Landscape Design: Designing for Local Action in Complex Worlds

Daniel A. Levinthal; Massimo Warglien

Organization Science, Vol. 10, No. 3, Special Issue: Application of Complexity Theory to Organization Science. (May - Jun., 1999), pp. 342-357.

Landscape Design: Designing for Local Action in Complex Worlds

Daniel A. Levinthal • Massimo Warglien
The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104
Universita' Ca' Foscari University of Venezia, Department of Business Economics and Management, 30123 Venezia, Italy

Abstract

In recent years, the management literature has increasingly emphasized the importance of self-organization and "local action" in contrast to prior traditions of engineering control and design. While processes of self-organization are quite powerful, they do not negate the possibility of design influences. They do, however, suggest that a new set of design tools or concepts may be useful. We address this issue by considering the problem of landscape design—the tuning of fitness landscapes on which actors adapt. We examine how alternative organizational designs influence actors' fitness landscapes and, in turn, the behavior that these alternative designs engender. Reducing interdependencies leads to robust designs that result in relatively stable and predictable behaviors. Designs that highlight interdependencies, such as cross-functional teams, lead to greater exploration of possible configurations of actions, though at the possible cost of coordination difficulties. Actors adapt not only on fixed landscapes, but also on surfaces that are deformed by others' actions. Such couupled landscapes have important implications for the emergence of cooperation in the face of social dilemmas. Finally, actors' perceptions of landscapes are influenced by the manner in which they are framed by devices such as strategy frameworks and managerial accounting systems.

(*Landscape Theory*; *Organizational Design*; *Coordination*; *Organizational Adaptation*)

Introduction

In recent years, the management literature has increasingly emphasized the importance of self-organization and "local action" in contrast to prior traditions of engineering control and design. Across the different facets of the literature, this shift in emphasis goes under a variety of labels. Whether termed empowerment, continuous learning, or intrapreneurship, the need to harness the knowledge and understanding of lower-level actors within organizations now seems to be widely acknowledged. This enthusiasm for self-organization and the potential wisdom of local action suggests that a revisitation of some basic issues of organizational design may be of value.

The very idea of designing a self-organizing system may seem an oxymoron. While processes of self-organization are quite powerful, they do not negate the possibility of design influences. They do, however, suggest that a new set of design tools or concepts may be useful. Self-organizing processes depend upon the context within which they arise. By manipulating the context, one may indirectly affect the dynamics of the process. Such efforts require a new set of design concepts that provide a language for addressing the seeming contradiction of designing for autonomous action and, in turn, pose a new set of design questions.

To develop such concepts, we draw from recent work on the theory of complex, self-organizing systems, particularly Stuart Kauffman's (1993) ideas of fitness landscapes. In this paper, we consider the issue of landscape design—the tuning of the fitness landscape on which actors adapt. The underlying idea is that by designing the surface on which adaptation processes take place, one may affect the quality of the adaptive process without the need to specify directly individual behavior. The approach is inherently dynamic, focusing on designs for local adaptation—the behavior of individuals that is guided by feedback from their particular task environment.

Landscape design influences behavior by influencing the relationships between individual action and payoffs, rather than directing individual's actions themselves. From this point of view, landscape design is similar in spirit to what economists call "mechanism design" (Hurwicz 1973). However, our approach differs from that of standard mechanism design with regard to the behavioral assumptions underlying individual action and, in turn, the importance of the dynamics of an actor's response to a given incentive structure. Mechanism design

1047-7039/99/1003/0342/$05.00

has primarily addressed the issue of incentive compatibility among rational actors. The most prominent application of such ideas to organizations occurs in the context of agency theory models (Baiman 1982, Levinthal 1988).

The analysis developed here differs from incentive design based on agency theory in its underlying postulates about individual choice processes. Building on the traditions of Simon (1955) and March and Simon (1958), we assume that individual behavior is adaptive and is driven by adaptive search processes. As a result, our design concern is the shaping of the local context in which individual actors adapt. This leads us to focus on the dynamics of the adaptation process instead of the equilibrium properties of incentive schemes. Because individual behavior is driven by adaptive search processes, the entire payoff surface is of importance in guiding behavior—not just the characterization of a global peak that defines the optimal response of an actor to a specified payoff structure. Landscape design is concerned with the behavioral "paths" actors take rather than the identification of particular peaks on a payoff surface.

The topography of the payoff surface, which we will term a fitness landscape, in turn, depends upon the degree to which the payoff to a given choice is dependent on other choices (Kauffman 1993). Increasing the density of interdependencies affects the complexity of the landscape and, consequently, the emergent patterns of behavior. Thus, structuring interdependencies is the key landscape design task. From this point of view, landscape design preserves the focus on interdependencies of traditional organizational design theories (Thompson 1967). As a result, landscape design sits at the crossroads of the often disjoined domains of mechanism design and organizational design.

Landscape design, however, differs from traditional organizational design in that it reverses the role of interdependencies. In the classic work of Simon (1982) and Thompson (1967), the critical design choices lie in decisions regarding the locus of interaction among elements of the organization. For Simon (1982) these issues are reflected in the idea of decompositions, while Thompson (1967) identified the various forms of interdependence (sequential, pooled, and reciprocal) that may be present in organizational action.

For traditional organizational design, the critical challenge is first to identify the interdependencies posed by the tasks and the technology of the organization and then to specify an organizational structure and coordination mechanisms consistent with the set of interdependencies that have been identified. Broadly speaking, the design problem is to chose the organizational structure so as to maximize the intensity of interactions within a unit of the organization (i.e., an unit of decomposition) and to minimize the interactions across elements of the system. The underlying structure of interdependencies is always taken as a given.

Our approach is similar in spirit to that of Simon and Thompson in that it emphasizes the centrality of interdependencies for design. At the same time, it implies a shift from designing on the basis of a given set of interdependencies to designing by manipulating the set of interdependencies. The design of self-organizing systems requires a tuning of interdependencies.

One may argue that the structure of interdependencies is given by the "world," and that a designer can at best modify the perception of actors of these inherent interdependencies by manipulating flows of information and patterns of communication—features on which traditional organization design focuses. We contend that while the external environment provides significant constraints, there are usually important degrees of freedom left to the designer. Indeed, a design problem arises *because* there are these degrees of freedom. Consider, for example, the design of a computer system. A computer system can be designed such that the operating system is tightly linked with the processor or the two subsystems can be designed in a relatively modular, independent manner (Baldwin and Clark, 1997). Either design choice can produce a computer system; however the two design choices may result in systems that differ in their efficiency at a particular task and in their robustness to changes in the task and to changes in the components of the overall system.

The following section develops the framework of a fitness landscape and discusses how this framework can be applied to problems of organizational design. We can think of a fitness landscape at various levels of analysis—individual, team, and organization. That is, the attribute that contributes to performance may be the set of choices that a single individual makes, a set of individuals (i.e., a team), or a broader system such as an organization.

Fundamentally, there are two types of landscape structures, depending on whether there is a single payoff associated with a set of actions or a distinct payoff surface for different actors. The former structure is akin to team theory (Marshack and Radner 1972) in that actors may make better or worse choices and there may be issues of coordination in the making of these choices, but incentive conflicts are absent. Within the class of single-payoff landscapes, landscapes may be single- or multipeak, depending on the specification of the interdependence among individuals. Introducing multiple payoff structures, that we frame in terms of so-called coupled landscapes (Kauffman 1993), allows us to introduce problems of conflicting interests and incentive alignment. We explore single-payoff, single-peak landscapes and then consider the implications of interactions within a given landscape. Subsequently, we consider the nature of the design problem in the context of multiple, coupled landscapes.

The landscape designer is faced with two distinct but related design challenges. First, it is important to understand the implications of a given landscape for behavior. The answer to this question, of course, depends on the process of search and adaptation in which actors engage as they attempt to move within (i.e., climb) a given landscape. Given this insight as to what behavior a given landscape may elicit, there is a subsequent question of what dynamics of behavior are desired. Is incremental behavior desirable or does the landscape designer wish to encourage wider search and exploration? As a result, a good landscape design depends upon the designer's belief about the organization's external context and, in particular, the designer's beliefs about what might comprise effective adaptive performance in a dynamic environment.

Interdependencies and Fitness Landscapes

The idea of a fitness landscape was introduced in the context of the biology literature by Sewell Wright (1932). The landscape is simply a mapping from an organism's genetic structure to its fitness level. This concept of a fitness landscape has rather natural analogues in the domain of social and economic phenomena.

