance), and the standard deviation in that performance. All of these features are remarkably insensitive to p. Because the landscape features we study depend little on p or, in the case of the number of local peaks, behave linearly in p, we focus below on the extreme cases, the local and random influence matrices, and not on matrices with intermediate values of p.

The first panel of Table 3 contains the number of local peaks for the other seven influence matrices. For reference's sake, we again include the results from the local and random matrices. The panel shows that *landscapes based on the same number of total interactions but different interaction patterns can contain dramatically different numbers of local peaks*. On $K = 2$ landscapes, for instance, the number of local peaks ranges from 3.4 for landscapes based on centralized influence matrices to 133.3 for landscapes based on dependent influence matrices. Similarly, for $K = 5$, the range is from 18.8 (centralized) to 238.2 (dependent).[6] The ends of these ranges differ markedly from the number of local peaks derived from the frequently used random influence matrix.

One immediate consequence of the different number of local peaks is that firms are much more likely to find the global peak in landscapes with centralized

[6] As a benchmark, note that a fully interdependent influence matrix, with $N = 12$ and $K = 11$, has 315.1 local peaks on average.

Table 3 Characteristics of Landscapes Based on Different Types of Influence Matrices

	Centralized	Hierarchical	Scale-free	Random	Preferential attachment	Local	Block-diagonal	Diagonal	Dependent
Panel 1. Average number of local peaks									
$K=1$	1.9	2.3	4.5	4.8	5.0	5.0	5.6	7.6	55.7
$K=2$	3.4	4.9	10.3	11.7	12.1	14.2	16.5	21.6	133.3
$K=3$	6.2	9.7	20.5	23.1	23.4	27.5	33.6	43.6	177.1
$K=4$	11.1	24.3	35.1	39.3	36.9	48.5	40.2	74.2	209.5
$K=5$	18.8	64.3	56.3	60.8	57.7	71.7	82.9	114.4	238.2
$K=6$	32.6			88.7	80.7	98.2	100.3		248.7
Panel 2. Fraction of low-exploration firms that reach global peak (%)									
$K=2$	56.4	42.2	35.1	29.4	27.5	24.2	28.9	19.8	3.3
$K=4$	25.3	16.7	17.0	12.0	9.6	9.7	9.4	8.0	1.5
Panel 3. Portion of all local peaks within Hamming distance of 4 of the global peak (%)									
$K=2$	6.4	17.9	31.9	30.1	33.1	37.5	49.9	29.2	30.8
$K=4$	10.6	15.7	19.2	20.6	20.0	20.6	26.5	19.1	24.5
Panel 4. Standard deviation of number of local peaks									
$K=1$	0.3	0.7	3.1	3.0	2.8	3.5	4.8	4.9	27.5
$K=2$	0.7	1.2	7.4	4.9	4.8	6.6	14.2	9.1	35.7
$K=3$	1.1	1.8	12.9	6.1	7.0	7.4	18.8	11.3	31.4
$K=4$	1.7	3.7	19.0	6.9	8.4	8.7	13.1	14.4	28.3
$K=5$	2.5	7.0	26.3	8.1	11.4	7.4	23.3	15.8	20.1
$K=6$	3.5			8.5	12.1	7.9	18.3		21.8
Panel 5. Average local-peak performance (as a portion of maximum performance attainable on landscape)									
$K=1$	0.957	0.952	0.931	0.951	0.945	0.947	0.949	0.933	0.880
$K=2$	0.921	0.919	0.897	0.909	0.904	0.907	0.912	0.892	0.846
$K=3$	0.902	0.889	0.871	0.879	0.879	0.882	0.889	0.866	0.834
$K=4$	0.883	0.856	0.855	0.866	0.862	0.860	0.871	0.844	0.830
$K=5$	0.865	0.836	0.844	0.853	0.847	0.849	0.857	0.830	0.822
$K=6$	0.852			0.844	0.840	0.837	0.843		0.818
Panel 6. Standard deviation of local-peak performance									
$K=1$	0.062	0.066	0.068	0.050	0.052	0.053	0.049	0.055	0.069
$K=2$	0.078	0.071	0.073	0.064	0.068	0.063	0.062	0.066	0.074
$K=3$	0.076	0.079	0.075	0.067	0.071	0.066	0.067	0.070	0.073
$K=4$	0.078	0.078	0.074	0.067	0.071	0.068	0.069	0.073	0.072
$K=5$	0.076	0.074	0.071	0.066	0.071	0.066	0.069	0.073	0.071
$K=6$	0.072			0.065	0.069	0.065	0.067		0.069

Note. Each result in Panels 1, 2, and 4–6 is an average over 200 landscapes. Each result in Panel 3 is an average over 50 landscapes.

interaction patterns than in landscapes that have dependent interaction patterns. Placing a firm on every point of the landscape and letting each firm continue to search for superior alternatives that differ from the firm's current choices in one decision until the firm has reached a local peak, we report in the second panel of Table 3 the fraction of firms that reach the global peak. In general, a pronounced negative relationship exists between the number of local peaks and the fraction of firms that reach the global peak. For instance, when $K=2$, 56.4% of firms reach the global peak on centralized landscapes, while only 3.3% reach the global peak on dependent landscapes.

An additional feature of interest concerns the clustering of local peaks. Are local peaks clustered around the global peak or are they spread out? As previous studies have argued (Kauffman 1993, Rivkin 2000), the answer to this question is interesting because it captures the degree to which knowledge of one good combination of choices conveys information about the whereabouts of other good combinations. In the third panel of Table 3, we report the fraction of local peaks that differ from the global peak along four or fewer decisions. For $K=2$ landscapes, we detect very different degrees of clustering of local peaks. Block-diagonal landscapes appear to be the most clustered, and centralized landscapes the most dispersed. For $K=4$ landscapes, the differences remain but are much smaller.

In Table 3, the fourth panel notes the standard deviation of the number of local peaks across landscapes. (Given 200 landscapes of each type, the standard error of each mean in the first panel would thus be given by the standard deviation divided by $\sqrt{200-1} \approx 14.1$.)

