# Organizational Sticking Points on NK Landscapes

JAN W. RIVKIN<sup>1</sup>

<sup>1</sup>Harvard Business School, 239 Morgan Hall, Boston, MA 02163; e-mail: jrivkin@hbs.edu

## NICOLAJ SIGGELKOW<sup>2</sup>

<sup>2</sup>Wharton School, 2017 Steinberg Hall-Dietrich Hall, Philadelphia, PA 19104; e-mail: siggelkow@wharton.upenn.edu

Received January 7, 2002; revised May 20, 2002; accepted May 22, 2002

Scholars studying human organizations have recently adopted the notion of fitness landscapes, a concept pioneered in the biological and physical sciences. Such scholars have generally assumed that organizations will migrate toward the local peaks of these landscapes, as biological and physical entities do. We use an agent-based simulation to show, to the contrary, that a hierarchical human organization may very well come to rest at a "sticking point" that is not a local peak on the fitness landscape of the overall organization. Three pervasive features of human organizations create the distinction between sticking points and local peaks: the delegation of choices to separate decision makers, interdependencies between the domains of those decision makers, and differences between local incentives and global incentives. Our results illustrate both that it is valuable to use tools developed to study one type of complex adaptive system in order to examine another type and that researchers must adapt the tools with care as they attempt to do so. © 2002 Wiley Periodicals, Inc.

Key Words: fitness landscape; local peak; sticking point; equilibrium; organizations

### **1. INTRODUCTION**

In the toolbox of complexity researchers, few implements have seen more use than the notion of a "fitness landscape." Evolutionary biologists [1–7] envision species that search a space of possible genotypes and, through a process of mutation, recombination, and selection, arrive atop a high peak in this space. Computer scientists and operations researchers [8–10] often approach difficult combinatorial optimization problems by constructing cost surfaces in the space of feasible solutions to the problems. Physicists [10,11] model the magnetic alloys known as spin glasses with landscapes that map energy levels to each possible configuration of spins.

In recent years, social scientists have adopted fitness landscapes to model human organizations and artifacts—firms [12–21], technologies [22–26], economies [12], and political systems [27] *inter alia*. The seminal work of Kauffman and his colleagues [3–5,7] has proven especially popular among social scientists. Kauffman's NK model provides a simple but flexible way to construct fitness landscapes in a

stochastic yet well-controlled manner. It allows the modeler to tune the degree of interaction among the "horizontal" dimensions of the landscape with ease.

In analyzing fitness landscapes, evolutionary biologists, computer scientists, operations researchers, and physicists have devoted special effort to characterizing the "local peaks" of such landscapes. A local peak is a configuration of components—a genotype, a combinatorial solution, or a spin configuration, for instance—that cannot be improved on by a change in a single component. Particularly when components interact richly in determining system fitness, such local peaks are numerous [6,7]. Researchers interested in fitness landscapes have striven to understand the number of local peaks, the number of local peaks attainable from a typical locale, the typical distances between local peaks, the average heights of the local peaks, the variance of the heights, and so forth [3–7,11,28].

This focus on local peaks is appropriate in the natural sciences. In biological systems, for instance, the variety that fuels evolution is often provided by mutations to single, randomly chosen genes. Moreover, selection is usually assumed to operate more quickly than mutation [2]. Consequently, the evolution of a species is accurately depicted as an undirected, upward climb on a fitness landscape. The climb comes to rest precisely atop one of the local peaks.

As social scientists have imported the notion of the fitness landscape, they have remained focused on the local peaks of landscapes. Of course, the human agents that social scientists study may engage in search routines more sophisticated than simple hill-climbing. They may undertake large-scale, intentional changes to imitate the actions of successful peers [18,20] or form simplified cognitive maps of the payoff surface [19], for instance. Landscape models built by social scientists incorporate such considerations. Even with enhanced search routines, however, human organizations in these models continue to come to rest, ultimately, on local peaks. At first glance, a focus on local peaks seems reasonable in the context of human organizations. After all, why should an organization stop its search efforts before it reaches a local peak on the fitness landscape of the overall organization, that is, when local opportunities for improvement remain untapped?

In this article, we argue that the focus on local peaks may be misplaced for hierarchical human organizations. *Such an organization may very well come to rest at a "sticking point" that is not a local peak on the fitness landscape of the organization as a whole.* To make this point, we develop an agent-based simulation model in which firms search on NK landscapes for effective sets of business decisions. Our key innovation is to give each firm an internal decision-making structure that resembles the structure reported in the extensive empirical literature on organizations. The structure includes an allocation of decisions to low-level "departmental" managers, an incentive system that guides departmental managers in making proposals to senior management, and a mechanism for senior management to review the proposed actions of subordinates. For firms with various internal structures, we characterize sets of "sticking points." A sticking point is a configuration of choices such that, once the firm arrives at the configuration, the firm will never deviate from it. In essence, a sticking point is an equilibrium in the game played among a firm's subordinate and senior managers.

The main result of the article is that, for the typical firm, the set of sticking points is neither equal to, nor a subset or superset of, the set of local peaks on the overall organization's landscape. First, local peaks might not be sticking points. That is, a firm might continue to move on the landscape even after it happens upon a local peak. Suppose, for instance, that the decisions that constitute a local peak imply poor performance for one of the firm's departments. A self-interested subordinate in charge of that department may be able to enact a parochially attractive change that forces the firm off the local peak and damages the performance of the firm as a whole. Second, firms might get stuck on locations that are not local peaks. It is even possible to have situations in which every member of management would like to make an incremental change, yet no change occurs. This might arise, for example, when (a) each subordinate wishes to make a parochially attractive move but senior management rejects the move because it undermines overall performance and (b) an incremental move that would be desirable for the total organization is not in the parochial interest of the relevant subordinate and thus is never proposed to senior management. One might initially think that the distinction between sticking points and local peaks arises solely because final decisions are made by managers who pursue something other than the fitness of the organization as a whole. In fact, the distinction persists even when the organization includes a senior manager who explicitly looks out for total fitness and must ratify every move.