In the domain of organizations, the set of elements that influence fitness may be interpreted in a variety of ways. For instance, the attributes that determine fitness for an organization may comprise the elements of its business strategy, its human resource policy, manufacturing system, and so on. Fitness can be represented by profit, or by a mix of variables related to the organization's goals. Alternatively, the fitness landscape can be applied as a mapping of the actions of a set of individuals and their collective performance. The same structure can also be viewed from an individual's perspective. In this case, one is modeling the payoff to a focal individual conditional on the actions of the other members of the group of actors under consideration. Applied in this manner, peaks in a fitness landscape would correspond to a best response to the actions of others, and a point on an individual peak would correspond to the notion of a Nash Equilibrium—an actor could not improve his payoff conditional on the current choices of the other actors (Axelrod et al. 1995).

Kauffman (1993) has shown how the topography of a fitness landscape is influenced by the degree to which the contribution toward fitness of genes is interdependent. Kauffman characterizes fitness landscapes with essentially two structural variables, *N*, the number of elements that characterize the entity (genes in the context of Kauffman's work, actions or policy choices in the context of organizations), and *K*, the number of elements of *N* with which a given attribute interacts. In biology, the notion that the fitness contribution of genes has such interdependence is referred to as epistatic interactions (Smith 1989). In particular, when there is no epistasis, the landscape tends to assumes a single-peak configuration (see Figure 1). As such interactions increase, the landscape becomes more rugged or multipeaked (see Figure 2).

Why does a lack of interdependence generate a single peak landscape? If the contribution of each actor to global fitness is independent from that of others, there is an optimal behavior independent of others' behavior. If we view the landscape as being comprised of policy choices, a lack of interdependence in the fitness landscape implies a situation of universal best-practices. Each actor improves the collective fitness by improving locally his or her own contribution to fitness, so that the global maximum can be reached from any point by walking toward adjacent locations yielding higher payoffs.[1]

What underlies the "ruggedness" of the landscape? Multiple peaks are the direct result of interdependence among a set of actors or policy choices. With a high degree of interdependence, a change in a single action may

Figure 1 **Single-Peaked Landscape**

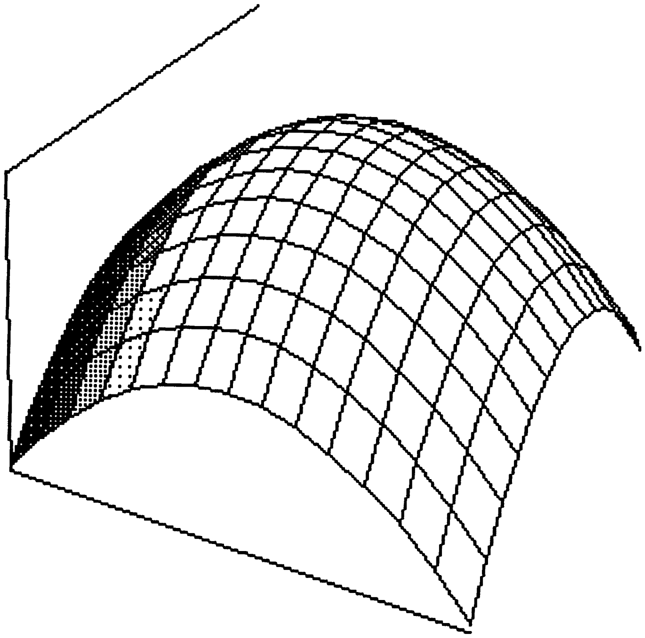

Figure 2 **Rugged Landscape**

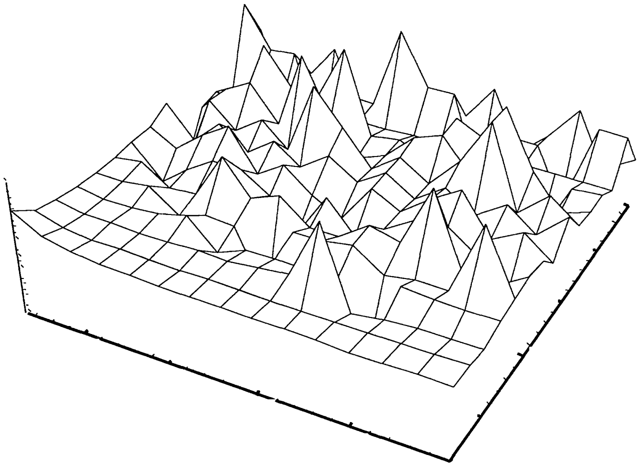

appear dysfunctional (i.e., diminishing performance or fitness) despite the fact that a simultaneous change in a large set of actions may enhance performance. For instance, adopting a just-in-time inventory system may diminish organizational effectiveness if there are not also concurrent changes in the management and production systems. However, with appropriate adjustments to other elements of the organization's policies, the change to a just-in-time inventory system may yield important performance benefits.[2]

The implications of a rugged landscape are very much a function of the search behavior of the actors moving on the landscape. If actors are omniscient and can readily search globally, then the behavior on a smooth surface such as depicted in Figure 1 or a rugged surface as suggested by Figure 2 would not fundamentally differ. In both instances, actors would identify the global peak. The analysis developed here is within the behavioral tradition of intelligent local search (March and Simon 1958, Cyert and March 1963). Actors are assumed to be intelligent, but that intelligence is local to their position on the landscape. Thus, actors are assumed to be able to identify the positive and negative gradients around their current position, but not capable of making similar judgements for more distant locales.

Local adaptation, in conjunction with a rugged fitness landscape, has important implications for the pattern of behaviors that emerge. Actors will engage in incremental "hill-climbing" leading them toward the closest peak. In mathematical terms, the hill is termed a basin of attraction and the peak an attractor (Weisbuch 1991) for the adaptive dynamics of the system. Those locations that lead to a common peak, or attractor, belong to a common "basin of attraction." When the combinatorial space of possible configurations of the elements is very large (as in the human genome or organizational design), the existence of hills dramatically reduces the number of possible configurations emerging from such adaptive processes. Simple adaptive walks to local peaks produce the emergence of quite ordered structures out of a vast number of possibilities; the more interdependencies in the entity, however, the larger the number of possible emergent configurations (Kauffman 1993).

Single-Peak Landscapes and Robust Design

When the elements of a system are independent, the fitness landscape they generate tends to be single peaked, with smooth surfaces leading gently to the maximum from any starting point. As the visual imagery suggests, the main feature of such landscapes is that they are easy to scale. Any kind of local experimentation and myopic adaptation will inevitably lead toward the top of the hill, with the only constraint being that climbers are able to discriminate, on average, between immediate neighboring locations that are uphill of the current location versus those that represent a downward move.

This immediately suggests the key attractiveness of such landscapes from a design point of view: designs that provide single-peak, smooth surfaces for adaptation will make the dynamics of autonomous agents highly predictable under a great variety of behavioral assumptions. For example, in such a setting the only difference between a fully omniscient actor and one with only "local" intelligence is the speed at which they reach the global optimum. A related point is that, in single-peaked landscapes, the independence of choices implies that local improvement in performance (fitness) always leads to global improvement. Furthermore, different starting points in the landscape will make no significant difference in the final outcomes of the adaptive process.

Although the requirement of perfect independence among the elements of the system may seem too demanding, landscape theory (Kauffman 1993) suggests that the main features of single-peak landscapes also can be preserved in settings of low, but not zero interdependence. When *K*, the parameter expressing interdependence, is low compared to *N*, the number of elements in a system, the landscape will tend to have a very large and smooth basin of attraction leading to the highest peak. Lower peaks will tend to lie in peripheral regions, with minor valleys separating them from the dominant hill. Thus, although adaptive walks toward the highest peak will no longer be guaranteed, the probability of reaching the top through local improvement remains quite high.

We define *robust* design as a design in which, thanks to low interdependence among the elements of a system, the asymptotic behavior of autonomous agents can be predetermined with high probability, independent of the knowledge of their specific choice process and their starting point in the landscape. Robust design is a familiar issue to product designers or urban planners, whose task is often to provide artifacts that can be used in predictable ways by agents whose intelligence, knowledge, and attention may vary and whose behavior cannot in any case be prescribed or controlled in detail. It is also central to economics. Markets are shown to be efficient under conditions of convexity and agents' independence (Arrow and Hahn, 1971), which in turn implies the existence of smooth, single-peaked landscapes.[3]

We suggest that robust design is also a key concern for organization design. Robust design is desirable when it is not clear where the best solution is located—thus search and adaptation are needed—but the designer wants to be reasonably confident that such a solution will be reached or at least approached. By designing low interdependence systems, the designer generates smooth adaptation surfaces that help ensure that local experimentation promotes global improvement. Furthermore, such designs reduce the cost of mistakes in experimentation. The smoothness of the landscape implies that a step in the wrong direction by a single actor will involve only a minor degradation of global performance, which makes the system resilient to errors and avoids the threatening effects of a high cost of mistakes. As a result, robust design, or low interdependence, is a critical property when continuous improvement policies are sought.