In general, for each pattern the standard deviation increases with the number of local peaks, except for the block-diagonal and the dependent patterns, for which the standard deviation declines slightly at high levels of K. For all patterns, the coefficient of variation (standard deviation/mean) declines as K increases.[7]

Much of the interest in the number of local peaks in the past literature has arisen because of performance implications: As K increases and the web of potential conflicting constraints thickens, local peaks proliferate, and it becomes harder for a searching firm to achieve a high elevation (Kauffman 1993, Levinthal 1997). It is possible, however, that the proliferation of local peaks that comes from the *pattern* of interactions, not the *number* of interactions (K), has little impact on local peak height. Perhaps, at a given level of K, the numerous peaks associated with, say, the dependent influence matrix are just as high as the few peaks of the centralized matrix. To examine this possibility, we report the average performance of local peaks in Panel 5 of Table 3 (and corresponding standard deviations in Panel 6). Holding the interaction pattern fixed, i.e., looking down each column of Panel 5, we see a negative impact of K on local peak height, familiar from prior literature on random interaction patterns. Likewise, holding K fixed, i.e., examining each row, we also see a strong negative correlation between the number of peaks and average peak performance (ranging from -0.95 for $K = 1$ to -0.80 for $K = 5$). Thus, the average height of peaks declines as the count of peaks rises, regardless of whether their number is increased by more interactions or by a change in interaction pattern.

One should note, though, that average peak height does not give a full picture of how well firms might perform on a particular landscape. Firms that search a landscape are, by the end of their search processes, not equally distributed across all local peaks because some peaks "draw in" more firms than do others.[8] The average performance of firms, then, is a function of both the average peak heights and the number of firms attracted to each peak, a feature that can be affected by the underlying interaction pattern. To take both effects into account, one needs to analyze the performance of firms that have searched the landscape (as we do in §6) rather than inferring performance differences from local peak heights.

We conclude our analysis of the features of performance landscapes by examining two influence matrices drawn from the literature on DSMs. Figures 2 and 3 replicate the DSMs of an automobile brake system and a cartridge design. Using each of these influence matrices, we create 50 performance landscapes and compute the number of local peaks that arise on average. The brake system is composed of $N = 13$ elements, while the cartridge project is composed of $N = 14$ elements. One can measure K for each matrix by counting the number of off-diagonal interactions and dividing by N; a random interaction matrix with this level of K would have the same number of total interactions. This yields $K = 3.8$ and $K = 2.5$ for the two DSMs, respectively. For the brake system, we find that 61.2 local peaks arise on average. This is statistically significantly higher than the 53.0 local peaks in random landscapes with $K = 4$ (and $N = 13$). For the cartridge system (with $K = 2.5$), we find 57.6 local peaks on average, which is statistically significantly higher than the 26.2 local peaks in random landscapes with $K = 2$ (and $N = 14$) and not statistically different from the 53.3 local peaks found on random landscapes with $K = 3$. Thus, in each case, the actual performance landscape appears to be more rugged than the random benchmark.

5. Intuition

Even if the total number of interactions among decisions is held constant, performance landscapes can differ markedly in the number of local peaks they contain. To understand what drives these differences, consider the two influence matrices that produce the fewest and the most peaks: the centralized and the dependent matrices, respectively. In particular, take the matrices shown in Figures 1G and 1J, for which $N = 12$ and the total number of interactions is the same as in a random matrix with $K = 2$. For each of these two, we describe the shapes of the resulting landscapes as well as the underlying intuition for the number of local peaks that arise.

The centralized matrix is distinguished by the large number of columns that contain only one x. These columns represent decisions that do not affect the contributions of other choices. The presence of such "uninfluential" decisions creates large smooth subspaces on each performance landscape—gently sloped plateaus—that limit the number of local peaks (for a related notion of neutral networks, see Lobo

[7] The largest variance is found for landscapes based on scale-free matrices. Here, however, the variance is boosted in part by the landscape-generating mechanism itself, which guarantees an average of $N * K$ off-diagonal interactions, but permits some variation in the number of interactions from one specific landscape to another.

[8] More technically, suppose that each firm searches the landscape by changing one decision at a time. A local peak's one-decision basin of attraction is the set of points from which a firm would be able to reach that peak by a series of upward steps, each involving adjustment to a single decision. Local peaks on a given landscape will differ in how large their one-decision basins of attraction are. When firms are assigned initial configurations of choices at random, local peaks with larger basins will "draw in" more firms than peaks with smaller basins.

et al. 2004). In Figure 1G, for instance, suppose that decisions 1, 2, and 3 have been set. The contribution of each remaining decision then depends only on the resolution of that decision itself. The best configuration of the remaining choices conditional on d_1, d_2, and d_3 is easy to find: Simply set d_4 to zero or one, whichever produces higher performance, and then do the same for $d_5, d_6, \ldots, d_{12}$. Because decisions 4–12 are uninfluential, the alteration of each does not affect the contributions of the other decisions, and this simple procedure produces the greatest possible performance conditional on decisions 1–3. Thus, for each possible configuration of $\{d_1\ d_2\ d_3\}$, there is a plateau that rises smoothly to a maximum, and the total number of local peaks can be no greater than eight, the number of different configurations of $\{d_1\ d_2\ d_3\}$ (Solow et al. 1999a). In fact, the number may be smaller than eight if the maximum point on any plateau is below an adjacent point on another plateau. The actual number of local peaks, on average, is 3.4 (Table 3). (For a visualization of a landscape produced by a centralized matrix that shows the presence of distinctive plateaus, see Figure EC.2 of online Appendix 1, which is provided in the e-companion.)[9]

In more intuitive terms: The presence of uninfluential decisions reduces the number of choices that threaten to confound the decision maker and face her with difficult trade-offs. In the matrix in Figure 1G, for instance, once decisions 1, 2, and 3 have been made, the remaining choices are obvious. The number of potentially conflicting constraints plunges, and this simplifies matters dramatically. As the effective dimensionality of the problem falls, broad exploration for solutions becomes less valuable (as we demonstrate directly in the next section).

