7-2007

Patterned Interactions in Complex Systems: Implications for Exploration

Jan W. Rivkin

Nicolaj Siggelkow
University of Pennsylvania

Recommended Citation

Rivkin, J. W., & Siggelkow, N. (2007). Patterned Interactions in Complex Systems: Implications for Exploration. *Management Science, 53* (7), 1068-1085. http://dx.doi.org/10.1287/mnsc.1060.0626

Patterned Interactions in Complex Systems: Implications for Exploration

Abstract

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.

Keywords

exploration, simulation model, pattern of interaction

Disciplines

Business Administration, Management, and Operations

**Patterned Interactions in Complex Systems:
Implications for Exploration***

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

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

* Special thanks to Howard Brenner for computer programming and to Monica Yan and Gavaskar Balasingam for research assistance. We are grateful to the Mack Center for Technological Innovation and the Division of Research of Harvard Business School for generous funding. Errors remain our own.

Patterned Interactions in Complex Systems:
Implications for Exploration

Abstract: 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, broader exploration is far more valuable 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 requires greater investment in broad exploration.

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 (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). But how 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 (CASs), 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 tradeoff – that are responsive and creative yet stable and orderly – neither frozen nor chaotic (e.g., Langton 1990; Kauffman 1993). Among the CAS 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).

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 epistatiscally affect them, essentially arbitrary interactions are possible" (Kauffman 1993: 41). In the context of organizational, technical, and social systems, however, recent empirical work has shown that interactions are often very patterned. Our paper exploits this newly-gained knowledge. Specifically, it examines how commonly observed patterns of interactions affect the proliferation of local optima and, accordingly, the appropriate amount of 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. Moreover, we identify easily observable characteristics of interaction patterns, beyond the overall degree of interaction, that allow one in many cases to look at two patterns of interaction and tell immediately which one generates more local optima and requires agents to explore more broadly for effective sets of choices. 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 1999; 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 the 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 ten 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 *power-law* 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, however, that interaction patterns affect other organizational phenomena as well, including the ability of a firm to adapt to environmental change, to find a valuable configuration of choices, to imitate the effective configurations of other firms, and to replicate one's own effective configurations (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). Our findings suggest how firms might design systems to be more readily searchable.

The paper is structured as follows: Section 2 describes in detail the ten interaction structures we analyze. Section 3 outlines how we create decision problems with these different underlying interaction structures, and it describes four types of organizations we employ to show the effects of interaction patterns on organizational search outcomes. The results in Section 4 characterize the local optima that arise from various interaction structures. 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 Section 6 shows. Section 7 concludes.

[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: 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."

2. Types of influence matrices

While the model we study is general enough to encompass a wide range of organizational, technical, and social systems, for expositional purpose we focus on firms as our system of interest. Following a long tradition in the organization literature (e.g., Learned, *et al.* 1961) that has gained energy recently from empirical, prescriptive, and computational studies (e.g., Siggelkow 2002; Porter 1996; Levinthal 1997), we conceptualize firms as systems of interdependent choices. Firms must make numerous decisions. Each firm must choose, for instance, how much to train its sales force, whether to field a broad product line or a narrow one, whether to pursue basic R&D or not, etc. A number of these decisions interact with each other. For instance, the value of having a well-trained sales force might increase as a firm broadens its product line.

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 payoff 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. Since each decision affects itself, all influence matrices have x's along their diagonal. Influence matrices can differ, however, in the total number of off-diagonal x's, i.e., in the number of interactions among the decisions, and in the patterns of these interactions. In the original NK set-up (Kauffman 1993), it was assumed that each decision is affected by exactly K other decisions, i.e., each row contained K off-diagonal x's. 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., Rivkin 2000), 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 re-arranged).[2] This reduces the number of patterns by a factor of N!. Yet, for N = 12 and K = 2, this still leaves $2.84*10^{17}$ different patterns. Given this vast space of possibilities, it is helpful to consider different *types* of interaction patterns. In particular, we focus on ten types that were culled from current work on networks, from studies of firms as 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, while an entry in row *i*, column *j*, would denote 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

[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.

system (element *i*). Table 1 examines all activity system maps that have been published in 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 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 ten different types of influence matrices we explore can be divided into two groups. For the first 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 subsequent 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 only depend on themselves.