Robust Design and Local Improvement in Manufacturing Process

The practice of Japanese manufacturing provides striking examples of robust design based on single-peaked landscapes in the organization of processes at the workplace level. This is not surprising given the emphasis of Japanese management on the global virtues of local adaptation. The basic design challenge is that of creating contexts that may naturally lead workers to adapt to their task environment. We will consider two examples, one at the microlevel of designing individual tasks, the other at the higher level of manufacturing process design.

An example of the exploration of single-peaked landscapes at the task design level is that of Baba-Yoke, literally meaning idiot-proof design (Shingo 1983).[4] The concept of Baba- (or Poka-) Yoke is that the process is arranged so that workers can naturally be driven toward the correct operations. Conceptually, this implies designing the task in ways that lead naturally to climbing the fitness hill with no false steps. Such a design requires two critical elements:

- Smooth, unique attractors for behavior. For example, in die change operations, using pear-shaped clamps that can smoothly be brought to fit in only one way and thereby drive even approximate movements into the right direction. In this case, the landscape is designed by the physical shape of the task environment.
- Unambiguous, rich feedback that helps an individual actor recognize at any moment the fitness surface and its gradient (c.f., the Japanese emphasis on rich visual feedback for operating controls).

In general, however, robust designs for local action do not literally imply a single peak in the fitness landscape, but rather a "tuning down" of interdependencies. Consider a classical hard coordination problem of synchronizing operations in order to reduce interoperational stocks (i.e., buffer inventories). In classical manufacturing practices, this problem is complex because it is coupled with the problem of fully utilizing machine capacity. This coupling introduces additional constraints that make the problem hard to solve, even with sophisticated operations research techniques. In addition, solutions to such a coupled problem are not robust—small changes in the production plans may generate severe losses of synchronization.

In Toyota-like production systems these two problems are decoupled by allowing machines to be unsaturated (i.e., not fully utilized) and by leveling production on how much needs to be produced (Ohno 1992). With production leveling obtained in this manner, "synchronizing becomes very easy" (Shingo, 1983, p. 45). In this case, synchronizing simply requires that each station adapt, via the pull method of kanban, to the production rhythm dictated by the downstream station. Each station becomes directly interdependent only with the two neighboring ones. The number of interdependencies is reduced to two from the vast combinatoric problem posed by the overall problem of capacity utilization. The principle of limited epistatic interactions also facilitates mutual help: if a worker is late, the preceding and following worker in the production process will help in order to maintain a steady production flow (Ohno 1992, Shingo 1983).

In sum, the Toyota production system heightens sequential interdependencies, both internal to the production process and externally with suppliers and customers, while globally there is a substantial reduction in interdependencies. Rather than a central planner trying to solve the problem of materials planning, a simple "pull" system is used built upon sequential interdependence and local action.

Rugged Landscapes—Problems of Diversity and Coordination

In general, landscapes are complex in that a number of local peaks are present. The presence of a rugged fitness landscape has important implications for the diversity of organizational forms that we observe (Levinthal 1997). As long as search processes are to some degree local, the set of policies at one point in time will significantly effect the subsequent path of organizational evolution. In this sense, imprinting effects (Stinchcombe 1965) persist even in the absence of organizational inertia.

Those features of adaptation in multipeaked landscapes can be illustrated not only by formal analysis or computer simulation (Levinthal 1997) but also by simple laboratory experiments involving human, rather than artificial, actors. Many experiments on coordination games provide striking examples. In the domain of game theory, coordination games typically present a multiplicity of equilibria. The "average opinion" game of Van Huyck et al. (1991) has been an important context for the study of coordination in experimental economics. In this experiment, players select a number between one and seven. The payoff is distributed to players according to a rule that punishes those players that have made choices deviating from the median choice of the group. The game is repeated with subsequent choices informed by feedback on the choices made by others in the prior round (sometimes this feedback is merely the median choice; in other settings, more detailed information about the distribution of choices is provided). The game clearly presents strong epistatic interactions in the payoff function and, as a result, there are multiple peaks in the collective payoff landscape. Table 1, taken from Van Huyck et al. (1991), shows the payoff matrix for a player of one such game. Seven peaks in the fitness landscape are generated by the payoff matrix.

The experimental results are consonant with the main properties of adaptation in rugged landscapes that we have sketched above. First, players initially start out of equilibrium (they are remote from the peaks). However, all experiments end up on one of the peaks. Thus, the property of peaks as attractors is well supported. Second, there is diversity of behavior. While all experiments, after some iterations, end up at a peak in the landscape, it is not always the same peak. There is diversity in the emergent configurations. Finally, the equilibrium emerging at the end of each sequence of play is *always* equal to the median of the distribution of choices made in the first period of the sequence. Thus, there is extreme path dependence in the adaptive dynamics of the game.

Table 1 Payoff from Median Value Game

Your Choice of χ	7	6	5	4	3	2	1
7	1.30	1.15	0.90	0.55	0.10	−0.45	−1.10
6	1.25	1.20	1.05	0.80	0.45	0.00	0.55
5	1.10	1.15	1.10	0.95	0.70	0.35	−0.10
4	0.85	1.00	1.05	1.00	0.85	0.60	0.25
3	0.50	0.75	0.90	0.95	0.90	0.75	0.50
2	0.05	0.40	0.65	0.80	0.85	0.80	0.65
1	−0.50	−0.05	0.30	0.55	0.70	0.75	0.70

Source: Van Hyck et al. 1991. Reprinted with permission.

Rugged Landscapes and Inducing Diversity

From a design perspective, the interest is not solely on the implications of a rugged landscape for organizational behavior, but on the question of what sort of landscape topography one might wish to construct. Why would an organizational designer purposely construct a rugged landscape? Such a construction would seem inferior to a simple, single-peak environment in which there are no difficulties in coordinating collective action. With a robust design, all individuals are inextricably guided to a mutually consistent set of actions.

This characterization of the virtue of a single-peak landscape also suggests its fundamental weakness from a design perspective. A single peak solves the problem of coordination, but it does so at the cost of diversity. The presence of a rugged fitness landscape leads to unpredictability in collective behavior. Put differently, and more positively, a rugged landscape encourages nonincremental search and exploration. This contrast is laid out in the context of a comparison of competing design principals—the hierarchy of design versus cross-functional integration.

Hierarchy of Design Versus Cross-Functional Integration

The evolution of new product design practices provides some insights into the benefits, and the complications, of making landscapes more rugged. Traditional approaches to design emphasize concepts of hierarchical decomposition (Von Hippel 1990, Alexander 1964, Simon 1982). The design problem is decomposed such that coordination among elements is minimized. The designers of each subassembly need only know the characterization of the interface with neighboring components.

The emerging new product design paradigm (Imai et al. 1985, Wheelwright and Clark 1992) emphasizes building contexts that channel autonomous, local action. In particular, the approach attempts to foster cross-functional, nonhierarchical interactions to achieve more innovative outcomes, as well as to speed the ultimate time-to-market of the new design. A key feature of this new set of practices is to bring together a rich set of constraints and design tradeoffs at the beginning of the design process (Womack et al. 1990, p. 115). This is achieved primarily through cross-functional teams that highlight the interdependencies among the technological, manufacturing, marketing, and other dimensions of the design process. This broadens the search process. In the words of Wheelwright and Clark (1992), it enlarges the mouth of the development funnel. Similarly, Imai et al. (1985) refer to cross-functional teams as variety amplifiers.

For the purposes of landscape design, cross-functional teams bring together multiple constraints, increase interdependencies in early design phases, and thus make the design landscape more rugged. Furthermore, the variety of functional background of team members makes it likely that different starting points are initially sampled. As a result, a variety of alternative designs is likely to emerge. In such a setting, the speed of an initial design effort is likely to be slow and laden with frustration as a result of conflicting constraints. As Wheelwright and Clark (1992) stress, getting convergence to a final set of specifications is far more challenging. At the same time, a cross-functional team is likely to consider and to some degree experiment with more alternative configurations. If the individual alternatives themselves are mutually consistent (i.e., most functional elements complement one another or are at least compatible with one another), then the choice set will comprise a set of local peaks in the fitness landscape. As a result, this initial period of casting about in a highly interactive (i.e., cross-functional) effort is likely to prove useful in identifying a new, globally attractive, though not necessarily optimal, configuration.

Search in rugged landscapes is basically the search for new sets of complementarities. This is necessary when "breakthrough" and "platform" innovations are desired and hence new configurations of elements responding to the constraints are needed. Of course, not all product development efforts require this sort of exploration of rugged landscapes. When the basic design is stable and only incremental refinements are needed, traditional decompositions are preferred (Wheelwright and Clark 1992, p. 175). Such decompositions lead to few peaks and a relatively smooth fitness landscape that is readily searched through a process of piecemeal adaptation.