In contrast to the centralized matrix, the dependent matrix is distinguished by the large number of rows that contain only one x and the small number of rows that contain many xs. In matrix 1J, for instance, each of decisions 1–9 makes a contribution to performance that is not influenced by other decisions, while decisions 10–12 are sensitive to many other choices. The "uninfluenced" decisions 1–9 create a distinctive topography: The performance contributions from these choices alone form a smooth, single-peaked surface, as would arise from an $N = 9$, $K = 0$ matrix. Consider two choice configurations that differ only in terms of one of these nine decisions. The performance of these adjacent points can differ from one another by no more than $1/N$, the maximum performance contribution of the decision that distinguishes those configurations. Accordingly, decisions 1–9 form a smooth underlying surface. Added onto that surface to form the complete performance landscape are the contributions of decisions 10–12. These contributions are very sensitive to many other choices: Indeed, the contributions of decisions 11 and 12 change from one randomly drawn number to another whenever *any* decision is altered. A change in a single decision can alter the total contributions of decisions 10–12 by as much as $3/N$. Naturally, the addition of relatively large random increments to a smooth underlying surface creates a landscape with many, many local peaks, akin to the dimpled surface of a golf ball.[10] (A landscape produced by a dependent matrix, depicted in Figure EC.3 of online Appendix 1, shows the very bumpy nature of the terrain.)

More intuitively, the concentration of many decisions' influences onto a handful of decisions creates the potential for many conflicting constraints and lots of internally consistent configurations of choices. From each of these consistent configurations, a change in one decision leads to lower performance, but changes in two or more decisions might cause performance to improve again. This is especially likely when many decisions are uninfluenced, causing all configurations to have a similar underlying level of performance and permitting small differences to create numerous local optima. As we show below, this increases the need for broad exploration to escape poor local optima and to find a good one.

The intuition for the centralized and dependent matrices lead us to a hypothesis: For a given number of total interactions in an influence matrix, *the number of local peaks declines with the number of uninfluential decisions* (*i.e., those with one* x *per column*) *and rises with the number of uninfluenced decisions* (*i.e., those with one* x *per row*). To examine this hypothesis further, we focus on $K = 3$, generate 50 influence matrices of each type shown in Table 3, count the number

[9] An electronic companion to this paper is available as part of the online version that can be found at http://mansci.journal.informs.org/.

[10] In contrast, if the underlying surface is already somewhat rugged, the perturbations caused by decisions that are affected by many other decisions create fewer additional local peaks. The following analysis confirms this intuition. An $N = 12$, $K = 0$ landscape is very smooth, containing only one peak. Its influence matrix contains xs only on the diagonal. If we fill one row of this influence matrix with xs, i.e., make one decision's contribution dependent on all other decisions, the average number of local peaks increases sharply to 58. Now start with an influence matrix in which decision 1 is affected by itself and decision 2, decision 2 is affected by itself and decision 3, etc. This influence pattern, which contains no uninfluenced decisions, leads to a performance landscape with nine local peaks. Filling one row of this influence matrix with xs increases the number of local peaks only to 39. An intriguing implication of this finding is that adding interactions to some influence matrices can reduce the number of local peaks by effectively changing the underlying pattern of interactions.

of uninfluential and uninfluenced decisions in each matrix, generate a performance landscape with each, and count the number of local peaks on each. This produces a sample of 450 landscapes (50 landscapes per type $*$ 9 types). We then use this sample to regress the number of local peaks on the number of uninfluential decisions and the number of uninfluenced decisions, and we obtain

$$\begin{aligned}&\text{number of local peaks}\\&\quad = 27.4 - 4.0 * \text{number of uninfluential decisions}\\&\qquad\qquad (t\text{-stat} = -4.9)\\&\qquad + 19.8 * \text{number of uninfluenced decisions.}\\&\qquad\qquad (t\text{-stat} = 14.9).\end{aligned}$$

The very large t-statistics confirm the power of these two variables to predict the number of local peaks. Indeed, the two variables explain 89.3% of the variance in the number of local peaks.

This suggests that one can inspect two influence matrices, count the number of uninfluential and uninfluenced decisions, and predict with accuracy which is likely to produce more local peaks and, accordingly, which will probably require more exploration. We return to the power and the limits of this hypothesis in the concluding section.

6. Value of Exploration

In prior sections, we have asserted that the proliferation of local peaks increases the value of, and the need for, broad exploration. Other research has shown this to be true when the proliferation comes from an increase in K (Kauffman 1993, Rivkin and Siggelkow 2003). Here, we illustrate that interaction patterns that produce more local peaks, even if K is fixed, also call for broader exploration in the long run. To do so, we examine the performance paths of three firms:

1. A *low-exploration firm* starts at a random choice configuration **d**, evaluates in each period a randomly chosen alternative **d**′ that differs from **d** in terms of one decision, and adopts **d**′ if it yields higher performance. The firm continues to do so each period until it can find no superior alternatives. At that point, it rests atop a local peak.

2. The *medium-exploration firm* allocates some of its search efforts to the consideration of more distant alternatives. Specifically, in each period it considers a randomly chosen alternative **d**′ that differs from **d** in terms of one or two decisions. For instance, in an $N = 12$ simulation, a medium-exploration firm at 000000000000 might evaluate the alternative 010000000100. We refer to this firm as having a search radius of two.

3. The *high-exploration firm* considers each period an alternative **d**′ that differs from **d** in terms of as many as five decisions (search radius $= 5$). Although we define "high-exploration" as the ability to change as many as five decisions, our results are not sensitive to this specific choice.[11]

All firms are given the same randomly chosen starting point and are allowed to search for better configurations for 2,500 periods. We calculate the performance of each firm relative to the global peak of the landscape, record the performance level of each firm over time, and then generate a new performance landscape with the same underlying influence matrix. For each type of influence matrix, we repeat this exercise 500 times.[12] The performance differences across firms capture the value of broader exploration in the face of each type of influence matrix. In particular, because the model assigns zero cost to search, the performance differences can be interpreted as the most a firm would be willing to pay for the broader search capability. (While we use a firm's search radius to manipulate its degree of exploration directly, exploration can also be influenced by the firm's organizational structure. For an example with results consistent with the subsequent findings of this section, see online Appendix 2, which is provided in the e-companion.)