Three pervasive features of human organizations drive the main result: the delegation of decisions to distinct managers, the interdependencies between those managers, and the deviation of local incentives from global incentives. Below, we discuss the implications of and rationale for each feature.

# 2. MODEL DESCRIPTION

Our model is conventional in how it constructs fitness landscapes, but innovative in how the firms that search for high ground on those landscapes make decisions.<sup>1</sup>

#### 2.1. Landscape Construction

Each firm in our model makes *N* decisions,  $\{d_1 \ d_2 \ \dots \ d_N\} = \mathbf{d}$ . For each decision, the firm is assumed to have two options,

which we denote by 0 and 1. For instance,  $d_1$  might be the decision of whether to offer a narrow range of products  $(d_1$ = 0) or a wide range of products  $(d_1 = 1)$ , whereas  $d_2$  might be the decision of whether to invest in flexible machine tools  $(d_2 = 1)$  or not  $(d_2 = 0)$ . We refer to the *N*-digit string **d** as the firm's choice configuration. We assign a fitness value to each of the  $2^N$  choice configurations as Kauffman [3–5,7] does: The fitness of each choice configuration **d** is assumed to be the average of the contributions that each individual decision makes. The contribution of each individual decision, in turn, is affected by how this decision is resolved (0 or 1) and by the resolution of K randomly chosen other decisions  $\mathbf{d}_{-i}$  (*K* is an integer between 0 and N - 1.) Denote with  $C_i(d_i, \mathbf{d}_{-i})$  the contribution of decision *i* given choice  $d_i$  for that decision and choices  $\mathbf{d}_{-i}$  for the other relevant decisions. Each of the  $2^{K+1}$  possible contributions of each decision is determined by a draw from a uniform distribution between 0 and 1. The fitness of each choice configuration **d** is then given by

$$P(\mathbf{d}) = \frac{\sum_{i=1}^{N} C_i (d_i; \mathbf{d}_{-i})}{N}$$

# 2.2. Organizational Decision-Making

Conventional hill-climbing search of the fitness landscape would be an appropriate model of managerial decisionmaking if all *N* decisions were made by a single, boundedly rational executive devoted to maximization of fitness. The extensive literature on organizational structure and decisionmaking, however, paints a very different picture of how decisions are made in real organizations. In reality, senior managers delegate primary responsibility for most decisions to an array of subordinates [29] who have considerable latitude to pursue their own parochial interests [30,31]. Senior managers retain the right to review the proposals that bubble up from the hierarchy below them, but they exercise greatest influence by controlling "rules of the game" such as incentive schemes [32].<sup>2</sup> We try to model the essence of this more realistic picture as follows.

## Allocation of Decisions

Each modeled firm has a management team consisting of a CEO and two subordinate managers: A and B. Manager A has primary responsibility for a subset of the *N* decisions, and Manager B has responsibility for the complementary subset. We use a string of a's and b's to designate a particular allocation of decisions. In a simulation with N = 6, for instance, the allocation abbbba would indicate that Manager A has responsibility for decisions 1 and 6, whereas Manager B oversees decisions 2, 3, 4, and 5.

#### Subordinate search

Search proceeds in a series of periods. In each period, each subordinate manager reconsiders the configuration of

choices in his department. Specifically, he compares the status quo set of choices to all "local" alternatives that involve a change in one of his decisions. In the N = 6 example mentioned above, suppose that the configuration of firm choices is 100111 at the beginning of a period. Thus, the configuration of choices in Manager B's department is 0011, and he contemplates a move to each of the four adjacent alternatives: 1011, 0111, 0001, and 0010. The configuration in Manager A's department is 11, and he evaluates 01 and 10.

#### Incentives

Each manager ranks the alternatives and the status quo from most preferred to least. In assessing alternatives, he puts primary weight on the performance of his department, but he may also consider the effects of his changes on the other department's performance. INCENT, a parameter that ranges from 0 to 1, captures the degree to which the subordinate cares about the ramifications of his actions on the other department. INCENT = 0 implies that each manager considers only effects within his department; it may reflect, for instance, a firm in which managers are paid strictly on the basis of local business unit profitability. INCENT = 1 implies that each manager is equally concerned with effects outside his department and genuinely wants to maximize firm-wide payoff; this may reflect a firm in which divisional officers' compensation is tied solely to overall corporate performance. Continuing the example above, subordinate Manager B evaluates a choice configuration **d** as follows:

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

In evaluating alternatives, each manager assumes that decisions in the other department will remain unchanged.

Thus, the subordinate is engaged in simple hill-climbing on a departmental "subscape"—that is, a landscape comprised of the decisions within his department. The subscapes are coupled to the degree that decisions in one department influence the contributions of decisions in the other. As decisions in one department change, the subscape of the other deforms [16] or, in the imagery of Kauffman and Johnsen's NKCS model [33], "dances." The parameter INCENT determines whether the elevations of the subscapes reflect the contributions of each department in isolation (INCENT = 0) or the elevations on the composite firm's landscape (INCENT = 1).

## Vertical hierarchy and proposals to the CEO

After evaluating alternatives and the status quo in his department, each subordinate sends up to the CEO the P proposals that he most prefers. A low level of *P* reflects a firm in which managers are expected to, or permitted to, narrow down options a great deal before turning to superiors. A high level of *P* reflects a firm in which senior managers want to review many alternatives themselves.