Even during the evolution of single platform development projects, however, search patterns change over time, moving toward more incremental search as the main architectural decisions have been made. This is consistent with landscape theory: once a promising region in the solution space is found, some independence among functional units can be restored and efforts can be directed toward the most promising peak. The nature of the search process induced by the landscape structure also has important implications for the role of management during the development of a new product. Path dependence makes management efforts more important and fruitful in the early phases of product development; moves away from the path are difficult and costly, and reorientation efforts often fruitless. This implies that in multipeaked landscapes, the usual distribution of management efforts over time (concentrated on the final steps of the project) has to be reversed (Wheelwright and Clark 1992, p. 32). Management efforts are more useful in the initial phases of search, while the ability to influence the project's outcome decreases as path dependence increases. At the same time, such management intervention need not address the details of the development effort (these are left to local action), but rather provide a sense of direction—a shared representation of what constitutes a fitter design. These efforts restrict the area of the landscape that is searched by narrowing the neck of the development funnel (Imai et al. 1985).

There are a variety of other organizational practices that reflect a similar contrast. The use of matrix structures, which force people to respond simultaneously to multiple, potentially conflicting, decision premises heightens the perception of interdependencies. Similarly, the timing of actions can be structured so as to heighten or reduce interdependencies. Asynchronous action implies that each single actor faces a simpler landscape, since the degrees of freedom of subsequent actors have been reduced by prior choices. Manipulating distance, social and geographic, among actors can serve to tune interdependencies. For instance, if individual payoffs are affected only by local, not global, externalities, the more others are effectively distant, and the lesser is K. By reducing distances, via technology or other means, one increases the apparent ruggedness of the landscape—suggesting why a "global village" is such a complex social structure.

Exploring Rugged Landscapes: Local Adaptation, Long Jumps, and Recombinations

A fundamental challenge in complex, adaptive processes is the dual imperative to exploit past experience and, in turn, current wisdom, while at the same time sustaining efforts at exploring alternative bases of action (Holland 1975). In a single-peak world, local, incremental adaptation is a satisfactory solution to this dilemma. The current policy reflects the wisdom of past experience, and search in the immediate neighborhood of the current set of actions will, ultimately, lead one to the global peak.

In a rugged landscape, such an incremental search procedure will lead only to the local peak closest to the starting point of the search process, regardless of its height relative to other peaks in the landscape. As a result of this locking-in to the first available solution, we observe a strong form of path dependence and, on average, modest performance.

One mechanism to overcome such competency traps (Levitt and March 1988) is to engage in long-jumps, random exploration of more distant portions of the landscape. However, such distant search efforts, by not exploiting the wisdom gained from past experience, are likely to result in a deterioration of performance. This may not be a serious liability in the initial phases of the exploration of a new landscape, when little is known about the landscape and the chances of finding more attractive positions by sampling new, distant places is high. However, as experience is acquired, random search risks wasting search effort and, perhaps more importantly, risks forgoing the benefits of known, attractive sets of behavior.

The problem of adaptation strategies in rugged landscapes can be thus reframed as a familiar dilemma: how to reap the benefits of exploration without losing the advantages of exploiting acquired knowledge (March 1991). An at least partial resolution of this dilemma lies in search strategies that recombine elements of existing "solutions" (Cohen 1981). Schumpeter's (1934) entrepreneurs do not randomly sample the space of alternatives; rather, they rather find new, unforeseen combinations of known but previously distant elements. In terms of our framework, they combine building blocks from distant regions of the landscape. Thus, there is a strong parallel to processes of biological recombination (Holland, 1975, Cohen 1981, Goldberg, 1989).

What are being recombined are partial solutions. What gets recombined depends of course upon the situation that is being analyzed. In a new product design context, it can be building blocks of the new product concept, or partial solutions to problems of product design, such as a good disk drive solution with a good processor. Organizational processes can be recombined as well, such as a good order fulfillment routine with a good system of inventory control. In pharmaceutical research, the elements being recombined may be the chemical building blocks of a molecule.

The critical point is that if there are complementarities (epistatic interdepedenencies) among elements, complementary elements tend to be found together in better solutions (Holland 1975). These building blocks can be recombined with other blocks of complementary elements from other solutions. Such a strategy has strong advantages in rugged landscapes. On the one hand, it allows whole blocks of existing solutions to be changed rather than modifying them piecemeal. On the other hand, it exploits knowledge embedded in the recombined alternatives. The highest peaks in a rugged landscape tend to be located rather close to one another (Kauffman 1993). As a result, if two peaks are found, it is likely that another peak lies in the region between them. Landscape theory predicts that building blocks of different good solutions to the same problem have a high probability of exhibiting complementarity among themselves and thereby jointly increasing fitness.

A recombination is a shift in position that exploits existing knowledge without being trapped by this knowledge. To return to the example of cross-functional teams, this argument suggests that such teams work, not because they are more "creative," in the sense that their random long-jumps have a better distribution of outcomes, but because they build on the diverse knowledge of their members and thus permit a search process that can simultaneously explore and exploit. Good adaptation strategies make use of the correlated structure of the rugged landscape to extract information from differentiated local knowledge.

Another important issue related to search strategies in rugged landscapes concerns the "seeding" of search, i.e., the problem of identifying the starting points of the search process. A key point about rugged landscapes is that multiple starting points are needed in order to broadly explore the landscape. Subsequent search is conducted by leveraging upon the understanding of the topography of the landscape that the initial sampling has generated—e.g., by recombining the tentative "solutions" of the more successful initial trials. Of course, if enough is known about the landscape, an organization can focus on specific areas of the landscape for the initial exploratory effort and exclude less successful paths.

An interesting analogy is provided by pharmaceutical drug design. When designers know little about which molecules are effective, they often resort to random search in a vast molecular region. As more knowledge is acquired through these initial tests, designers select a few subregions and start making incremental improvements, escalating local hills. Finally, they select the best hill and make further local refinements ("optimization") of the molecule. In contrast, when some preliminary knowledge exists, smaller regions of the space are explored at the beginning and search can focus very rapidly on the few more promising hills.

Coupled Landscapes

The analysis to this point has considered potentially complex, but fixed landscape topographies. Actors attempting to climb a given landscape may find, however, that the landscape shifts and deforms as the result of the adaptive efforts of other actors with whom they are linked. For instance, the payoff to one actor deciding on the extent to which to act in a cooperative manner may depend on the behavior of other actors. Landscapes may also be coupled as a result of processes of mutual adaptation. For instance, as an organization modifies its policies and actions, its suppliers and/or customers may modify their actions to be consonant with these policies and, in turn, may change the payoff associated with these policy choices.

Formally, the idea of coupled landscapes is captured by Kauffman's *NK(C)* model (Kauffman 1993). The parameter C refers to the number of the N elements that are linked across entities, in contrast to the parameter K which indicates the degree of interrelatedness within a single entity. If an organization is characterized by N-attributes, then the payoff to the organization associated with a set of attributes may be independent of the characteristics of other organizations in the population (i.e., $C = 0$) or it may depend on a variety of attributes of other organizations (i.e., $C > 0$).

For instance, the payoff associated with shifting some elements of the production process may not depend on the set of attributes characterizing other organizations; however, the value of different product positioning strategies may well depend on the positioning strategies of other organizations. Similarly, individuals within an organization may climb their own landscape surface that is deformed by the choices of other actors within the organization. Such a structure captures the problem of inducing cooperative behavior among self-interested individuals. In contrast, in our prior analysis of individual adaptation within an organization, we assumed that individuals climb a shared, organizational fitness landscape.

Coupled landscapes result in quite complex dynamics as movement by one actor shifts the landscape topography for other actors. Actors are engaged in hill-climbing, but climbing a continually shifting landscape. We explore such dynamic processes internal to the firm using the problem of cooperation and social dilemmas (Dawes 1980) as a motivating example. Subsequently, we use this framework to examine processes of co-adaptation between the organization and its customers.

The problem of cooperation is a quintessential organizational challenge. For an organization to be a meaningful and valuable unit of analysis, it must be the case that there are benefits to coordinated collective action. It is the role of organizational design to provide the context for realizing these collective benefits. Conflict, however, is an inevitable feature of organizational life (March and Simon 1958). The risks of free riding and intraorganizational opportunistic behavior imply a constant threat of collective action failures for which hierarchical monitoring and sanctioning can only partially mitigate (Crozier 1964, Tirole 1986). The value of collective action in the face of these incentive conflicts implies that the design of mechanisms to facilitate cooperation is a central challenge of organizational design.