Random. In a random influence matrix, exactly K *x*'s 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 set-up most commonly used in the organization literature (e.g., Westhoff, Yarbrough, and Yarbrough 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., decision 1 would be 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. Though not new, the notion of small-world networks has attracted renewed attention due to recent theoretical advances (Milgram 1967; 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, Yoo, and Baker 2003), memberships in underwriting syndicates (Baum, Shipilov, and Rowley 2003), firm-alliance networks (Schilling and Phelps 2004), career networks of artists (Uzzi, Spiro, and Delis 2002), 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.

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 back 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 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 many other decisions. 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 than others. For example, Figure 2 displays the DSM of an automobile brake system as reported by Black, Fine, and Sachs (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 imbalances 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 that contain 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 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 *x*'s 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 *x*'s that are already in that column. In particular, if D_j is the number of *x*'s in column *j* and S is the total number of *x*'s 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 a *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.

Power law. 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.) A number of networks have been shown to have degree distributions that follow a power law (Albert, Jeong, and Barabási 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 closely follows a power law.

We create a power law influence matrix in two steps. First, we initialize the matrix with *x*'s along the main diagonal. Second, in each column, M off-diagonal *x*'s 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 *x*'s in each influence

[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.

matrix equals to 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.

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: 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 *x*'s along the main diagonal, this matrix is created by adding *x*'s 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 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 *x*'s along the main diagonal, we create a hierarchical influence matrix by adding *x*'s 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 high-ranked decisions never affect low-ranked decisions, yet decision 1 is not necessarily the most central decision (as in the hierarchical structure). Starting with *x*'s along the main diagonal, this matrix is created by randomly adding *x*'s 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 (1999).

Dependent. The dependent influence matrix captures an instance in which a handful of decisions are affected by virtually every other decision the firm makes, yet those decisions exert little influence themselves. We construct such a matrix by transposing the centralized influence matrix. See Figure 1J.

3. Creation of performance landscapes and firms that search on them

Firms are assumed to make N binary decisions about how to configure their 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 $\mathbf{d} = d_1 d_2 \ldots d_N$ with each d_i either 0 or 1.

Once a particular influence matrix is chosen, the computer generates a performance landscape based on this influence matrix. Specifically, it assigns a payoff to each of the 2^N possible configurations of choices. The contribution C_i of each decision to overall firm value is affected by other decisions: $C_i = C_i(d_i;\text{ other } d_j\text{'s})$, where the identity of the "*j*'s" (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 relevant other d_j's, a contribution is drawn at random from a uniform U[0, 1] distribution. The overall payoff associated with a configuration is the average over the N contributions:

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

In Section 4, we will describe a number of properties of different performance landscapes. In this discussion, a key construct is the concept of a "local peak." A local peak is a configuration **d** such that no configuration **d'** exists that differs from **d** in only one decision and has higher performance than **d**.

In Section 6, we will examine how patterns of interactions affect the value of greater exploration. To do so, we will analyze the performance of two types of firms:

a) The *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.

b) The *high-exploration firm* allocates some of its search efforts to the consideration of more distant alternatives. Specifically, it considers each period an alternative **d'** that differs from **d** in terms of one *or two* decisions. For instance, in an N = 4 simulation, a high-exploration firm at 0000 might evaluate the alternative 0110.

Section 6 also illustrates the impact of interaction patterns on the relative performance of firms that differ in their organizational designs. This involves two other types of firms:

a) In the *decentralized firm*, decisions are split between two managers, A and B. Manager A is responsible for the first N/2 decisions, while manager B is responsible for the remaining N/2 decisions. In each period, each manager evaluates local alternatives for her "department." Continuing with the N = 4 example, in assessing any alternative **d**, manager A would consider $P_A(\mathbf{d}) = C_1(\mathbf{d}) + C_2(\mathbf{d})$, while manager B would consider $P_B(\mathbf{d}) = C_3(d) + C_4(\mathbf{d})$. In evaluating alternatives, each manager assumes that choices in the other department will not change. After evaluating alternatives, each manager implements the alternative that she finds best (or maintains the status quo if no evaluated alternative has higher performance).