We consider two types of CEOs. The first type simply rubberstamps the subordinates' proposals; that is, she accepts Manager A's favorite proposal and Manager B's favorite without review. The second type of CEO actively exercises discretion. She considers all possible combinations of proposals from below, assesses the composite alternatives in light of the interests of the firm as a whole, compares them to the composite status quo, and selects the combination that yields the best payoff for the firm. This combination is implemented and becomes the launching point for further search in the next period.

One can conceive of the CEO and her subordinates as being engaged in coevolutionary search processes. Subordinates search on subscapes whose elevations resemble the elevations of the overall firm's landscape to a degree determined by their incentives. Each subordinate's subscape deforms as changes are enacted in the other department, and the paths available to each may be constrained by the CEO. The CEO searches on a landscape whose elevation reflects the performance of the total organization, but her movement is constrained by the proposals she receives. Modeled firms differ from one another in their organizational arrangements: the grouping of decisions into departments, the amount of information conveyed to senior management (P), the degree to which the CEO acts on that information (rubberstamping vs. active), and the incentives that managers have to consider effects beyond their domains (INCENT).

In a special case, our model becomes very similar to an important prior model of coevolution: Kauffman and Johnsen use their NKCS model to examine the coevolution of S species with internal epistatic interaction parameterized by K and a degree of cross-species interaction tuned by C. Each species engages in hill-climbing on its own landscape, and each species' landscape "dances" as other species evolve. When the CEO rubberstamps proposals and INCENT = 0, our subordinates engage in a similar uncoordinated, yet interdependent search. Outside of this special case, however, our model differs substantially from the NKCS model; it adds pivotal features that reflect human organization and that have limited parallels in biological coevolution. When active, the CEO implements wellmatched proposals from subordinates and protects the interests of the firm as a whole. When INCENT > 0, each subordinate cares about the impact of his proposals on the other department. Appropriate for a biological context, the NKCS model does not incorporate such a vertical hierarchy or incentive system.

#### **3. ORGANIZATIONAL STICKING POINTS**

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

In landscape models, it is common to think of firms or other searching entities as getting stuck on local peaks and only on such peaks; that is, the sets of sticking points and local peaks are identical. In our model, firm choices are the result of search efforts by three distinct agents, two of whom may pursue interests other than total firm fitness and all of whom shape and constrain the search of others. In this section, we ask whether the sticking points that arise from this process correspond to the local peaks of the overall organization. In short, we find that sticking points and local peaks often differ. They are identical only under special assumptions concerning decision allocation, interactions, incentives, and information flow.

#### 3.1. An Example

A simple example with N = 4, K = 3, and decision allocation aabb illustrates that (a) not every local peak is a sticking point and (b) a sticking point might not be a local peak for any agent. Suppose the contributions of each subordinate's choices are those shown in Table 1.

(a) Not every local peak is a sticking point. First consider the case in which INCENT = 0 and the CEO rubberstamps the proposals made by subordinates. On the fitness landscape shown in Table 1, the choice configuration 1111 is a local peak for the firm. (Indeed, it is the global maximum of the landscape with total performance 1.00.) A firm that attains this optimal set of choices will not maintain it, however. Manager A will propose a change from 11 to 01 for his department; such a change promises to raise his department's contribution from 0.50 to 0.60. Likewise, Manager B will propose a change from 11 to 01 for his department. Because the CEO rubberstamps the proposals, the firm will move away from the local peak to choice configuration 0101. From this choice configuration, the firm will not move because neither manager can find a local improvement given that the other manager does not change his decision. As a result, 0101 is a sticking point for the firm. In this case, 01 is a local peak on each manager's subscape.<sup>3</sup>

(b) Not every sticking point is a local peak. Next consider a firm in which INCENT remains 0, but the CEO actively reviews proposals. In particular, each manager sends up his most preferred proposal (P = 1). Suppose a firm is initially

# TABLE 1

Sticking Point Example

A's Choices for Decisions 1 and 2	B's Choices for Decisions 3 and 4	A's Contribution $[(C_1 + C_2)/4]$	B's Contribution $[(C_3 + C_4)/4]$	Total Performance $[(C_1 + C_2 + C_3 + C_4)/4]$
00	00	0.25	0.40	0.65
10	00	0.30	0.20	0.50
01	00	0.20	0.50	0.70
11	00	0.10	0.10	0.20
00	01	0.05	0.45	0.50
10	01	0.20	0.10	0.30
01	01	0.15	0.55	0.70
11	01	0.10	0.60	0.70
00	10	0.10	0.10	0.20
10	10	0.10	0.10	0.20
01	10	0.10	0.10	0.20
11	10	0.10	0.10	0.20
00	11	0.10	0.10	0.20
10	11	0.10	0.10	0.20
01	11	0.60	0.10	0.70
11	11	0.50	0.50	1.00

located at 0000 with performance 0.65. Clearly 0000 is not a local peak for the firm as a whole because 0100, just one change away, yields a higher total performance of 0.70. Manager A, however, will never propose a change from 00 to 01 for his department. Evaluating his options 01 and 10 (given Manager B's current choices of 00), Manager A prefers option 10 and proposes 10 to the CEO. Manager B, evaluating his options, most prefers option 01 and proposes this option to the CEO. The CEO, in turn, evaluates the composite options, 1001, 0001, and 1000, and finds the status-quo 0000 superior to all of the alternatives. Consequently, the firm will not change its choice configuration; 0000 is a sticking point for the firm. In this case, 00 is not a local peak in either manager's subscape. Each manager wants to make an incremental move that would increase his department's performance, but the CEO, who takes into account the externalities that the managers exert on each other, forbids these moves.