Social dilemmas present a particular challenge to collective action in that the payoff structure differs for the individual and for the group. Individuals have a strong incentive to free ride and deviate from socially optimal behavior. While life is studded with social dilemma situations, the picture we observe in reality is not so somber: cooperation seems to arise even in situations where economic rationality would not predict it. Consistent with this sense of everyday life, laboratory evidence shows that in social dilemmas pure free riding rarely occurs; although, at the same time, the amount of cooperation achieved is typically far from the social optimum (Dawes and Thaler 1989, Holt 1995, Ledyard 1995). Furthermore, the amount of cooperation varies systematically with factors such as the incentive structure, sanctioning opportunities, experience, and communication (Ledyard 1995).

Thus, since the problem of avoiding or mitigating free riding is a central theme in the design of group incentives (Nalbatnian and Schotter 1997), both the observation of everyday life and laboratory evidence suggest an interesting design problem: when facing social dilemma situations, is it necessary to realign individual incentives and collective interest, or may other factors act as substitutes for incentive alignment? The economics literature focuses on efforts of incentive alignment, but such efforts are not always successful and are often quite costly.

One classic form of public good problems, or social dilemmas, is the collective use of a common, exhaustible resource. For such a public good, the social returns derived from increased use rise until a saturation level is reached, beyond which returns from the common resource decline. For example, the use of e-mail within an organization is increasingly productive until an overcrowding level is reached, beyond which "infoglut" makes electronic communication a source of painful loss of time and productivity. Collecting orders from sales persons is useful until production capacity is saturated and clients start being hurt by increasing delays in supply. Photocopying papers in a university department is useful until reproducing everything passing through professors' hands creates long queues at the copying machines and ever growing stacks of documents to read.

Experiments have designed many common resource dilemmas of this kind. One well-known experiment in common resources is the Common Pool Resource (CPR) game of Ostrom et al. (1992). In their experiment, each of the subjects is given an endowment of m tokens. The individual can invest his tokens either in a constant return asset or in exploiting a "common pool resource." The common pool resource provides a payoff function with diminishing returns—the payoff from the common resource increases with investment in the resource, but at a diminishing rate. This diminishing return reflects saturation effects like the overcrowding of common meadows by farmers' goats or infoglut in e-mail. An individual's share of the collective payoff from the CPR is simply their share of the investment in exploiting the CPR (if you have more goats grazing the common meadow than your neighboring farmer has, then your individual benefit from the meadow is proportionally higher).

The socially optimum level of exploitation of the CPR implies some moderation in individual exploitation of the common resource; nevertheless, at the socially optimum level of exploitation, there are private incentives for each individual to invest more in the CPR. As a result, the common resource will be overexploited and the Nash equilibrium of the game corresponds to a socially suboptimal solution. The CPR game represents a typical social dilemma of the "tragedy of the commons" (Hardin 1968).

As with most social dilemmas, however, experimental evidence shows a rather different picture from the one predicted by game theory. Manipulation of experimental conditions (Ostrom et al. 1992, Rocco and Warglien 1996) shows that some degree of cooperation can be achieved even when it is not predicted by the Nash equilibrium. In particular, groups of subjects exhibit cascades of cooperative behavior in which cooperative acts in one period elicit greater cooperation in subsequent periods, as well as dramatic collapses from high levels of mutual cooperation (Ostrom et al. 1992, Rocco and Warglien 1996). We make sense of such dynamics using the tool of coupled landscapes and consider some of the associated design implications.

Positive Cascades and Slippery Peaks

Consider, for simplicity, a two-player CPR game. Figure 3 portrays the individual payoff landscape of player 1 when player 2 employs the Nash equilibrium strategy. The landscape has a peak (E)—the maximal payoff for player 1 conditional on player 2's exploitation of the common resource (i.e., the Nash equilibrium). Suppose, however, that the agents engage in a local search process. As a result, there is a positive probability that both agents may simultaneously decrease their level of exploitation (i.e., move left on the curve depicted in Figure 3). When this happens, both individual landscapes shift: the landscape moves up, reflecting the benefits of cooperation, although the peak moves rightward to A. As a result, both individuals experience a higher payoff corresponding to the point A′. This increase in payoff reinforces the tendency to cooperate (i.e., go left on the curve), increasing the probability that in the next iteration both actors will again decrease their level of exploitation. In that case, the landscape shifts upwards again, while the peak (B) shifts further to the right. Kauffman (1993) refers to such mutual deformation of coupled, individual fitness landscapes as "dancing landscapes."

Agents are not really climbing uphill on their profit landscape. In fact, they are going downhill on their individual landscapes. Despite this, they experience an "illusory hill climbing" process because as the level of exploitation is decreased, individual hills rise up. If actors do not consider the entire topography of the landscape, but only the payoff implications of their local action (i.e., they are myopically adaptive), then the positive dynamic of an upward shifting landscape can overwhelm the marginal calculations of selfish behavior (Rossi et al. 1997).

Figure 4 depicts the joint payoff landscape generated by all possible moves of the two agents, with the Nash equilibrium and the social optimum identified. This landscape provides insights regarding the dynamics of cooperative behavior. The marginal returns to cooperation decrease as agents move toward the social optimum. Indeed, these payoffs begin to decrease after the social optimum is passed. As a result, as the social optimum is approached, increases in payoffs are no longer possible. The possibility of further virtuous reinforcing cycles diminishes, and the risks increase of generating negative self-reinforcing dynamics that drive behavior toward the Nash equilibrium.

Thus, sets of actors will tend to wander off of the social optimum and occasionally experience a cascade of destructive competitive dynamics. However, the greater gradient in individual payoffs around the Nash equilibrium implies that such a competitive catastrophe will be followed by a resurgence of cooperation and a collective bootstrapping back toward the cooperative equilibrium.

The problem of mechanism design is generally interpreted as a problem of aligning the peak of individual landscapes with the global peak as in the case of social dilemmas, or the peak for the incentive designer as in agency models. However, if actors move on individual landscapes through mechanisms other than the immediate

Figure 3 **Dancing Landscapes with Public Goods**

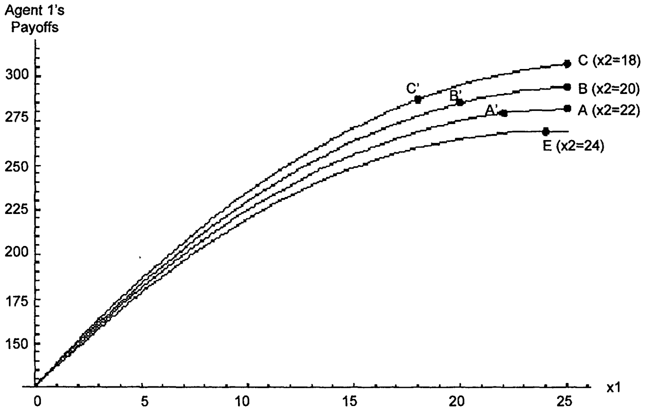

Figure 4 **Collective Landscape for Public Goods**

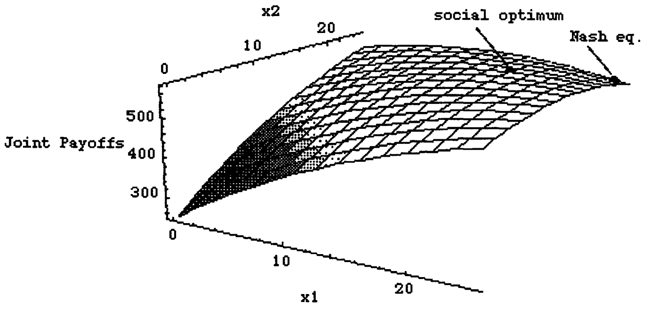

identification of a global peak, then the manipulation of local payoff surfaces may impact an actor's behavior.

One important tool of gradient management is the size of the group among which the interactions take place. A direct impact of group size, that impacts both local gradients and individual optima, is the effect on the marginal return to cooperation. In addition, however, increasing the number of players decreases the probability that they will jointly decrease their level of exploitation and increases the likelihood that at least one of them will act opportunistically. Indeed, it is a typical result of experimental research that the breakdown of cooperation results from the defection of a single individual that then triggers a chain of subsequent defections.

There is an implicit coordination problem in escaping social dilemmas (Camerer and Knez 1997). Thus, small groups are the key loci of collective action within organizations. This does not imply that large scale cooperation cannot occur, but rather that it is more likely to be induced as the aggregation of cooperation within smaller groups.[5]

In the context of illusionary hill-climbing on landscapes, the emergence of cooperation requires that actors initially move in the same direction. The relevance of the coordination problem in social dilemmas is well exemplified by Argyris and Schon's (1978) treatment of change and learning in organizations. For Argyris and Schon, rigidity to organizational change stems, in large part, from individuals' reluctance to recognize the flaws in the organization's current strategy. Such public declarations are dangerous to one's career prospects, as well as posing a threat of psychological dissonance for those individuals involved in the setting of the prior policy. Therefore, individual incentives that inhibit change and learning from experience may lead to persistence in a failing policy.