b) In the *hierarchical firm*, decisions are again split up and department managers assess alternatives as in the decentralized firm. In this firm, however, each department manager is required to send her most preferred alternative to a CEO. The CEO, in turn, evaluates all possible combinations of departmental proposals and implements the combination that is best for the firm. For instance, if the status quo is 0000, and Manager A proposes 10 for decisions 1 and 2 and Manager B proposes 01 for decisions 3 and 4, then the CEO evaluates 0000, 1000, 0001, and 1001 and picks the configuration that has the highest value of $P(\cdot)$. This choice configuration would then be the starting point for search for both department managers in the next period.

4. Landscape characterization

We use each of the ten different influence matrices to generate performance landscapes and determine a number of topological characteristics of the resulting landscapes. For all simulations, we consider the case of N = 12. 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 set-up of 1 through 6.[4]

One of the key characteristics of a landscape is the number of local peaks it contains. 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 given a *fixed* value of K, i.e., holding the total number of interactions constant, that are to be found in landscapes with different underlying patterns of interactions.

Table 2 reports the number of local peaks for the random, the local, and the small-world matrices. Recall that the small-world set-up involves the parameter, p, the probability of non-local interactions, and it includes as special cases the local influence matrix ($p = 0$) and the random influence matrix ($p = 1$). Two patterns in Table 2 are noteworthy. First, as the interaction structure becomes increasingly random (i.e., as p increases), the number of local peaks declines. The change in number of local peaks is, however, rather modest – a decrease of 11-24% 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 -0.72 to -0.93 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 non-linear 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 non-linearly to p, the number of local peaks in performance landscapes based on small-world influence matrices behaves rather smoothly as p is changed. Since the landscape features we study (including ones discussed below) behave linearly in p, we will 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 table shows

[4] Note, for K > 5 (given N =12), it is not possible to construct diagonal, hierarchical, or power law influence matrices. As a result, since we are interested in comparisons across influence matrices, we do not investigate values larger than K = 6.

that *landscapes based on the same number of total interactions but different interaction patterns can contain dramatically different numbers of local peaks*. For instance, on K = 2 landscapes, the number of local peaks ranges from 3.4 for landscapes based on centralized influence matrices to 129.0 for landscapes based on dependent influence matrices. Similarly, for K = 5, the range is from 18.8 (centralized) to 242.0 (dependent).[5] Note that the high and low 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 interaction patterns than in landscapes that have dependent interaction patterns. Placing one low-exploration firm on every point of the landscape and letting them engage in incremental search until they have 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, in the K = 2 landscapes, 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 bottom 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 the be the most clustered and centralized landscapes the most dispersed. For K = 4 landscapes, the differences remain but are much smaller.

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

[5] As a benchmark, note that a fully interdependent influence matrix, with K = 11, has 315.1 local peaks on average.

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 significantly higher than the 53.0 local peaks in random landscapes with K = 4 (and N = 13). For the cartridge system, we find 57.6 local peaks on average, which compares to 53.3 local peaks found on random landscapes with K = 3 (and N = 14). 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, Miller, and Fontana (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 0 or 1, whichever produces higher performance, and then do the same for d_5, d_6, …, 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 $\{d_1\ d_2\ d_3\}$. 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).

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 tradeoffs. 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 *x*'s. 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 topology: the performance contribution from these choices alone form a smooth, single-peaked surface, as would arise from a 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, they change from one randomly drawn contribution 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.[6]

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, in order to escape poor local optima and 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 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 x 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:

Number of local peaks = 27.4 – 4.0 x number of uninfluential decisions + 19.8 x number of uninfluenced decisions
(t-stat = -4.9) (t-stat = 14.9)

[6] 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 x's only on the diagonal. If we fill one row of this influence matrix with x's, i.e., make one decision's contribution dependent on all other decisions, the number of local peaks increases sharply to 58. Now start with an influence matrix in which decision 1 is affected by decision 2, decision 2 is affected by decision 3, etc. This influence pattern, which contains no uninfluenced decisions, leads to a performance landscape with 9 local peaks. Filling one row of this influence matrix with x's increases the number of local peaks only to 39.