## 3.2. Simulation Results

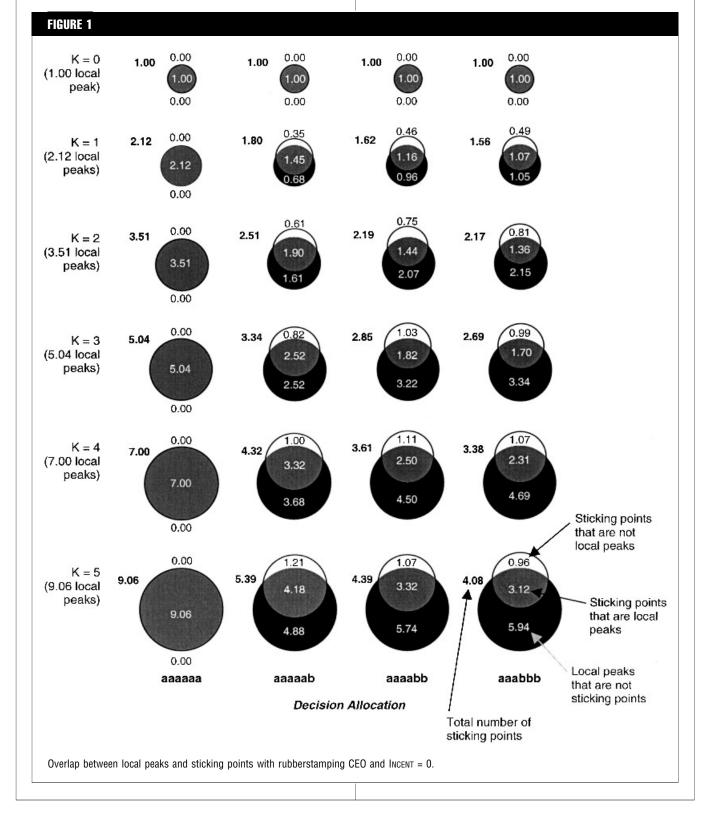
To shed further light on the differences between sticking points and local peaks, we identify the full sets of sticking points and local peaks on a large number of landscapes under various organizational arrangements, and we show the overlap between the sets. In the following analyses, each firm faces six decisions (N = 6). Each reported result is an average over 500 landscapes; because our interest is in sticking points and local peaks, we average over landscapes rather than over walks on landscapes.

We begin with a situation in which each firm has a rubberstamping CEO and INCENT = 1. For such a firm, the sets of sticking points and local peaks are identical regardless of the allocation of decisions between Managers A and B (i.e., aaaaaa, aaaaab, aaaabb, or aaabbb). A simple reductio ad absurdum argument shows why the sticking points and local peaks are the same. First, assume for a moment that the firm is stuck with a choice configuration that is *not* a local peak. Then there exists at least one alternative configuration involving a single change that improves the performance of the firm as a whole. Because each subordinate judges alternatives in light of firm-wide implications (INCENT = 1) and takes the existing choices of the other subordinate as given, the subordinate in charge of the performance-improving decision will propose the change, and the rubberstamping CEO will approve and enact the proposal. The firm moves, contradicting the initial assumption that the firm's starting configuration was a sticking point. Thus every sticking point must be a local peak.

Next assume that there exists a local peak that is not a sticking point. Because the local peak is not a sticking point, some manager must be eager to move off this local peak. For this to be so, there must be an alternative configuration involving a single change from the local peak that improves the performance of the firm as a whole (because INCENT = 1). This contradicts the assumption that the firm was initially located on a local peak. Hence, to avoid a contradiction, it must be true that every local peak is a sticking point.

So far, the sets of sticking points and local peaks are identical. A seemingly small change in organizational arrangements, however, causes the two sets to diverge substantially. Suppose that firms continue to have rubber-stamping CEOs, but now INCENT = 0: subordinates are free to

pursue parochial interests. Figure 1 displays Venn diagrams for sticking points and local peaks for various levels of K and allocations of decisions. The upper crescent, lower crescent, and intersection in each diagram show, respectively, the number of sticking points that are not local peaks, the number of local peaks that are not sticking points, and the number of configurations that are both sticking points and local peaks.



The total number of local peaks, reported on the far left edge of the diagram, displays a result familiar from the work of Kauffman [7]: as interactions among decisions proliferate, so do local peaks. Sticking points remain identical to local peaks so long as one subordinate has responsibility for all decisions (aaaaaa, the left column of Venn diagrams) or no interactions cross-departmental borders (K = 0, the top row). However, once decisions are divided between subordinates whose fates are intertwined, sticking points and local peaks become distinct. A local peak fails to be sticky if either subordinate manager perceives and pursues a change that is in his parochial interests even though it undermines overall fitness. This is especially common when crossdepartment interactions are prevalent: on high-K landscapes and for firms that give each manager purview over three decisions. For instance, for an organization that assigns three decisions to each manager, only 34% of local peaks are sticking points on a K = 5 landscape. In addition, sticking points that are not local peaks emerge. Such points are truces between the two subordinates in which changes that would benefit the firm as a whole never get proposed because they are not in the parochial interest of the manager who would have to suggest them. For example, for the same organization as above, 24% of its sticking points are not local peaks.

The addition of an active CEO to whom subordinates send their most preferred proposal (P = 1) causes a proliferation of sticking points and makes local peaks very nearly a proper subset of the set of sticking points (Figure 2). Three distinct effects account for this impact of the active CEO. First, because she evaluates alternatives from the perspective of the firm as a whole and always has the prerogative to "freeze" a firm at its old configuration, she takes the local peaks that were not sticking points in Figure 1 and makes virtually all of them sticky. (The next paragraph explains why we say "virtually all.") Second, she expands the set of sticking points that are not local peaks. To be a member of this set in Figure 1, a configuration had to offer no parochially attractive alternative to Manager A or to Manager B. Such configurations continue to be sticking points in Figure 2, but they are now joined by configurations that offer parochially attractive but globally harmful alternatives. The CEO, protector of global performance, rejects movement from such configurations to such alternatives. (The configuration 0000 discussed in Section 3.1.b provides an example.) The interplay between parochial interests and the prerogative of the CEO to reject parochial moves leads to new sticking points that are not local peaks. We call this second effect of the active CEO the "parochialism-versus-prerogative effect." For instance, for firms that assign three decisions to each manager, more than 40% of all sticking points are not local peaks for any K > 0.