We discuss long-run performance differences and then turn to performance levels over time. Table 4 reports the average performance advantage of the medium-exploration firm over the low-exploration firm in Period 2,500 (Panel 1) and the advantage of the high-exploration firm over the low-exploration firm at the same time (Panel 2). The results reveal three striking patterns concerning the value of broader exploration. First, as one would expect in a setup where exploration is made costless, medium- and high-exploration firms have markedly better long-run performance than low-exploration firms for all levels of K and for all types of influence matrices. The performance advantage of the medium-exploration firm ranges from 1.0% of 10.1% of the peak heights shown in Panel 5 of Table 3, and the advantage of the high-exploration firm ranges from 2.2% to 16.4%. All differences are highly statistically significant with $p < 0.001$. Second, within each type of influence matrix (i.e., for each column of the table), higher levels of K make

[11] Because each firm evaluates alternatives chosen at random from those within its search radius, the relative frequency of search for more distant alternatives is dictated by the frequency with which such alternatives exist. The medium-exploration firm, for instance, has 12 alternatives that differ from the status quo in terms of one decision and 66 that differ in terms of two decisions. Hence, it considers a one-decision change $12/(12+66) = 15\%$ of the time and a two-decision change 85% of the time. The high-exploration firm considers a one-, two-, three-, four-, and five-decision change 1%, 4%, 14%, 31%, and 50% of the time, respectively.

[12] We examine 2,500 periods because, by that time, modeled firms have largely exhausted their improvement opportunities. A sample size of 500 landscapes of each type produces performance differences across firms that are statistically significant.

Table 4 Long-Run Value of Broader Exploration

	Centralized	Hierarchical	Scale-free	Random	Preferential attachment	Local	Block-diagonal	Diagonal	Dependent
Panel 1. Performance difference between medium-exploration firm and low-exploration firm									
$K = 1$	0.010	0.017	0.020	0.023	0.022	0.026	0.034	0.027	0.041
	0.050	0.052	0.046	0.039	0.039	0.037	0.041	0.042	0.050
$K = 2$	0.027	0.027	0.035	0.033	0.035	0.045	0.058	0.046	0.072
	0.071	0.066	0.064	0.051	0.056	0.054	0.054	0.064	0.072
$K = 3$	0.025	0.042	0.041	0.036	0.039	0.050	0.068	0.049	0.076
	0.071	0.075	0.075	0.061	0.063	0.068	0.066	0.074	0.072
$K = 4$	0.037	0.049	0.045	0.041	0.043	0.052	0.072	0.047	0.080
	0.073	0.075	0.077	0.070	0.073	0.071	0.072	0.073	0.077
$K = 5$	0.043	0.058	0.043	0.050	0.048	0.053	0.077	0.063	0.079
	0.076	0.080	0.073	0.071	0.079	0.069	0.071	0.076	0.070
$K = 6$	0.050			0.047	0.058	0.059	0.070		0.083
	0.075			0.071	0.075	0.076	0.071		0.078
Panel 2. Performance difference between high-exploration firm and low-exploration firm									
$K = 1$	0.021	0.027	0.026	0.025	0.028	0.027	0.032	0.034	0.044
	0.051	0.052	0.044	0.040	0.040	0.037	0.041	0.041	0.048
$K = 2$	0.049	0.050	0.054	0.044	0.051	0.054	0.063	0.063	0.084
	0.067	0.064	0.060	0.053	0.055	0.050	0.053	0.061	0.069
$K = 3$	0.056	0.071	0.068	0.062	0.064	0.070	0.082	0.077	0.101
	0.067	0.073	0.072	0.058	0.060	0.061	0.064	0.067	0.069
$K = 4$	0.072	0.089	0.078	0.075	0.082	0.082	0.102	0.087	0.116
	0.068	0.070	0.071	0.064	0.066	0.063	0.067	0.066	0.068
$K = 5$	0.089	0.105	0.090	0.091	0.091	0.089	0.114	0.109	0.126
	0.073	0.072	0.069	0.066	0.071	0.064	0.064	0.072	0.065
$K = 6$	0.096			0.094	0.103	0.106	0.116		0.134
	0.070			0.063	0.073	0.068	0.069		0.069

Notes. Each cell contains the performance difference in Period 2,500 between a firm that engages in either medium exploration (evaluates alternatives that differ in up to two decisions from the status quo) or high exploration (evaluates alternatives that differ in up to five decisions from the status quo) and a firm that engages in low exploration (i.e., evaluates only alternatives that differ in one decision from the status quo). Performance is measured relative to the highest performance possible in each landscape. Performance differences are averages over 500 landscapes. Standard deviations are given under each performance difference. Standard errors are smaller by a factor of $\sqrt{500-1}$ ($\approx$22). All performance differences in this table are statistically significant with $p < 0.001$.

broad exploration more valuable—a finding in line with the prior research mentioned above.

Third and crucially, the long-run value of broad exploration varies significantly across types of influence matrices even if the total number of interactions is held constant (i.e., for each row of the table). Moreover, the within-row differences correspond closely to differences in the number of local peaks. As the number of local peaks increases, the value of broad exploration increases even if K is fixed.[13] For $K = 2$, for instance, the centralized matrix produces only 3.4 local peaks on average, and the performance advantage of the medium-exploration firm is merely 0.027, while the dependent matrix generates many more local peaks, 133.3, and the value of broader exploration is statistically significantly higher at 0.072 ($p < 0.001$). Indeed, the number of local peaks appears to do as good a job as K at predicting the value of broad exploration. Casual inspection of the top panels of Tables 3 and 4 supports this notion. It is easy to find pairs of landscapes that have approximately the same number of local peaks and lead to the same performance advantage for the medium-exploration firm despite differences in K. Consider, for instance, the random matrix with $K = 2$ and the centralized matrix with $K = 4$; the scale-free matrix with $K = 3$ and the centralized matrix with $K = 5$; and the dependent matrix with $K = 1$ and the scale-free matrix with $K = 5$. More rigorously, an analysis of the value of broad exploration in Table 4, the number of local peaks in Table 3, and K reveals that (a) the number of local peaks is a strong and robust predictor of the long-run value of broad exploration, with a correlation coefficient of 0.81 for the medium-exploration firm and 0.77 for the high-exploration firm; and (b) the number of local peaks explains as much of the variance in the value of broad exploration as does K alone.

[13] For the diagonal influence matrix, the benefit of broader search is somewhat lower than one would expect given its number of local peaks. For the hierarchical matrix, it is slightly larger.

Overall, the long-run results lend strong support to two related notions. First, the marginal value of broader exploration is greater in the long run on landscapes with more local peaks. Second, the investment in broader exploration that a firm can justify depends on the pattern—not just the density—of interactions among choices. Our findings validate the prior literature's results on the impact of K alone on the value of exploration, but they highlight an important caveat: Those results implicitly assume that the underlying interaction pattern remains constant as K changes.