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. Performance consequences

In prior sections, we have asserted that the proliferation of local peaks increases the value of, and need for, broad exploration. Other researchers have 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. To do so, we conduct the following simulation: On each performance landscape we place a low-exploration firm, able to evaluate only the nearest alternatives to the status quo, and a high-exploration firm, able to evaluate alternatives that differ from the status quo in up to two decisions. Both firms are given the same, randomly chosen starting point and are allowed to search for better configurations for 300 periods. By then, both firms have exhausted opportunities for improvement. We calculate the performance of each firm relative to the global peak of the landscape, record the performance difference between the two firms, and then generate a new performance landscape with the same underlying influence matrix. For each influence matrix, we repeat this exercise 1,000 times. The average performance difference captures the value of broader exploration in the face of each type of influence matrix.

Those differences, reported in Table 4, reveal three striking patterns concerning the value of broader exploration. First, as one would expect in a set-up where exploration is made costless, high-exploration firms have significantly better performance than low-exploration firms for all levels of K and for all types of influence matrices. Second, within each type of influence matrix (i.e., for each column of the table), higher levels of K make broad exploration more valuable – a finding in line with the prior research mentioned above.

Third and crucially, the 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.[7] For K = 2, for instance, the centralized matrix produces only 3.4 local peaks on average and the value of broader exploration is merely 0.023, while the dependent matrix generates many more local peaks, 129.0, and the value of broader exploration is statistically significantly higher at 0.068. Indeed, the number of local peaks appears to do a better job than K at predicting the value of broad exploration. Casual inspection of Table 4 and the top panel of Table 3 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 benefit of broader search despite differences in K. Consider, for instance, the random matrix with K = 2 and the centralized matrix with K = 4; the power-law matrix with K = 4 and the centralized matrix with K = 6; and the diagonal matrix with K = 4 and the preferential attachment matrix with K = 6. 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 value of broad exploration, with a correlation coefficient of 0.827, and (b) the number of local peaks explains more of the variance in the value of broad exploration than does K alone. Overall, the results lend strong support to the notion that the marginal value of broader exploration is greater on landscapes with more local peaks.

Prior research efforts – both empirical classics like Burns and Stalker (1961) and kindred simulations like Rivkin and Siggelkow (2003) – have shown that organizational designs differ in how much exploration they encourage. That finding, coupled with this paper's argument about influence matrices and the value of exploration, suggests a speculation: a change in influence matrix may call for a change in

[7] The only exception to this relationship is the diagonal pattern, for which the benefit of broader search is somewhat lower than one would expect given its number of local peaks.

organizational design, even holding the total number of interactions constant. We support this speculation with an example. Specifically, we place one decentralized and one hierarchical firm on a random starting point on each landscape. In each firm, managers evaluate two alternatives per period. In the decentralized firm, each division manager implements the alternative she finds best for her department; in the hierarchical firm, each division manager sends her most preferred alternative to the CEO who then evaluates all possible combinations of proposals and implements the best combination she finds. We measure the performance of firms at the end of period 300 and, in Table 5, report the average performance differences from 1,000 landscapes of each type. A positive difference denotes that the decentralized firm outperforms the hierarchical firm; a negative difference implies that the hierarchical firm has the higher performance.

The column labeled "random" replicates the finding of Rivkin and Siggelkow (2003), which used random influence matrices. For low levels of K, the hierarchical firm significantly outperforms the decentralized firm, while for high levels of K, the decentralized firm significantly outperforms the hierarchical firm. Similar trends in the performance differences arise in all columns; as K increases, the benefit of the decentralized structure increases. These trends reflect the impact of organizational design on exploration: the autonomy of departmental managers in the decentralized firm permits broad exploration, while in a hierarchical firm, the oversight of a CEO – who refuses to ratify moves that reduce performance even temporarily – confines exploration. As K increases and local peaks proliferate, broad exploration becomes more valuable on the margin, and the benefit of the widely-exploring decentralized design increases.