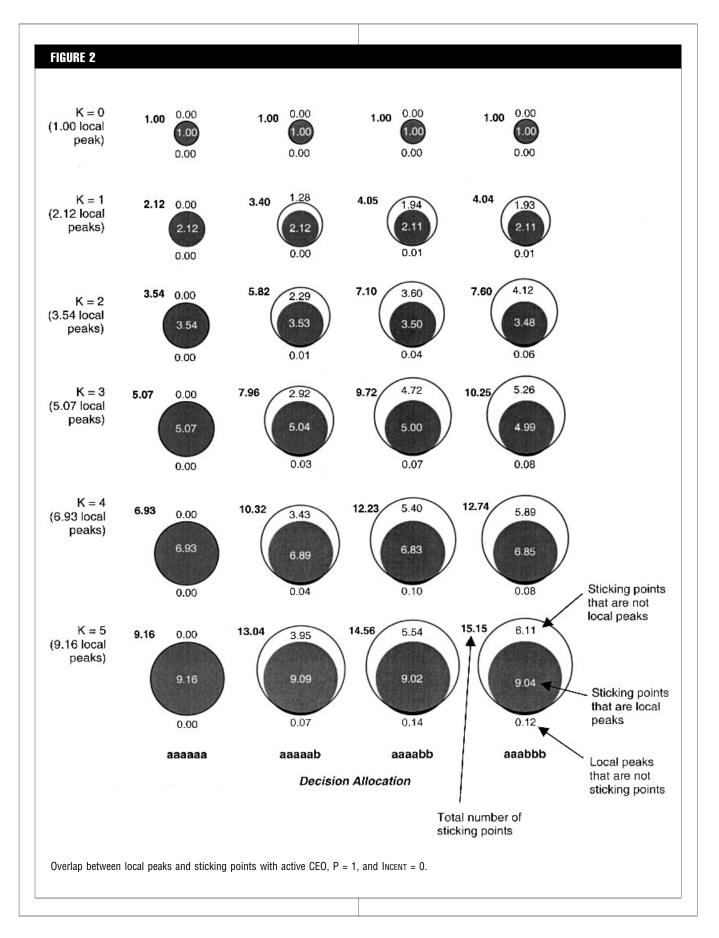
The third effect of the active CEO is to take some local peaks and make them "unsticky." Because the CEO receives

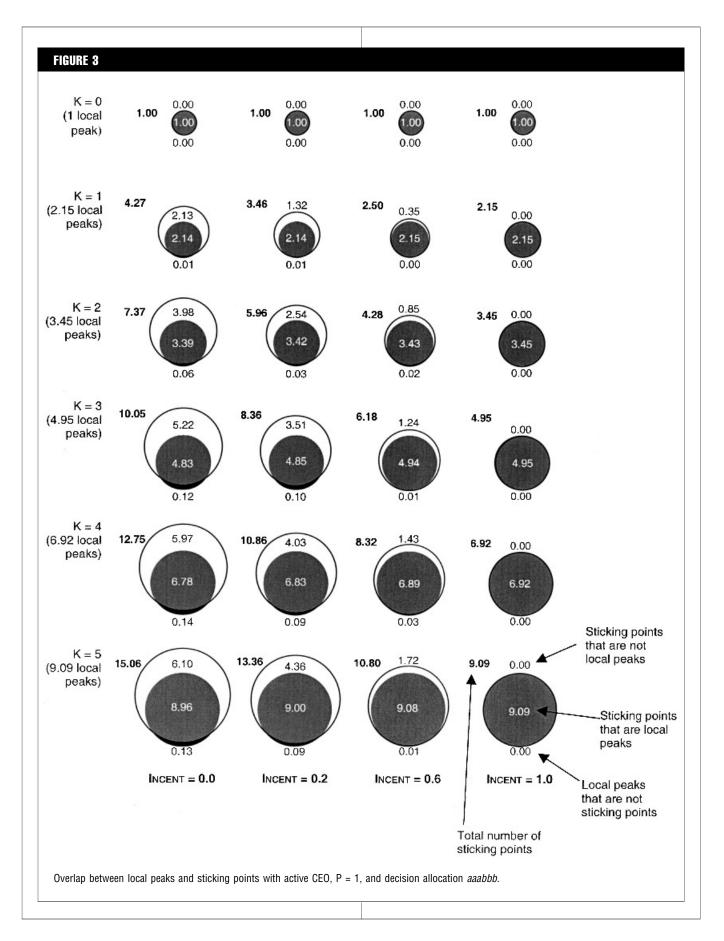
single-change proposals from each department, she is able to evaluate a limited number of alternatives that involve simultaneous change in two decisions—composite alternatives that were never evaluated under the rubberstamping CEO. If a local peak can be improved upon by the simultaneous change of one choice in each department, the peak will no longer constitute a sticking point. This "recombinatory effect" of the active CEO accounts for the very small number of local peaks are not sticking points in Figure 2.