Interesting dynamics underlie the long-run results we have just described. To be concrete, we focus on the difference between a high-exploration and a low-exploration firm. In comparing the performance of these firms over time, one can identify three phases with distinct dynamics. During the earliest periods, a high-exploration firm can enjoy an advantage over the low-exploration firm for the following reason. Both firms start from the same randomly chosen configuration of choices, and each evaluates an alternative within its search radius, accepting the alternative only if it yields better performance. The average height of points within a search radius of five of a randomly chosen point is, in expectation, the same as the average height of points within a search radius of one, but the variance of heights in the first set of points is greater. Because firms accept only improvements and start with the same performance level, the first improvements of the low- and high-exploration firms are essentially random draws from the upper tails of two distributions with equal means but different variances. The draw is greater in expectation for the firm that experiences the higher variance, the high-exploration firm. Intuitively, an upward long jump from a randomly chosen location brings a greater chance of finding a dramatic improvement than does a tweak from that location.

Subsequently, the low-exploration firm scales a nearby peak, exploiting the local correlation of the landscape (i.e., the fact that above-average locations tend to be surrounded by other above-average locations). Meanwhile, the high-exploration firm wastes many periods weighing distant alternatives that yield little or no performance improvement. As a result, the low-exploration firm enjoys a period of advantage in this second phase.

Lastly, however, the low-exploration firm exhausts its improvement opportunities and gets stuck on an average local peak, while the high-exploration firm long continues to discover better alternatives. Eventually, it reaches a relatively high local peak and achieves the higher performance evidenced in Table 4.

These three phases can be seen in Figure 4A, which shows the average difference in cumulative performance between a firm with search radius five and a firm with search radius one on a landscape with an underlying centralized interaction pattern and $K = 1$. In early periods, the high-exploration firm outperforms the low-exploration firm, leading to an ever-increasing cumulative performance advantage until Period 15. (Please note that the x-axis has a logarithmic scale.) Starting from that point, the low-

Figure 4 **Cumulative Performance Differences Between High- and Low-Exploration Firms**

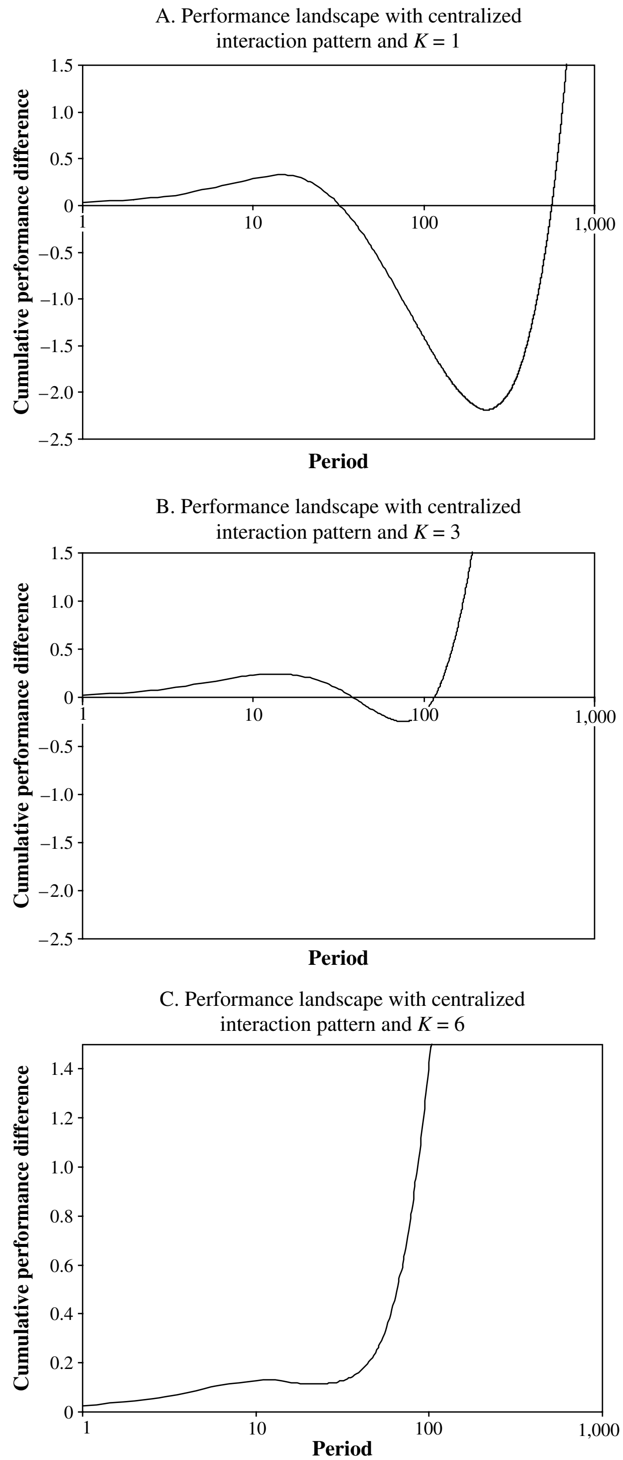

exploration firm outperforms the high-exploration firm until about Period 230. During this time, the low-exploration firm outperforms the high-exploration firm to such a degree that even the cumulative performance turns to the low-exploration firm's advantage. After Period 230, the high-exploration firm's broader search starts to pay off, leading to higher performance and higher cumulative performance after about Period 560.

As the number of interdependencies increases (Figures 4B and 4C), the size of the effects in each of the three phases changes. (Similar dynamics can be seen for the other interaction patterns.) Recall, that the key driver of the performance difference in the first phase was that configurations further away from a randomly chosen starting point have greater performance variance than do configurations that are close by. As interactions increase and the landscape becomes more rugged, the variance increases for all radii. However, it increases more for configurations close by than for configurations far away. As a consequence, the difference between the variances of performances that local and broad searches can access decreases, and the first effect becomes smaller as K increases.[14]

The performance advantage of the low-exploration firm in the second phase has its root in the low-exploration firm's ability to reach a local peak relatively quickly. As K rises and the average height of the (nearby) local peak decreases, one would expect this performance advantage to decline.[15] Figures 4B and 4C bear out this intuition: The temporary performance advantage of the low-exploration firm in Phase 2 declines as K increases. Lastly, as argued above, the long-run value of broad exploration increases with the number of local peaks. Consequently, the cumulative performance advantage of the high-exploration firm sets in earlier as K increases.