The interesting new finding is again in the rows. For K = 2, for instance, the hierarchical structure leads to significantly higher performance than the decentralized structure when the underlying interaction structure is random (difference = -0.0070), but the performance ranking between these two firms reverses when the underlying interaction structure is dependent (difference = +0.0321). This result is consistent with the prior intuition: for the same level of K, landscapes with dependent interaction structures have many more local peaks than landscapes with random interaction structures (Table 3) so we would expect the benefit of the broader exploration generated by the decentralized structure to arise at lower levels of K if the firm is on a dependent landscape than if the firm is on a random landscape. This (admittedly stylized) example illustrates how the appropriate choice of organizational design may hinge on the pattern of interactions among decisions, not just the total number of interactions. Overall, the performance advantage of the decentralized firm, with its broad exploration, is closely correlated to the number of local peaks. Indeed, the correlation coefficient between the entries in the top panel of Table 3 (the number of local peaks) and the figures in Table 5 (the decentralized firm's performance advantage) is a large and highly significant 0.915.

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 (Rivkin and Siggelkow 2003), for instance.

The main thrust of the present 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 ten patterns (including the random benchmark) and examined the characteristics of landscapes produced by each. We found that underlying interaction patterns affect landscape topology 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 non-linear, changes in the number of local peaks as the probability of non-local 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.[8] This suggests a practical rule of

[8] 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* re-randomizes 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 in order to anticipate the number of local peaks. For instance, one might have to calculate

thumb for individuals who are deciding how much to invest in 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, managers are well advised to devote more resources to exploration and to adopt organizational designs that encourage wider exploration. Though 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

---

how concentrated influence is in, and on, a handful of decisions. This is a speculation that deserves investigation in future research.

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.

Our results also have system-design implications since 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. 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 system-wide 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.

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

A. Random

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

B. Local

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

C. Small-World

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

D. Block-Diagonal

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

E. PREFERENTIAL ATTACHMENT

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

F. POWER-LAW

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

G. CENTRALIZED

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

H. HIERARCHICAL

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

I. DIAGONAL

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

J. DEPENDENT

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

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, Fine, and Sachs (1990) as in: http://www.dsmweb.org/Scrap_book/Brake.htm

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 (1999), as in http://www.dsmweb.org/Scrap_book/Cartridge.htm

TABLE 1: CHARACTERISTICS OF ACTUAL DESIGN STRUCTURE MATRICES AND ACTIVITY SYSTEMS

Example	N	K*
Design structure matrices:		
Automobile break system (Black, Fine, and Sachs 1990)	13	3.8
Kodak cartridge development process (Ulrich and Eppinger 1999)	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 1999)	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.

TABLE 2: NUMBER OF LOCAL PEAKS FOR SMALL-WORLD INFLUENCE MATRICES

p =	0.0 (local)	0.05	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0 (random)
K = 1	5.0	5.1	4.8	4.6	4.7	4.2	4.5	4.5	4.6	4.7	4.3	4.2
K = 2	14.0	14.0	13.6	12.5	12.3	11.3	12.6	11.6	10.9	11.6	11.1	10.7
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	48.0	46.8	45.0	43.8	42.6	41.7	41.1	40.7	39.9	39.1	39.5	39.7
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	99.6	96.1	95.2	92.2	92.1	89.9	89.3	89.2	88.5	88.6	88.3	87.9

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

TABLE 3: CHARACTERISTICS OF LANDSCAPES BASED ON DIFFERENT TYPES OF INFLUENCE MATRICES

	Centralized	Hierarchical	Power law	Random	Preferential attachment	Local	Block-diagonal	Diagonal	Dependent
1. Average number of local peaks									
K = 1	1.9	2.3	4.5	4.8	5.0	5.0	5.6	7.6	55.4
K = 2	3.4	4.9	10.5	10.7	12.8	14.0	14.9	20.1	129.0
K = 3	6.2	9.7	20.5	23.1	23.4	27.5	33.6	43.6	177.4
K = 4	10.9	24.9	31.1	39.7	37.1	48.0	40.2	75.7	206.0
K = 5	18.8	64.3	56.3	60.8	57.7	71.7	82.9	114.4	242.0
K = 6	32.2			87.9	80.9	99.6	102.7		248.0
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
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

Each result in Panel 1 and 3 is an average over 200 landscapes. Each result in Panel 2 is an average over 50 landscapes.