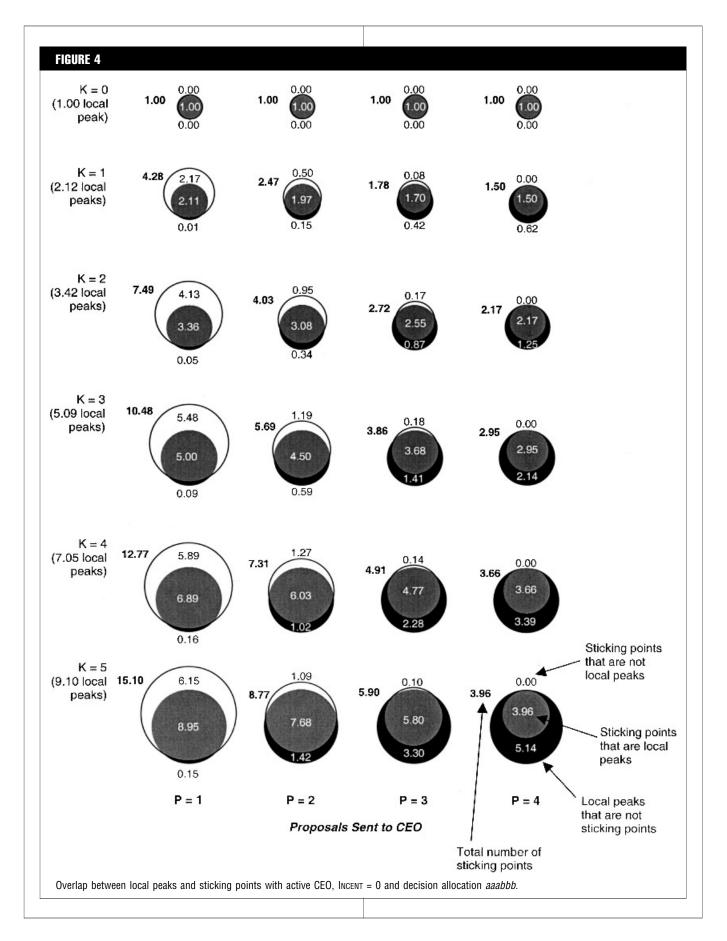
The difference between sticking points and local peaks in Figure 2 depends crucially on the fact that department managers pursue interests other than firm-level performance. Keeping the active CEO in place but allowing INCENT to rise from 0 to 1 (Figure 3), we see parochialism disappear and, with it, the parochialism-versus-prerogative effect and the sticking points that are not local peaks. Indeed at INCENT = 1, the sets of sticking points and local peaks become identical, by the following reasoning. If a firm with INCENT = 1 is on a sticking point, neither manager can find a single-change move that improves performance for the firm as a whole. Because the two managers together cover all possible single-change moves, every sticking point is a local peak. Similarly, if such a firm is on a local peak, neither manager is able to find a single-change move that improves performance, neither subordinate proposes a change to the status quo, and the CEO with P = 1 reviews nothing but an endorsement of the status quo. Hence the firm sticks at any local peak. In short, the parochialism which creates sticking points that are not local peaks in Figure 2 disappears when subordinates have high INCENT.<sup>4</sup>

The effect of parochialism also disappears when the flow of information through the hierarchy increases. The pursuit of parochial objectives depends, in part, on the ability of subordinates to hide globally attractive alternatives from the CEO. This becomes difficult when *P* is high. In Figure 4, we return to INCENT = 0 but force subordinates to submit more and more proposals to the CEO. As P increases, the number of sticking points that are not local peaks declines. In the extreme, if a subordinate is forced to send all alternatives to the CEO (P = 4), the CEO personally sees and assesses all alternatives that involve a change in any single decision. As a result, each sticking point has to be a local peak for such a firm. At the same time, with high P, the recombinatory effect of the CEO comes to the fore. With P = 4, the CEO not only sees all alternatives that involve a change to a single decision, but also reviews all two-move alternatives that involve a single change in each department. As a result, there exist many local peaks on which the firm does not get stuck. For instance, on a landscape with K = 5,56% of local peaks are not sticking points when P = 4.

Although the effect of parochialism, that is, the creation of sticking points that are not local peaks, is eliminated by firm-level incentives ( $I_{NCENT} = 1$ ), the recombinatory effect of an active CEO remains as long as more than just the most







preferred alternative is sent to the CEO. As a result, not every local peak is a sticking point if P > 1, even if INCENT = 1. In the extreme, if all alternatives are sent to the CEO (P = 4), the level of INCENT becomes irrelevant, as subordinates do not rule out any alternatives based on their evaluations (results available on request).

# 4. DISCUSSION

In sum, we see a difference between the set of local peaks and the set of sticking points when:

- 1. primary responsibility for decision-making is divided between the two subordinates (i.e., *not* aaaaaa); and
- 2. decisions in one department affect the contributions of decisions in the other (K > 0); and either
- 3a. the interests of departmental managers are not perfectly aligned with the interests of the firm as a whole (INCENT < 1); or</li>
- 3b. if the interests of departmental managers are perfectly aligned with the interests of the firm as a whole (INCENT = 1), the CEO receives more than one proposal from each subordinate (P > 1).

We argue that this set of conditions is typically met in all but the smallest organizations. The cognitive limits of human managers force the first condition to be true in organizations of even modest size [34]. In line with the second condition, a burgeoning literature in economics and management demonstrates that managerial decisions commonly interact with one another [35-38]. Supporting condition 3a, extensive research on the "principal-agent" problem shows that interests often diverge in different parts of an organization [39]. Even if one could make all managers pursue the fitness of the firm as a whole (INCENT = 1), it is not clear than one would want to do so. Rewards based on local performance may, for instance, evoke greater effort [40-42]. A common tenet of organizational design-that one should make a manager accountable only for decisions for which he is responsible [43-44]-prescribes a value of INCENT well below 1. Concerning condition 3b, we acknowledge that, because of communication costs and limits on the processing power of senior managers, firms may limit the flow of information in a hierarchy. Rarely, however, would a subordinate be allowed to present only his favorite option to senior management.