[14] Performance correlation between adjacent configurations is high when many decisions are unaffected by other decisions. As K increases, this becomes less likely and performance variance among nearby configurations increases. Configurations that are further apart from each other are already less correlated at lower levels of K. Hence, increasing the number of interdependencies tends to affect their correlation less. For instance, for landscapes with centralized interaction patterns, the difference between the standard deviation of performance values of configurations that lie within a five-step radius of a focal point and the standard deviation of performance of configurations that lie within a one-step radius of a focal point decreases from 0.023 for $K = 1$ to 0.008 for $K = 6$. This decrease is driven largely by an increase in the standard deviation of performance of configurations that lie within a one-step radius.

[15] It is possible that the short-term advantage of the high-exploration firm and the intermediate performance advantage of the low-performance firm depend on the interaction pattern per se, not just K. We have not yet spotted robust relationships along these lines, but we consider this an intriguing avenue for future research.

The upshot of these dynamics is that the appropriate investment in exploration depends not only on the number and pattern of interactions, but also on the time horizon of a management team. The swings in advantage shown in Figure 4 work against simple heuristics (such as "shorter horizons call for more exploitation and less exploration"). By reproducing Table 4 for earlier periods or by comparing cumulative performance over different durations, one can generate a wide variety of results without robust patterns. Consequently, we emphasize performance consequences in the long run, where patterns are clear and stable. There, we find that the proliferation of local peaks, due to the number or pattern of interactions, boosts the value of broad exploration.

7. Discussion and Conclusion

In management science, the study of complex systems has recently gained momentum as simulation tools, originally developed in biology and physics, have been applied to organizational, social, and technological settings. This paper aims to make such simulation models more realistic by incorporating into one particular model some of our empirical knowledge of such settings. Many simulation models in this field of inquiry have two parts: a problem space (a performance landscape, an environment, etc.), and entities that search (or move, or live) in the problem space. The early models in this genre were—as a natural starting point—fairly simplistic in both respects. The original NK model, for instance, which formed the starting point for many applications in the organization literature, assumed performance landscapes in which the interactions among elements were determined randomly and entities in which change occurred only through incremental, local search. While the latter assumption was appropriate for biological systems that evolve by mutations to single, randomly chosen genes, it is dubious for organizational, social, and technological systems in which human agents can employ more sophisticated forms of search. A number of studies have attempted to model search more realistically, incorporating cognition (Gavetti and Levinthal 2000) or internal organizational structure (Siggelkow and Rivkin 2005), for instance.

The main thrust of this paper was to infuse more realism into the first part of these simulation models, the creation of performance landscapes. The random interaction assumption has often been justified by pleading ignorance of what true interaction patterns look like. That plea is implausible, we feel, in the settings that interest management scientists, thanks to recent empirical studies. These studies show that the interactions among activities, product elements, decisions, and decision makers are not random, but follow distinctive patterns. We identified 10 patterns (including the random benchmark) and examined the characteristics of landscapes produced by each. We found that underlying interaction patterns affect landscape topography substantially even if the total number of interactions is held constant. In particular, dependent, diagonal, and, to a lesser degree, local and block-diagonal interaction patterns tend to generate performance landscapes with substantially more local peaks than the random interaction pattern, while centralized and hierarchical interaction patterns typically lead to substantially fewer local peaks. Interestingly, small-world type interaction patterns exhibit linear, rather than nonlinear, changes in the number of local peaks as the probability of nonlocal interaction is changed.

The interaction patterns that produce very few local peaks are marked by a handful of highly influential decisions and a large number of uninfluential decisions. These patterns produce landscapes that are easy to search: Once the handful of core decisions are made, other choices fall into place naturally. As a result, the decision maker faces a problem whose true dimensionality is modest. In contrast, interaction patterns with a handful of highly sensitive decisions and a large number of uninfluenced decisions tend to produce many local peaks. The uninfluenced decisions produce a smooth underlying surface that is made very rugged by the handful of sensitive decisions. For a given level of K, we can explain a remarkably high portion of the variance in the number of local peaks—nearly 90%—by reference to the number of uninfluential and uninfluenced decisions.[16] This suggests a practical rule of thumb for individuals who are deciding how much to invest in long-run exploratory efforts. Relatively little exploration is required in systems where a handful of core decisions influence a large number of peripheral, otherwise-independent choices. More exploration is necessary in systems where a large number of independent decisions converge to influence a handful of choices.

We have emphasized the implications of these results for the allocation of resources toward exploration versus exploitation. When facing interaction patterns that create many local peaks and when focusing on long-run consequences, managers are well advised to devote more resources to exploration. Although our simulation results focus on the value of broad exploration, we believe they also have ramifications for other organizational phenomena. For instance, prior research efforts with related models have shown that the proliferation of local optima makes it difficult for organizations to adjust successfully in the face of environmental change (Levinthal 1997), to imitate the successes of others (Rivkin 2000), and to replicate their own successes (Rivkin 2001). These research efforts have focused on increases in the total number of interactions as the reason for the proliferation of local peaks, but proliferation caused by differences in influence matrices should have similar effects. Thus, we see interaction patterns affecting not just the appropriate degree of exploration, but also the likely success of change, imitation, and replication efforts.

Similar logic suggests a cautionary word about previous studies that have examined only random influence matrices. Most of these studies were concerned with effects that arise as the number of interactions, K, increases. Our results imply that comparative static results with respect to changes in K, such as "imitation becomes more difficult as K increases," continue to hold as long as the underlying interaction pattern remains fixed. The results also show, however, that K is not the only factor that determines landscape characteristics and consequent competitive phenomena. For instance, a firm that has based its competitive advantage on a set of choices with high K and a centralized interaction pattern may find that its advantage is eroded by imitation more easily than if it had a lower value of K, but a diagonal interaction pattern.