TABLE 4: VALUE OF BROADER EXPLORATION

	Centralized	Hierarchical	Power law	Random	Preferential attachment	Local	Block-diagonal	Diagonal	Dependent
K = 1	0.012	0.017	0.021	0.024	0.026	0.024	0.030	0.029	0.041
K = 2	0.023	0.028	0.029	0.032	0.039	0.045	0.057	0.039	0.068
K = 3	0.029	0.038	0.040	0.039	0.040	0.051	0.068	0.047	0.077
K = 4	0.037	0.046	0.048	0.040	0.045	0.050	0.060	0.050	0.081
K = 5	0.040	0.058	0.047	0.045	0.049	0.053	0.072	0.062	0.079
K = 6	0.048			0.043	0.052	0.057	0.070		0.079

Each cell contains the performance difference in period 300 between a firm that engages in broad exploration (i.e., evaluates alternatives that differ in up to two decisions from the status quo) and a firm that engages in narrow 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 1,000 landscapes. All performance differences in this table are statistically significant with $p < 0.001$.

TABLE 5: PERFORMANCE ADVANTAGE OF DECENTRALIZED FIRM

	Centralized	Hierarchical	Power law	Random	Preferential attachment	Local	Block-diagonal	Diagonal	Dependent
K = 1	-0.0007	0.0006	-0.0016	**-0.0076**	*-0.0057*	-0.0032	-0.0012	0.0007	**0.0173**
K = 2	-0.0011	0.0045	*0.0065*	*-0.0070*	*-0.0058*	-0.0014	-0.0007	*0.0053*	**0.0321**
K = 3	0.0058	0.0000	0.0054	-0.0059	0.0019	0.0039	-0.0000	**0.0099**	**0.0362**
K = 4	0.0000	0.0043	*0.0070*	0.0022	*0.0070*	0.0049	-0.0001	**0.0131**	**0.0409**
K = 5	0.0027	**0.0207**	*0.0080*	*0.0089*	*0.0066*	*0.0065*	0.0023	**0.0169**	**0.0464**
K = 6	0.0012			**0.0134**	**0.0101**	**0.0144**	0.0010		**0.0419**

Each cell contains the performance difference in period 300 between a firm that is completely decentralized and a firm that employs an active hierarchy. Performance is measured relative to the highest performance possible in each landscape. Performance differences are averages over 1,000 landscapes. Differences in italics are significant at a level of 0.05; differences in bold are significant at a level of 0.01 or better.

References

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

Albert, R. and A.-L. Barabási. 2002. Statistical mechanics of complex networks. *Reviews of Modern Physics* **74** 47-97.

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

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

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

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

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

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

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

Benner, M. J. and M. L. Tushman. 2003. Exploitation, exploration, and process management: The productivity dilemma revisited. *Academy of Management Review* **28** 238-256.

Black, T. A., C. H. Fine, and E. M. Sachs. 1990. A method for systems design using precedence relationships: An application to automotive bake systems. Working paper # 3208, MIT Sloan School of Management, Cambridge, MA.

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

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

Dong, Q. 1999. Representing information flow and knowledge management in product design using the design structure matrix. ME Department SM Thesis, M.I.T., Cambridge, MA.

Eisenhardt, K. M. and S. L. Brown. 1999. Patching: Restitching business portfolios in dynamic markets. *Harvard Business Review* **77**(May-June) 72-82.

Eppinger, S. D., D. E. Whitney, R. P. Smith, and D. A. Gebala. 1994. A model-based method for organizing tasks in product development. *Research in Engineering Design* **6** 1-13.

Erdös, P. and A. Rényi . 1959. On random graphs. *Publicationes Mathematicae* **6** 290-297.

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

Fleming, L. 2001. Recombinant uncertainty in technological search. *Management Science* **47** 117-132.