If this set of conditions is often met, it may be common for a hierarchical human organization to become stuck at a point that is not a local peak on the fitness landscape of the overall organization. A focus on local peaks of whole entities, sensible for biological and other systems, may not be appropriate for human organizations. This raises a general dilemma concerning the study of complex adaptive systems. On one hand, it is powerful to transport intellectual goods across discipline borders—to use tools developed in one domain to deepen our understanding of others. The application of fitness landscapes, for instance, has clearly enhanced our understanding of how human organizations search for good solutions to the complex problems they face [12–27]. On the other hand, models and concepts developed in one context must be tailored with great care to fit other settings. Features such as hierarchy, authority, and incentives, which may not have counterparts in biological or physical systems, are central to human organizations and need to be incorporated into researchers' tools. We hope that our work illustrates one way to incorporate such features.

A natural extension of our work is to characterize sticking points more fully, for example, to document their average heights under various organizational arrangements. Local peaks need not be especially effective sets of choices, especially when *K* is high [7], and we find that local peaks may be higher *or* lower than sticking points on average, depending on the particulars of landscapes and organizations. In addition, we find in related research [45] that many modeled firms wander incessantly from one configuration to another and never reach any equilibrium whatsoever. Clearly it is important not only to characterize sticking points, but also to examine whether firms actually come to rest at such points.

Our work could also be extended by exploring the relationship between sticking points and local peaks for firms that face more decisions, consist of more than two layers of hierarchy, and have more than two subordinates per boss all features of many real organizations. We suspect that an increase in the number of searching actors will only amplify the difference between sticking points and overall local peaks. Such an increase will make it more likely that some agent will be able and eager to block a move toward a local peak or to enact a move away from a peak.

Our simulations to date have taken the organizational structure of each firm as fixed. In reality, of course, firms can and do adjust their structures. Subordinates who act against the interest of the firm, for instance, may be denied responsibility for important decisions or even be replaced. An interesting extension of our work is to let organizational structure itself evolve, by shifting decision-making responsibility from one subordinate to another, by adjusting the flow of information to the CEO or her level of activity, or by altering the incentive system. A natural question is whether organizations can evolve through incremental change toward effective designs—that is, whether the landscape in the space of possible organizational structures is smooth or rugged.

#### **ACKNOWLEDGMENTS**

The authors thank two anonymous referees for helpful comments, Howard Brenner for computer programming heroics, and David Ma and Kevin Wang for excellent research assistance. They are grateful to the Mack Center for Technological Innovation and the Division of Research of Harvard Business School for generous funding.