The sensitivity of landscape topography to the structure of the influence matrix suggests that other aspects of the NK model might be fruitfully investigated in future research. In the original NK model, chance plays a powerful role in two places: in generating the influence matrix and in assigning contributions C_i to sets of choices. Because it relies heavily on chance and imposes little structure, the NK model can reveal general properties of a wide class of systems, but it can say little about specific systems. Our paper shows that, if we impose structure on the generation of influence matrices, we can refine predictions about landscape topography. Similarly, if one were to assign performance contributions in a more structured way, one might obtain results that differed meaningfully from the results of the original NK model. Prior research supports this speculation. Solow et al. (1999b), for instance, find that assigning contributions in alternative ways can attenuate a

[16] We suspect that our ability to explain so much of the variance by looking only at polar cases, wholly uninfluential decisions, and completely uninfluenced decisions, reflects an extreme assumption of the NK model: A change in any influential decision *completely* rerandomizes the contribution of a focal decision. Under a less extreme assumption, a change in an influential decision would alter the focal decision's contribution, but not completely. In such a setting, one might have to take into account more than simply the number of wholly uninfluential and completely uninfluenced decisions to anticipate the number of local peaks. For instance, one might have to calculate how concentrated influence is in, and on, a handful of decisions. This is a speculation that deserves investigation in future research.

central finding of Kauffman's early work, the "complexity catastrophe" (Kauffman and Levin 1987). In an application to political science, Axelrod and Bennett (1993) use specific ethnic, religious, territorial, ideological, and historical data rather than arbitrary figures to generate a landscape of potential World War II alignments. The use of actual data reduces the number of local peaks—stable configurations of nations—from 209 to two. A similar study of alliances in a very different setting—alliances of firms vying to establish Unix standards for workstations—shows that actual data can reduce the number of local peaks (in this case, Nash-equilibrium alliances) from 256 to as few as two, depending on the choice of unmeasured rivalry parameters (Axelrod et al. 1995). Such studies suggest that it would be worthwhile to investigate the impact of more structured contributions in the NK model as well. To do so, however, one must find sensible ways to characterize relations among contributions, something akin to the 10 influence matrices that we culled from prior research.

Our results also have system-design implications because firms can sometimes influence the interaction patterns they face rather than take them as given (Levinthal and Warglien 1999, Baldwin and Clark 2000). Because optimization of high-dimensional systems with many interdependencies is usually a difficult task, it may be very helpful to design a system in a way that smoothes performance landscapes and facilitates the search for good solutions. A management team might accomplish this by altering the pattern of interactions among elements in a system, even if the total number of interactions among the elements cannot be reduced. (For an interesting example of a redesign of a computer software architecture along these lines, see MacCormack et al. 2006.) Smoothing of a landscape may also make the system more robust—able to recover effectively after a perturbation in the mapping from choices to performance. On the other hand, if competitors can reproduce a firm's design of interactions, smoothing might make local search a more powerful means for rivals to rediscover a firm's configuration of choices and to copy its successes.

By managerial intervention or by the selective force of births and deaths of systems, the patterns of interactions present in organizational, social, and technological systems are likely to evolve. An exciting question for future research is: What interaction patterns will prevail over time? Or perhaps a contingent question is appropriate: What conditions encourage the emergence of which kinds of interaction patterns? Simon (1962) makes a strong argument for nearly decomposable systems, on the strength of their ability to improve module by module rather than in systemwide fashion. Patterns of interaction, however, may affect not only the power of exploration across discrete modules, but also the ability of managers to explore possibilities within each module. Our results show that the pattern of interactions among decisions can dramatically alter the search challenge that managers face. Patterns that improve "searchability" may very well prevail in ecological competition among interaction patterns.

8. Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at http://mansci.journal.informs.org/.

Acknowledgments

The authors express special thanks to Howard Brenner for computer programming, to Monica Yan, Gavaskar Balasingam, and John Lafkas for research assistance, and to Brian Uzzi and anonymous referees for helpful comments. The authors are grateful to the Mack Center for Technological Innovation and the Division of Research of Harvard Business School for generous funding. Errors remain the authors' responsibility.

References

Abernathy, W. J. 1978. *The Productivity Dilemma: Roadblock to Innovation in the Automobile Industry*. John Hopkins University Press, Baltimore, MD.

Albert, R., A.-L. Barabási. 2002. Statistical mechanics of complex networks. *Rev. Modern Phys.* **74** 47–97.

Albert, R., H. Jeong, A.-L. Barabási. 1999. Diameter of the world-wide web. *Nature* **401** 130–131.

Argyris, C., D. A. Schön. 1978. *Organizational Learning*. Addison-Wesley, Reading, MA.

Ashby, W. R. 1960. *Design for a Brain*, 2nd ed. Wiley, New York.

Axelrod, R., D. S. Bennett. 1993. A landscape theory of aggregation. *British J. Political Sci.* **23** 211–233.

Axelrod, R., W. Mitchell, R. E. Thomas, D. S. Bennett, E. Bruderer. 1995. Coalition formation in standard-setting alliances. *Management Sci.* **41** 1493–1508.

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

Barabási, A.-L. 2002. *Linked: The New Science of Networks*. Perseus, Cambridge, MA.

Barabási, A.-L., R. Albert. 1999. Emergence of scaling in random networks. *Science* **286** 509–512.

Baum, J. A. C., A. V. Shipilov, T. J. Rowley. 2003. Where do small worlds come from? *Indust. Corporate Change* **12** 697–725.

Benner, M. J., M. L. Tushman. 2003. Exploitation, exploration, and process management: The productivity dilemma revisited. *Acad. Management Rev.* **28** 238–256.

Black, T. A., C. H. Fine, E. M. Sachs. 1990. A method for systems design using precedence relationships: An application to automotive brake systems. Working Paper 3208, Sloan School of Management, Massachusetts Institute of Technology, Cambridge, MA.

Burns, T., G. M. Stalker. 1961. *The Management of Innovation*. Tavistock, London, UK.

Davis, G. F., M. Yoo, W. E. Baker. 2003. The small world of the American corporate elite. *Strategic Organ.* **1** 301–326.

Dong, Q. 1999. Representing information flow and knowledge management in product design using the design structure matrix. Master's thesis, Mechanical Engineering Department, Massachusetts Institute of Technology, Cambridge, MA.

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.

Erdős, P., A. Rényi. 1959. On random graphs. *Publicationes Mathematicae* **6** 290–297.

Ethiraj, S. K., D. A. Levinthal. 2004. Modularity and innovation in complex systems. *Management Sci.* **50** 159–173.

Fleming, L. 2001. Recombinant uncertainty in technological search. *Management Sci.* **47** 117–132.