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

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

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

Hannan, M. T. and J. Freeman. 1984. Structural inertia and organizational change. *American Sociological Review* **49** 149-164.

Kauffman, S. and S. Levin. 1987. Toward a general theory of adaptive walks on rugged landscapes. *Journal of Theoretical Biology* **128** 11-45.

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

Kauffman, S. A. and E. D. Weinberger. 1989. The NK model of rugged fitness landscapes and its application to maturation of the immune response. *Journal of Theoretical Biology* **141** 211-245.

Kogut, B. and G. Walker. 2001. The small world of Germany and the durability of national networks. *American Sociological Review* **66** 317-335.

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

Learned, E. P., C. R. Christensen, K. R. Andrews, and W. D. Guth. 1961. *Business Policy: Text and Cases.* Irwin, Homewood, IL.

Levinthal, D. A. 1997. Adaptation on rugged landscapes. *Management Science* **43** 934-950.

Levinthal, D. A. and M. Warglien. 1999. Landscape design: Designing for local action in complex worlds. *Organization Science* **10** 342-357.

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

March, J. G. 1991. Exploration and exploitation in organizational learning. *Organization Science* **2** 71-87.

Marengo, L., G. Dosi, P. Legrenzi, and C. Pasquali. 2000. The structure of problem-solving knowledge and the structure of organizations. *Industrial and 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. SDM Thesis, M.I.T., Cambridge, MA.

McKelvey, B. 1999. Avoiding complexity catastrophe in coevolutionary pockets: Strategies for rugged landscapes. *Organization Science* **10** 294-321.

Milgram, S. 1967. The small world problem. *Psychology Today* **2** 60-67.

Newman, M. E. J. 2001. The structure of scientific collaboration networks. *Proceedings of the National Academy of Sciences* **98** 404-409.

Osborne, S. M. 1993. Product development cycle time characterization through modeling of process iteration. Master's Thesis (Management/Engineering), M.I.T., Cambridge, MA.

Pimmler, T. U. and S. D. Eppinger. 1994. Integration analysis of product decompositions. working paper no. 3690-94, M.I.T., Sloan School of Management, Cambridge, MA.

Porter, M. E. 1996. What is strategy? *Harvard Business Review* **74**(6) 61-78.

Rivkin, J. W. 2000. Imitation of complex strategies. *Management Science* **46** 824-844.

Rivkin, J. W. 2001. Reproducing knowledge: Replication without imitation at moderate complexity. *Organization Science* **12** 274-293.

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

Schilling, M. A. 2000. Towards a general modular systems theory and its application to interfirm product modularity. *Academy of Management Review* **25** 312-334.

Schilling, M. A. and 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. *Academy of Management Journal* **44** 838-857.

Siggelkow, N. 2002. Evolution toward fit. *Administrative Science Quarterly* **47** 125-159.

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

Siggelkow, N. and J. W. Rivkin. 2004. Speed and search: Designing organizations for turbulence and complexity. working paper, Wharton School, University of Pennsylvania, Philadelphia.

Simon, H. A. 1962. The architecture of complexity. *Proceedings of the American Philosophical Society* **106** 467-482.

Sorenson, O. 2002. Interorganizational complexity and computation. J. A. C. Baum, (ed.) *Companion to Organizations.* Blackwell, Oxford.

Steward, D. V. 1981. The design structure system: A method for managing the design of complex systems. *IEEE Transactions on Engineering 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. and C. A. O'Reilly. 1996. Ambidextrous organizations: Managing evolutionary and revolutionary change. *California Management Review* **38**(4) 8-30.

Ulrich, K. T. and S. D. Eppinger. 1999. *Product Design and Development (2nd ed.).* McGraw-Hill, New York.

Uzzi, B., J. Spiro, and D. Delis. 2002. Emergence: The origin and evolution of career networks in the Broadway musical industry, 1877 to 1995. working paper, Kellogg School of Management, Northwestern University, Evanston, Ill.

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

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

Westhoff, F. H., B. V. Yarbrough, and R. M. Yarbrough. 1996. Complexity, organization, and Stuart Kauffman's The Origins of Order. *Journal of Economic Behavior and Organization* **29** 1-25.

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