#### REFERENCES

- 1. Wright, S. Evolution in Mendelian populations. Genetics 1931, 16, 97–159.
- 2. Gillespie, J.H. Molecular evolution over the mutational landscape. Evolution 1984, 38, 1116–1129.
- 3. Kauffman, S.; Levin, S. Toward a general theory of adaptive walks on rugged landscapes, J Theor Biol 1987, 128, 11-45.
- 4. Kauffman, S.A.; Weinberger, E.D.; Perelson, A.S. Maturation of the immune response via adaptive walks on affinity landscapes. In: Theoretical Immunology, Part I; Santa Fe Institute Studies in the Sciences of Complexity. Perelson, A.S., Ed.; Addison-Wesley: New York, 1988.
- 5. Kauffman, S.A.; Weinberger, E.D. The N-k model of rugged fitness landscapes and its application to maturation of the immune response. J Theor Biol 1989, 141, 211–245.
- 6. Macken, C.A.; Hagan, P.S.; Perelson, A.S. Evolutionary walks on rugged landscapes. SIAM J Appl Math 1991, 51, 799-827.
- 7. Kauffman, S.A. The Origins of Order: Self-organization and Selection in Evolution. Oxford University Press: New York, 1993.
- 8. Holland, J. Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence. University of Michigan Press: Ann Arbor, 1975.
- 9. Kirkpatrick, S.; Gelatt, Jr., C.D.; Vecchi, M.P. Optimization by simulated annealing. Science 1983, 220, 671-680.
- 10. Palmer, R. Statistical mechanics approaches to complex optimization problems. In: The Economy as an Evolving Complex System. Anderson, P.W.; Arrow, K.J.; Pines D., Eds.; Addison-Wesley: Reading, 1988.
- 11. Weinberger, E.D. Local properties of Kauffman's N-k model: A tunably rugged energy landscape. Phys Rev A 1991, 44, 6399-6413.
- 12. Kauffman, S.A. At home in the universe: The search for laws of self-organization and complexity. Oxford University Press: New York, 1995.
- Westhoff, F.H.; Yarbrough, B.V.; Yarbrough, R.M. Complexity, organization, and Stuart Kauffman's The Origins of Order. J Econ Behav Org 1996, 29, 1–25.
- 14. Levinthal, D.A. Adaptation on rugged landscapes. Management Sci 1997, 43, 934-950.
- 15. Sorenson, O.J. The complexity catastrophe in the computer industry: Interdependence and adaptability in organizational evolution. Ph.D. Dissertation, Stanford University: Stanford, 1997.
- 16. Levinthal, D.A.; Warglien, M. Landscape design: Designing for local action in complex worlds. Org Sci 1999, 10, 342-357.
- 17. McKelvey, B. Avoiding complexity catastrophe in coevolutionary pockets: Strategies for rugged landscapes. Org Sci 1999, 10, 294-321.
- 18. Rivkin, J.W. Imitation of complex strategies. Management Sci 2000, 46, 824-844.
- 19. Gavetti, G.; Levinthal, D. Looking forward and looking backward: Cognitive and experiential search. Admin Sci Quart 2000, 45, 113–137.
- 20. Rivkin, J.W. Reproducing knowledge: Replication without imitation at moderate complexity. Org Sci 2001, 12, 274-293.
- 21. Levitan, B.; Lobo, J.; Kauffman, S.A.; Schuler, R. Optimal organizational size in a stochastic environment with externalities. Working paper 99-04-024. Santa Fe Institute: 1999.
- 22. Kauffman, S.A.; Macready, W.G. Technological evolution and adaptive organizations. Complexity 1995, 1, 26-43.
- 23. Lobo, J.; Macready, W.G. Landscapes: A natural extension of search theory. Working paper 99-05-037. Santa Fe Institute, 1999.
- 24. Baldwin, C.Y.; Clark, K.B. Design Rules: The Power of Modularity. MIT Press: Cambridge, 2000.
- 25. Kauffman, S.A.; Lobo, J.; Macready, W.G. Optimal search on a technology landscape. J Econ Behav Org 2000, 43, 141-166.
- 26. Fleming, L.; Sorenson, O. Technology as a complex adaptive system: Evidence from patent data. Res Policy 2001, 30, 1019–1039.
- 27. Kollman, K.; Miller, J.H.; Page, S.E. Adaptive parties in spatial elections. Am Polit Sci Rev 1992, 86, 929-937.
- Solow, D.; Burnetas, A.; Tsai, M.-C.; Greenspan, N.S. Understanding and attenuating the complexity catastrophe in Kauffman's NK model of genome evolution. Complexity 1999, 5(1), 53–66.
- 29. Mintzberg, H. The Structuring of Organizations. Prentice Hall: Englewood Cliffs, NJ, 1979.
- 30. Jensen, M.C.; Meckling, W.H. Theory of the firm: Managerial behavior, agency costs and ownership structure. J Financial Econ 3: 1976, pp. 305–360.
- 31. Bower, J.L. Managing the Resource Allocation Process. Harvard Business School Press: Boston, 1970.
- Jensen, M.C.; Meckling, W.H.; Baker, G.P.; Wruck, K.H. Coordination, control, and the management of organizations: Course notes. Working paper 98-098. Harvard Business School: Boston, 1999.
- Kauffman, S.A.; Johnsen, S. Coevolution to the edge of chaos: Coupled fitness landscapes, poised states, and coevolutionary avalanches. J Theor Biol 1991, 149, 467–506.
- 34. Simon, H.A. Models of Man: Social and Rational; Mathematical Essays on Rational Human Behavior in a Social Setting. Wiley: New York, 1957.
- 35. Milgrom, P.R.; Roberts, J. The economics of modern manufacturing: Technology, strategy, and organization. Am Econ Rev 1990, 80, 511–528.
- 36. Milgrom, P.R.; Roberts, J. Economics, Organization and Management. Prentice Hall: Englewood Cliffs, NJ, 1992.
- 37. Porter, M.E. What is strategy? Harvard Bus Rev 1996, 74(6), 61-78.
- 38. Siggelkow, N. Change in the presence of fit: The rise, the fall, and the renaissance of Liz Claiborne. Acad Manage J 2001, 44, 838-857.
- 39. Levinthal, D.A. A survey of agency models of organizations. J Econ Behav Org 1988, 9, 153-185.
- 40. Holmström, B.; Milgrom, P.R. The firm as an incentive system. Am Econ Rev 1994, 84, 972-991.
- 41. Aghion, P.; Tirole, J. Formal and real authority in organizations. J Pol Econ 1997, 105, 1-29.
- 42. Maskin, E.; Qian, Y.; Xu, C. Incentives, information, and organizational form. Rev Econ Studies 2000, 67, 359-378.
- 43. Holmström, B. Moral hazard and observability. Bell J Econ 1979, 10, 74-91.
- 44. Holmström, B. Moral hazard in teams. Bell J Econ 1982, 13, 324-340.
- 45. Rivkin, W.; Siggelkow, N. Balancing search and stability: Interdependencies among elements of organizational design. Working paper. Harvard Business School, Boston, 2002.

#### NOTES

- 1. The simulation software was written in Visual C++ with MFC and run on conventional Pentium processor-based personal computers.
- 2. In his classic field research on the resource allocation process, Bower [31] is especially vivid. "[P]lanning begins at some low level and moves up toward division management" (p. 42), and "it is rare for presidential review to result in substantive change in the plan" (p. 43). "Given the prevalence of 'bottom-up' planning, it is not impossible (nor really unusual) for interdependent . . . subunits to develop plans that are inconsistent with each other" (p. 47). "The managers who make the individual investment and/or planning decisions respond to unique and personal sets of incentives. These incentives are not necessarily or likely to be the same as those of managers at any other place in the organization" (p. 20). Nonetheless, "there are levers which management has at its disposal to influence the . . . process. . . . [O]rganization, measures, and rewards—what we will call the corporate structure—indicate to a manager what 'the corporation wants of me,' and hence play a critical role in shaping the decision rules a manager uses to organize the demands of his job" (p. 54).
- 3. With INCENT = 0 and a rubberstamping CEO, one might ask whether the independent subordinate managers still constitute a firm. In this article, we take firm boundaries as given. A promising avenue for future research is to use NK methods to examine the challenging issue of firm scope. Note, however, that the above results do not hinge on INCENT being equal to 0. As long as INCENT < 0.25, the global peak 1111 continues not to be a sticking point, whereas 0101 remains a sticking point. With INCENT > 0, it is more natural to think of the subordinates as members of the same firm.
- 4. The small recombinatory effect of the CEO that was present in Figure 2 disappears as well. If a firm with P = 1 and INCENT = 1 starts at a local peak, then each subordinate proposes only the departmental status quo to the CEO, the CEO never considers alternatives that involve simultaneous change in two decisions, and the firm does not move.