Argyris and Schon's action consultant, as well as other forms of communal discovery, help to create a positive dynamic of dancing landscapes. The testing of beliefs with close associates allows one to gauge the commonality of these ideas in a relatively safe manner. This may help coordinate initial moves in a collectively more desirable direction, which in turn may generate a self-sustaining dynamic of cooperation. Such a dynamic is similar to Weick's notion of "small wins" (1984) and to the evolution of cooperation in Axelrod (1984).

Coupled Landscapes and Coadaptation

There has been considerable attention given in recent years to the importance of firms being closely tied both to their customers (Day 1990) and suppliers (Dyer 1994). In regard to both sets of relationships the motivation revolves around the desire to make products and services more adapted to their context (the focal firm in the case of supply relations or end-customers in the marketing context).

Traditional ideas of marketing strategy are premised on notions of product positioning (Kotler 1997). The "landscape" for the firm consists of a positioning map, with attributes on the axes and the density of consumers' ideal product or preference contours determining the topography of the landscape. The challenge of positioning is then to identify peaks that are not already crowded by competing products.

We suggest the additional possibility of "interactive marketing." The landscape need not be fixed, but there may be interactions between customers' preferences and practices on the one side and product features on the other side. As a result, there is some coevolution of offerings and preferences such that, over time, there is an increasing degree of interdependence between the product and consumer.

This dynamic underlies the phenomenon of mass-customization (Pine 1993). The landscape of product offerings becomes quite rugged, with a large number of peaks mirroring the preferences of particular sets of consumers. This heightened linkage between firms and consumers creates an enormous increase in opportunities for differentiation. Potentially, this increased complexity of the landscape of product offerings could overwhelm the capacity of the organization's production process to manage such complexity. The solution to such a "complexity catastrophe" is to reduce the linkage between products and the production process. This reduced linkage is not achieved via buffering, as suggested by Thompson (1967), by rather by means of the modularization of the production process (Baldwin and Clark 1997). Thus, interdependencies are increased for firm-consumer linkages, while interdependencies are reduced internally, or at least not increased proportionately with greater product diversity.

Discussion: Landscapes—Communication and Representations

Coordination and Communication on Rugged Landscapes

A central theme in this discussion is the challenge of coordinating autonomous, but interdependent actors. Communication is an obvious candidate for supporting coordination efforts. Thus, we should expect communication to be especially valuable in rugged landscapes.

Even when the coordination problem is simple enough so that actors can see all these peaks in the payoff landscape, experimental evidence suggests that in many cases actors may fail to coordinate on the highest peak and can be trapped by their initial moves in suboptimal peaks. However, the introduction of communication can substantially alter the coordination process. An interesting case is the one in which there are two hills: one has a lower peak but is smoother (even if some actor deviates from the peak there is only a small loss for other actors), while the other has a higher peak but is steeper (there are higher losses if some actor deviates from the peak). This portrays a familiar situation in which the best outcome can be attained at the price of a higher risk in case of coordination failure. Laboratory experiments show that without communication, actors tend to select the less efficient, but less risky alternative; however, introducing communication enables actors to coordinate most of the time on the most attractive alternative (Cooper et al. 1992).

This powerful effect of communication for coordination efforts is much greater than would be suggested by rational choice theories of behavior. When communication is costless and nonbinding (what is termed "cheap talk"), it may fail to sustain coordination among rational actors (Farrell 1995). However, in a world of adaptive behavior, communication can be more robust in fostering coordination. The broad slopes of a rugged landscape induce strong path dependencies. Initial, partially believed communication may bring a sufficient number of actors to initially coordinate so as to sustain a subsequent process of convergence towards full coordination on cooperative behavior.

Strategy and Landscapes Representations

Similarly, a clearly articulated business strategy has the important virtue of directing disparate elements of an organization in a coherent, self-reinforcing direction (Porter 1996). The topography of a landscape is quite sensitive to the representation of actions and, in turn, to one's understanding of action-outcome linkages (Gavetti and Levinthal 1999). To naïve strategists, the world may look quite rugged, with local peaks of possible constellations everywhere they look. However, to strategists with more powerful analytical structures with which to inform their representation of the fitness landscape, a few striking possible sets of actions may be readily apparent.

The firm may, for instance, consider three distinct bundles of policies associated with ideas of cost-leadership, differentiation, and focus (Porter 1980). The articulation of these generic strategies transforms what is otherwise an extraordinary rugged terrain of a vast number of local peaks, to a relatively smooth surface with perhaps three local peaks corresponding to the coherent articulation of each of the three generic strategies.

Powerful analytical representations of the fitness landscape that reduce the dimensionality and, in turn, the complexity of the space, provide a strong guide to action (Gavetti and Levinthal 1999). Obviously, to the extent that the representation does not capture the essential structure of the true fitness landscape, it will be a misleading guide.

However, strategy as a sense of vision for the organization can also have a value even when it is wrong, by creating the perception of payoff gradients and thereby mobilizing action. It can serve to create an illusionary surface—a social construction as opposed to a caricature of an actual fitness landscape. As an extreme example, consider the story retold by Weick (1987) of an army group stranded in a storm-covered Alps. The group despairs until one member of the party discovers a map. With this new guidance in hand, calm prevails and they find their way to safety. After arriving in safety, they find that the map depicted the Pyrenees and not the Alps as they had believed.

How then could such a map be of any value? The map created a postcard landscape—an image of reality that, in retrospect, was found to have only a spurious connection to the actual topography. Nonetheless, the belief in the map gave the group a perceived gradient with which to motivate and guide their collective action. In the absence of any knowledge of the actual topography, the landscape would appear completely unstructured to the group. Effectively, such a landscape appears flat to the actors since movement in any direction has no clear payoff implications. Faced with such a landscape, the only choices are to remain fixed or wander aimlessly. It was the amplification of this otherwise flat topography with the misconstrued map that provided the basis for action.

Credit Assignment and Accounting Systems

A critical challenge in managing interdependent systems is the problem of credit assignment (Holland 1975)—how can credit for overall performance be allocated to individual actors or actions. In large measure, this is the role of managerial accounting systems. In many cases, the payoff that the firm receives from the market provides little feedback to particular organizational units. Fundamentally, there are two bases for this loose coupling between actions and market payoffs. One is temporal. Actions taken in the present may have performance impacts at distant points in time. The other is spatial. The payoff to the actions of one element of the organization may be dependent on the actions of other elements of the organization.

New accounting practices, such as Activity-Based Costing (Cooper and Kaplan 1991), single out different decompositions of experience (e.g., moving from a cost center decomposition to process-based decomposition). Of course, different decompositions project different postcard landscapes and thus drive different adaptation processes, i.e., people climb different hills. Different decompositions also imply that different interactions are being taken into account, and thus may effect the ruggedness of perceived landscapes. In general, postcard landscapes resulting from partial decompositions tend to take into account less interactions than are present in the true, underlying fitness landscape. As a result, these decompositions tend to make the landscape smoother, and more readily scaleable. However, this smoothing of the topography comes at the cost of potentially misdirecting individual and collective efforts (Kerr 1975, Baker et al. 1994).

At the same time, efforts to build a more fully integrated accounting system, such as the balanced scorecard (Kaplan and Norton 1992), may be dysfunctional for a locally adaptive system. An accounting system that reflects the full array of interaction effects may be so rugged as to inhibit effective local adaptation. For an omniscient actor, performance can only improve with the accuracy with which the accounting system corresponds to the actual landscape. However, for an actor with more limited vision, the increased accuracy may come at a cost of exacerbating the ruggedness of the landscape to such a degree that local adaptation is made less effective.

Conclusion

While ideas of self-organization have captured enormous attention and excitement, these enthusiasms have left practitioners and organizational theorists with a puzzle. How are such self-organizing systems to be controlled and directed? Is top management to step aside and watch the "beaker" of self-organizing systems operate and observe the resulting emergent structures?

Economists have long had an answer to such questions—incentive structures that reward, but do not necessarily direct, local action. We share the economist's interest in the design of contexts that guide local action, but we differ in our premises about individual behavior. The same informational limitations that necessitate that higher-level actors (i.e., principals) delegate authority to lower-level actors also implies the boundedness of the decision-making process of these lower-level actors. Thus, our concern with context, or landscape design, necessarily involves a consideration of issues of individual and collective search processes, as well as organizational structure and incentive systems.