Gavetti, G., D. Levinthal. 2000. Looking forward and looking backward: Cognitive and experiential search. *Admin. Sci. Quart.* **45** 113–137.

Ghemawat, P., D. Levinthal. 2000. Choice structures and business strategy. Working Paper 01-012, Harvard Business School, Boston, MA.

Ghemawat, P., J. E. Ricart i Costa. 1993. The organizational tension between static and dynamic efficiency. *Strategic Management J.* **14** 59–73.

Guimera, R., B. Uzzi, J. Spiro, L. A. Nunes Amaral. 2005. Team assembly mechanisms determine collaboration network structure and team performance. *Science* **308** 697–702.

Hannan, M. T., J. Freeman. 1984. Structural inertia and organizational change. *Amer. Sociol. Rev.* **49** 149–164.

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

Kauffman, S. A., S. Levin. 1987. Toward a general theory of adaptive walks on rugged landscapes. *J. Theoret. Biol.* **128** 11–45.

Kauffman, S. A., E. D. Weinberger. 1989. The NK model of rugged fitness landscapes and its application to maturation of the immune response. *J. Theoret. Biol.* **141** 211–245.

Kogut, B., G. Walker. 2001. The small world of Germany and the durability of national networks. *Amer. Sociol. Rev.* **66** 317–335.

Langton, C. G. 1990. Computation at the edge of chaos: Phase transition and emergent computation. *Physica D* **42** 12–37.

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.

Levinthal, D. A., M. Warglien. 1999. Landscape design: Designing for local action in complex worlds. *Organ. Sci.* **10** 342–357.

Lobo, J., J. H. Miller, W. Fontana. 2004. Neutrality in technological landscapes. Working paper, Santa Fe Institute, Santa Fe, NM.

MacCormack, A., J. Rusnak, C. Baldwin. 2006. Exploring the structure of complex software designs: An empirical study of open source and proprietary code. *Management Sci.* **52** 1015–1030.

March, J. G. 1991. Exploration and exploitation in organizational learning. *Organ. Sci.* **2** 71–87.

Marengo, L., G. Dosi, P. Legrenzi, C. Pasquali. 2000. The structure of problem-solving knowledge and the structure of organizations. *Indust. Corporate Change* **9** 757–788.

Mascoli, G. J. 1999. A systems engineering approach to aero engine development in a highly distributed engineering and manufacturing environment. System, design and management thesis, Massachusetts Institute of Technology, Cambridge, MA.

McKelvey, B. 1999. Avoiding complexity catastrophe in coevolutionary pockets: Strategies for rugged landscapes. *Org. Sci.* **10** 294–321.

Milgram, S. 1967. The small world problem. *Psych. Today* **2** 60–67.

Newman, M. E. J. 2001. The structure of scientific collaboration networks. *Proc. Natl. Acad. Sci. USA* **98** 404–409.

Osborne, S. M. 1993. Product development cycle time characterization through modeling of process iteration. Master's thesis (management/engineering), Massachusetts Institute of Technology, Cambridge, MA.

Pimmler, T. U., S. D. Eppinger. 1994. Integration analysis of product decompositions. Working Paper 3690-94, Sloan School of Management, Massachusetts Institute of Technology, Cambridge, MA.

Porter, M. E. 1996. What is strategy? *Harvard Bus. Rev.* **74**(6) 61–78.

Rivkin, J. W. 2000. Imitation of complex strategies. *Management Sci.* **46** 824–844.

Rivkin, J. W. 2001. Reproducing knowledge: Replication without imitation at moderate complexity. *Organ. Sci.* **12** 274–293.

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

Schilling, M. A. 2000. Towards a general modular systems theory and its application to interfirm product modularity. *Acad. Management Rev.* **25** 312–334.

Schilling, M. A., C. Phelps. 2004. Interfirm collaboration networks: The impact of small world connectivity on firm innovation. Working paper, New York University, 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.

Siggelkow, N. 2002. Evolution toward fit. *Admin. Sci. Quart.* **47** 125–159.

Siggelkow, N., D. Levinthal. 2003. Temporarily divide to conquer: Centralized, decentralized, and reintegrated organizational approaches to exploration and adaptation. *Organ. Sci.* **14** 650–669.

Siggelkow, N., J. W. Rivkin. 2005. Speed and search: Designing organizations for turbulence and complexity. *Organ. Sci.* **16** 101–122.

Simon, H. A. 1962. The architecture of complexity. *Proc. Amer. Philos. Soc.* **106** 467–482.

Solow, D., A. Burnetas, T. Roeder, N. S. Greenspan. 1999a. Evolutionary consequences of selected locus-specific variations in epistasis and fitness contribution in Kauffman's NK model. *J. Theoret. Biol.* **196** 181–196.

Solow, D., A. Burnetas, M.-C. Tsai, N. S. Greenspan. 1999b. Understanding and attenuating the complexity catastrophe in Kauffman's NK a model of genome evolution. *Complexity* **5**(1) 53–66.

Sorenson, O. 2002. Interorganizational complexity and computation. J. A. C. Baum, ed. *Companion to Organizations*. Blackwell, Oxford, UK, 664–685.

Steward, D. V. 1981. The design structure system: A method for managing the design of complex systems. *IEEE Trans. Engrg. Management* **28**(3) 71–74.

Strogatz, S. H. 2001. Exploring complex networks. *Nature* **410** 268–276.

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.

Ulrich, K. T., S. D. Eppinger. 2007. *Product Design and Development*, 4th ed. McGraw-Hill, New York.

Uzzi, B., J. Spiro. 2005. Collaboration and creativity: The small world problem. *Amer. J. Sociol.* **111** 447–504.

Wasserman, S., K. Faust. 1994. *Social Network Analysis*. Cambridge University Press, New York.

Watts, D., S. H. Strogatz. 1998. Collective dynamics of "small-world" networks. *Nature* **393** 440–442.

Westhoff, F. H., B. V. Yarbrough, R. M. Yarbrough. 1996. Complexity, organization, and Stuart Kauffman's *The Origins of Order. J. Econom. Behav. Org.* **29** 1–25.

Wheelwright, S. C., K. B. Clark. 1992. *Revolutionizing Product Development: Quantum Leaps in Speed, Efficiency, and Quality*. Free Press, New York.