The principal design tool on which we focus is the manipulation of interdependencies. Rather than take interdependencies as fixed and exogenous and, as a result, the constraint under which questions of organizational design are considered, we characterize how the manipulation of interdependencies influences the process of organizational adaptation. Given the preliminary nature of this work, what we offer at this point are a new set of concepts and tools with which to consider questions of organizational design rather than a set of recipes for optimal organizational design.

The *NK(C)* structure captures, at an abstract level, the fundamental problem of organizational design. Organizations are, at their most fundamental level, about the patterned interaction among individuals (Weick 1979). At the same time, obviously we are abstracting from much of the richness of these linkages among individuals. Thus, the *NK* model offers both the power and limitations of a canonical representation.

Some of the basic principles of landscape design that come out of this analysis, can be summarized as follows.

The diversity of behavior rises with interdependencies. This implies that heightening interdependencies is an attractive feature when nonincremental search is desired. Robust designs with low interdependence and, as a result, with clear attractors are useful when predictability is desired.

- Local, incremental adaptation requires tuning down interdependencies with other actors. In the absence of tuning down interdependencies, local adaptation is likely to get trapped at a local peak conditioned on others' behavior. Thus, harnessing the power of local adaptation not only requires autonomy, it also requires low interdependence.
- Recombination of behaviors is a useful search mechanism on a multi-peak landscape. Incremental search is sufficient with low interdependence; however, with a high degree of interdependence, incremental search is not an effective mechanism.
- Landscape design is concerned with rewards (topography) far from collective and individual peaks, while classical incentive alignment focuses on the marginal incentives around the global optimum. It is for this reason that issues of small wins and positive reinforcement are important in the effective dynamics of coupled behavior.
- For coupled landscapes, the synchronization of behavior is critical. Synchronization of cooperative behavior creates a positive, reinforcing dynamic.

Good landscape designers do not direct the flow of traffic by means of directives, such as "don't walk on the grass"; rather, they induce behavior by laying out attractive pathways. Similarly, the design of an organizational landscape implies the creation of contexts that facilitate the intelligent adaptation of individual actors. Patterns of interaction are manipulated so as to induce effective collective adaptation. Put simply, effective self-organization requires good design.

Acknowledgments

Research support has been provided by the Sol Snider Center for Entrepreneurship and by MURST and CNR grants. The authors have benefitted substantially from the thoughtful comments of Michael Cohen, Andrea Bonaccorsi, the senior editor, Phil Anderson, and those of the three reviewers, as well as seminar participants at the Haas School of Business, University of California, Berkeley, and the University of Pisa.

Endnotes

[1] There may be exceptions if the individual contribution to fitness, although independent, is multipeaked in itself (e.g., there are two equally fit ways to perform a task). Even in this case, however, individual improvements would always yield collective improvements.

[2] In the context of economic analyses, the analogous concept ois that of multiple equilibria. A set of actions is in equilibrium if no individual has an incentive to change his or her behavior; however, there may exist other constellations of actions that also comprise an equilibrium set of behaviors, possibly even yielding higher payoffs to the set of individuals.

[3] Ironically, given the endless debate within economics as to the appropriate assumptions about individual behavior (Friedman 1953, Simon 1978), it is precisely in such settings of robust economic design that there are no enduring differences between actors who are fully rational and actors whose intelligence is only local. Recently, Goode and Sunder (1993) have shown how even "zero-intelligence agents" can converge toward market equilibrium in well-structured markets.

[4] At first blush, ideas of "idiot-proof" design may appear to be a reembodiment of Taylor's trained ape metaphor (Philip 1927). While it is true that both systems try to put a high degree of intelligence in the task environment, this equating of the two systems is not valid. Tayloristic environments are designed for workers trained to repeat the same, predesigned sequence of elementary movements; in contrast, Toyota-style environments are conceived as a mechanism to foster simple adaptation by individual actors. Tayloristic environments attempt to obtain predictable behavior by oversimplifying the knowledge required of workers (among other design features), while the ultimate goal of robust design is to provide a context in which the individual actor can engage in effective adaptation in which the ultimate outcomes are not specified, or even known by the designer.

[5] See Glance and Huberman (1994) for a related point.

References

Alexander, Christopher. 1964. *Notes on the Synthesis of Form.* Harvard University Press, Cambridge, MA.

Argyris, Chris, Donald Schon. 1978. *Organizational Learning.* Addison-Wesley, Wesley, MA.

Arrow, Kenneth, Frank Hahn. 1971. *General Competitive Analysis.* Holden-Day, San Francisco.

Axelrod, Robert. 1984. *The Evolution of Cooperation.* Basic Books, New York.

——, Will Mitchell, Robert E. Thomas, Scott Bennett, Erhard Bruderer. 1995. Coalition formation in standard-setting alliances. *Management Sci.* **41** 1493–1508.

Baiman, Stanley. 1982. Agency research in managerial accounting: A survey. *J. Accounting Literature* **1** 154–213.

Baker, George, Robert Gibbons, Kevin J. Murphy. 1994. Subjective performance measures in optimal incentive contracts. *Quart. J. Econom.* **109** 1125–1156.

Baldwin, Carliss, Kim Clark. 1997. Managing in an age of modularity. *Harvard Bus. Rev.* **75** 84–93.

Camerer, Colin, Marc Knez. 1997. Coordination in Organization: a Game-Theoretic Perspective. In Z. Shapira, ed. *Organizational Decision Making.* Cambridge University Press, Cambridge, England.

Cohen, Michael. 1981. The power of parallel thinking. *J. Econom. Behavior Organ.* **1** 285–306.

Cooper, Robin, Robert Kaplan. 1991. Profit priorities from activity-based costing. *Harvard Bus. Rev.* **69** 130–135.

Cooper, R., D. DeJong, R. Forsythe, T. Ross. 1992. Communication in coordination games. *Quart. J. Econom.* **107** 739–771.

Crozier, Michael. 1964. *The Bureaucratic Phenomenon.* University of Chicago Press, Chicago.

Cyert, Richard, James G. March. 1963. *A Behavioral Theory of the Firm.* Prentice-Hall, Englewood Cliffs, NJ.

Dawes, Robin M. 1980. Social Dilemmas. *Annual Rev. Psych.* **31** 169–193.

——, Richard Thaler. 1988. Cooperation. *J. Econom. Perspectives* **2** (3) 187–97.

Day, George S. 1990. *Market Driven Strategies.* Free Press, New York.

Dyer, Jeffrey H. 1994. Specialized supplier networks as a source of competitive advantage: Evidence from the auto industry. *Strategic Management J.* **17** 271–291.

Farrell, Joseph. 1995. Talk is cheap. *Amer. Econom. Rev.* **85** 186–190.

Friedman, Milton. 1953. *Essays in Positive Economics.* University of Chicago Press, Chicago.

Gavetti, Giovanni, Daniel Levinthal. 1999. Looking forward and looking Backward: Cognitive and Experiential Search. Working Paper, University of Pennsylvania.

Glance, Natalie, Bernardo Huberman. 1994. Dynamics of social dilemmas. *Sci. Amer.* **270** 76–81.

Goldberg, David. 1989. *Genetic Algorithms in Search, Optimization and Machine Learning.* Addison-Wesley Press, Reading, MA.

Goode, Dhananjay, Shyam Sunder. 1993. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *J. Political Econom.* **101** 119–137.

Hardin, G. 1968. The Tragedy of the Commons. *Sci.* **162** 1243–1248.

Holland, John H. 1975. *Adaptation in Natural and Artificial Systems.* University of Michigan Press, Ann Arbor?

Holt, C. A. 1995. Industrial organization: a survey of laboratory research. J. H., Kagel, A. E. Roth, eds., *The Handbook of Experimental Economics.* Princeton University Press, Princeton, NJ.

Hurwicz, L. 1973. The design of mechanisms for resource allocation. *Amer. Econom. Rev.* **63** 1–30.

Imai K., Nonaka I., Takeuchi H. 1985. Managing the new product-development process: How Japanese companies learn and unlearn. K. Clark, R. Hayes, C. Lorenz, eds. *The Uneasy Alliance.* Harvard University Press, Cambridge, MA.

Kaplan, Robert, David Norton. 1992. The balanced scorecard—measuring the drivers of performance. *Harvard Bus. Rev.* **70** 71–79.

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

Kerr, Steven. 1975. On the folly of rewarding A, while hoping for B. *Acad. Management J.* **18** 769–83.

Kotler, Philip. 1997. *Marketing Management.* Prentice-Hall, New York.

Ledyard J. O. 1995. Public goods: A survey of experimental research. J. H. Kagel, A. E. Roth, eds., *The Handbook of Experimental Economics.* Princeton University Press, Princeton, NJ.

Levinthal, Daniel. 1988. A survey of agency models of organizations. *J. Econom. Behavior Organ.* **9** 153–185.

——. 1997. Adaptation on rugged landscapes. *Management Sci.* **43** 934–950.

Levitt, Barbara, James G. March. 1988. Organization learning. *Annual Rev. Sociology,* **14** 319–340.

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

——, Herbert A. Simon. 1958. *Organizations.* John Wiley, New York.

Marschak, Jacob, Roy Radner. 1972. *Economic Theory of Teams.* Yale University Press, New Haven, CT.

Milgrom, Paul, John Roberts. 1990. The economics of modern manufacturing. *Amer. Econom. Rev.* **80** 511–528.

Nalbantian, H. R., A. Schotter. 1997. Productivity under group incentives: An experimental study. *Amer. Econom. Rev.* **87** 314–41.

Philip, Andre. 1927. *Le Probleme Ouvrier aux Etats Units.* Alcan, Paris.

Pine, J. B., II. 1993. *Mass Customization: The New Frontier in Business Competition.* Harvard Business School Press, Cambridge, MA.

Porter, Michael E. 1980. *Competitive Strategy.* Free Press, New York.

——. 1996. What is strategy? *Harvard Bus. Rev.* **74** 61–78.

Ohno, Kenichi. 1992. *Dynamism of Japanese Manufacturing: Evidence from the Postwar Period.* Brookings Institute, Washington, DC.

Ostrom, Elinor, James Walker, Roy Gardner. 1992. Covenants with and without a sword: Self-governance is possible. *Amer. Polit. Sci. Rev.* **86** 404–417.

Rocco E., M. Warglien. 1996. Computer-mediated communication and the emergence of "electronic opportunism." Computable and Experimental Economics Laboratory, University of Trento, WP 1996-01.

Rossi A., M. Warglien, E. Zaninotto. 1997. Cooperation as illusory hill-climbing: Coadaptation and search in social dilemmas. R. Conte, R. Hegselmann, P. Terna, eds. *Simulating Social Phenomena.* Springer, Berlin.

Schumpeter, Joseph. 1934. *The Theory of Economic Development.* Harvard University Press, Cambridge, MA.

Shingo, S. 1983. *Maitrise de la Production et Méthode Kanban.* Les Editions d'Organisation, Paris.

Simon, Herbert A. 1955. A behavior model of rational choice. *Quart. J. Econom.* **69** 99–118.

——. 1978. Rationality as process and product of thought. *Amer. Econom. Rev.* **68** 1–16.

——. 1982. *The Sciences of the Artificial.* MIT Press, Cambridge, MA.

——, Albert Ando. 1961. Aggregation of variables in dynamic systems. *Econometrica* **29** 111–138.

Smith, J. M. 1989. *Evolutionary Genetics.* Oxford University Press, New York.

Spence, A. Michael. 1974. *Market Signaling: Informational Transfer in Hiring and Related Screening Processes.* Harvard University Press, Cambridge, MA.

Stinchcombe, Arthur. 1965. Social structure and organizations. J. March, ed. *Handbook of Organizations.* Rand McNally, Chicago.

Thompson, James D. 1967. *Organizations in Action.* McGraw-Hill, New York.

Tirole, J. 1986. Hierarchies and bureaucracies: On the role of collusion in organization. *J. Law, Econom. Organ.* (Fall) 181–214.

Van Huyck, J. B., R. C. Battalio, R. C. Beil. 1991. Strategic uncertainty, equilibrium selection, and coordination failure in average opinion games. *Quart. J. Econom.* **106** 885–911.

Von Hippel, Eric. 1990. Task partitioning: An innovation process variable. *Res. Policy* **19** 407–418.

Warglien, M. 1995. Hierarchical selection and organization adaptation. *Indus. and Corporate Change* **4** 161–187.

Weick, Karl E. 1979. *The Social Psychology of Organizing.* McGraw-Hill, New York.

——. 1984. Small wins: Redefining the scale of social problems. *Amer. Psychologist* **39** 40–49.

——. 1987. Substitutes for corporate strategy. D. Teece, ed. *The Competitive Challenge.* Ballinger Publishing, Cambridge, MA.

Weisbuch, Gerard. 1991. *Complex Systems Dynamics.* Addison Wesley Publishing, Redwood City, CA.

Wheelwright S. C., Kim B. Clark. 1992. *Revolutionizing Product Development.* The Free Press, New York.

Womack J. P., D. T. Jones, D. Roos. 1990. *The Machine that Changed the World.* Rawson Associates, New York.

Wright, Sewall. 1932. The roles of mutation, inbreeding, crossbreeding and selection in evolution. *Proc. XI Internat. Congress of Genetics* **1** 356–66.

Accepted by Phil Anderson; received October 1997. This paper has been with the authors for one revision.

LINKED CITATIONS

- Page 1 of 2 -

You have printed the following article:

Landscape Design: Designing for Local Action in Complex Worlds
Daniel A. Levinthal; Massimo Warglien
Organization Science, Vol. 10, No. 3, Special Issue: Application of Complexity Theory to Organization Science. (May - Jun., 1999), pp. 342-357.
Stable URL:
http://links.jstor.org/sici?sici=1047-7039%28199905%2F06%2910%3A3%3C342%3ALDDFLA%3E2.0.CO%3B2-6

Endnotes

[3] **A Behavioral Model of Rational Choice**
Herbert A. Simon
The Quarterly Journal of Economics, Vol. 69, No. 1. (Feb., 1955), pp. 99-118.
Stable URL:
http://links.jstor.org/sici?sici=0033-5533%28195502%2969%3A1%3C99%3AABMORC%3E2.0.CO%3B2-A

References

Subjective Performance Measures in Optimal Incentive Contracts
George Baker; Robert Gibbons; Kevin J. Murphy
The Quarterly Journal of Economics, Vol. 109, No. 4. (Nov., 1994), pp. 1125-1156.
Stable URL:
http://links.jstor.org/sici?sici=0033-5533%28199411%29109%3A4%3C1125%3ASPMIOI%3E2.0.CO%3B2-L

Communication in Coordination Games
Russell Cooper; Douglas V. DeJong; Robert Forsythe; Thomas W. Ross
The Quarterly Journal of Economics, Vol. 107, No. 2. (May, 1992), pp. 739-771.
Stable URL:
http://links.jstor.org/sici?sici=0033-5533%28199205%29107%3A2%3C739%3ACICG%3E2.0.CO%3B2-2

NOTE: *The reference numbering from the original has been maintained in this citation list.*

LINKED CITATIONS

Specialized Supplier Networks as a Source of Competitive Advantage: Evidence from the Auto Industry

Jeffrey H. Dyer

Strategic Management Journal, Vol. 17, No. 4. (Apr., 1996), pp. 271-291.

Stable URL:

http://links.jstor.org/sici?sici=0143-2095%28199604%2917%3A4%3C271%3ASSNAAS%3E2.0.CO%3B2-I

The Tragedy of the Commons

Garrett Hardin

Science, New Series, Vol. 162, No. 3859. (Dec. 13, 1968), pp. 1243-1248.

Stable URL:

http://links.jstor.org/sici?sici=0036-8075%2819681213%293%3A162%3A3859%3C1243%3ATTOTC%3E2.0.CO%3B2-N

Exploration and Exploitation in Organizational Learning

James G. March

Organization Science, Vol. 2, No. 1, Special Issue: Organizational Learning: Papers in Honor of (and by) James G. March. (1991), pp. 71-87.

Stable URL:

http://links.jstor.org/sici?sici=1047-7039%281991%292%3A1%3C71%3AEAEIOL%3E2.0.CO%3B2-N

A Behavioral Model of Rational Choice

Herbert A. Simon

The Quarterly Journal of Economics, Vol. 69, No. 1. (Feb., 1955), pp. 99-118.

Stable URL:

http://links.jstor.org/sici?sici=0033-5533%28195502%2969%3A1%3C99%3AABMORC%3E2.0.CO%3B2-A

Aggregation of Variables in Dynamic Systems

Herbert A. Simon; Albert Ando

Econometrica, Vol. 29, No. 2. (Apr., 1961), pp. 111-138.

Stable URL:

http://links.jstor.org/sici?sici=0012-9682%28196104%2929%3A2%3C111%3AAOVIDS%3E2.0.CO%3B2-Z

Strategic Uncertainty, Equilibrium Selection, and Coordination Failure in Average Opinion Games

John B. Van Huyck; Raymond C. Battalio; Richard O. Beil

The Quarterly Journal of Economics, Vol. 106, No. 3. (Aug., 1991), pp. 885-910.

Stable URL:

http://links.jstor.org/sici?sici=0033-5533%28199108%29106%3A3%3C885%3ASUESAC%3E2.0.CO%3B2-0

NOTE: *The reference numbering from the original has been maintained in this citation list.*