University of Pennsylvania
ScholarlyCommons

Management Papers

Wharton Faculty Research

9-2008

Choice Interactions and Business Strategy

Pankaj Ghemawat

Daniel A. Levinthal
University of Pennsylvania

Follow this and additional works at: https://repository.upenn.edu/mgmt_papers

Part of the [Business Administration, Management, and Operations Commons](#), and the [Strategic Management Policy Commons](#)

Recommended Citation

Ghemawat, P., & Levinthal, D. A. (2008). Choice Interactions and Business Strategy. *Management Science*, 54 (9), 1638-1651. <http://dx.doi.org/10.1287/mnsc.1080.0883>

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/mgmt_papers/139
For more information, please contact repository@pobox.upenn.edu.

Choice Interactions and Business Strategy

Abstract

Choice settings are strategic to the extent that they entail cross-sectional or intertemporal linkages. These same factors may impose daunting demands on decision makers. We develop a graph-theoretic generalization of the NK model of fitness landscapes to model the way in which policy choices may be more or less strategic. We use this structure to examine, through simulation, how fully articulated a strategy or set of policy choices must be to achieve a high level of performance and how feasible it is to offset past strategic mistakes through tactical adjustments (instead of alignment). Our analysis highlights the role of asymmetry in the interaction of strategic choices and in particular the degree to which choices vary in terms of being influential, dependent, or autonomous from other choices.

Keywords

strategic choice, activity systems, fitness landscapes, choice interactions, path dependence

Disciplines

Business Administration, Management, and Operations | Strategic Management Policy

Choice Interactions and Business Strategy

Prof. Pankaj Ghemawat
Morgan Hall 227
Harvard Business School
Boston, MA 02163
617-495-6286
pghemawat@hbs.edu

Prof. Daniel Levinthal
2035 Steinberg-Dietrich Hall
The Wharton School
University of Pennsylvania 19104
215-898-6826
Levinthal@wharton.upenn.edu

October 2002

We wish to thank seminar participants at Harvard Business School, INSEAD, the University of Illinois at Urbana-Champaign and the University of Southern California as well as the Department Editor, Rebecca Henderson, and three anonymous reviewers for their thoughtful comments on a prior draft.

Choice Interactions and Business Strategy

Abstract:

Strategists have tended to explain sustained performance differences across firms in terms of two types of interactions among choices: cross-sectional interactions and longitudinal ones. We explore the interplay between these two sorts of forces first in a qualitative manner drawing on case of the case of the Vanguard Mutual Fund. We then develop a graph-theoretic generalization of the NK model of Kaufman (1993) in order to examine these questions in a more structured manner. We use this structure to examine, through simulation, how fully articulated a strategy or set of policy choices must be to achieve a high level of performance, and how feasible it is to offset past strategic mistakes through tactical adjustments (instead of alignment). Our analysis highlights the role of asymmetry in the interaction of strategic choices and in particular the degree to which choices vary in terms of being influential, contingent, or autonomous from other choices.

Choice Interactions and Business Strategy

1. Introduction

Differences in perspective on how choices interact to determine performance underlie some of the fundamental debates in the strategy field. Thus, cross-sectionally minded researchers stress that strategies (should) consist, at any point in time, of tightly linked, complementary policy choices across a firm's full array of operational possibilities (Porter, 1996). Longitudinally-minded researchers, such as those who stress the importance of firms' resource stocks, strategic commitments or capability development trajectories (Wernerfelt [1984], Ghemawat [1991], and Teece and Pisano [1994]), place more emphasis on temporal, as opposed to cross-sectional, interactions among choices. Strategy process researchers pose even more complex temporal interactions among choices, the resolution of which is complicated by cognitive as well as motivational constraints.

This paper develops and analyzes a simple formal structure that connects choices to performance in a structured way that captures both important elements of the cross-sectional interactions among policy choices and simple temporal linkages as well. Procedurally, this structure is embedded in a context that allows for ex-ante strategic choice *and* the subsequent emergence of strategic positions through processes of incremental search (Gavetti and Levinthal, 2000). This exercise lets us explore, among other things, the completeness with which strategies must be specified to achieve satisfactory performance and the performance costs of irreversible mistakes in choices of different sorts.

Section 2 anchors our discussion of interactions among choices, both cross-sectionally and inter-temporally, in the context of a particular case, that of the Vanguard mutual fund company. The case study highlights the cross-sectional multiplicity of choices, the critical role that the temporal ordering of choices can play in organizing them, as well as the basic idea that not all choices should be treated as being equally important: some are more "strategic" than others.

With this case as backdrop, section 3 develops a stylized formal structure that captures important elements of the cross-sectional structure of policy choices and, at the same time, lets us consider simple temporal linkages as well. This involves a generalization of the framework of

NK fitness landscapes introduced by Kauffman (1993).¹ The paper examines not just the random interaction structure explored by Kauffman (1993) and others (Levinthal, 1997; Rivkin, 2000), but also two structured interaction patterns: one in which choices nest hierarchically in their influence upon one another and one in which choices vary in the degree to which they are central or peripheral to one another.

These alternative structures capture a sense of some choices as being more or less “strategic” than others. However, the degree to which one choice or another is strategic may have a temporal dimension as well. The paper focuses on perhaps the most basic temporal quality of a choice—the reversibility or irreversibility of the corresponding policy decision. This potential temporal rigidity, in conjunction with cross-sectional linkages among choices, imposes important constraints on the pattern of strategic adaptation that we explore with the generalized NK structure.

Section 4 consists of two sets of analyses. One set of analyses focuses on the question of the completeness with which a strategy must be articulated for it to lead to good performance. Discussions of cross-sectional linkages often presume that a coherent system of policy choices is arrived at by some process of *a priori* theorizing. Given the rich and complex web of interactions in such a system, the power of such *a priori* theories would seem to be limited. A more plausible characterization is that a firm makes a few choices about how it will compete and these choices, in turn, influence subsequent decisions. A central question regarding the emergence of a coherent, and possibly profitable, activity system is how fully specified these initial choices have to be. Is it possible to specify a few key choices and for other policies to emerge through some more incremental process? Or, does the rich set of interrelationships among policy choices require rather more complete specification *a priori*?

The second set of analyses in section 4 examines the downside rather than the upside of the effect of initial positioning in policy space. In a dynamic world in which ideal policy sets change with time, how constraining and damaging to performance are precommitments through

¹ The nature of NK fitness landscapes is described in detail in section 3. For now, it suffices to note that it consists of a simple analytical structure that allows the modeler to “tune” the degree of interaction among policy choices in determining the payoff to a vector of choices, with the parameter N denoting the number of choices and K the number of these that interrelate in determining performance.

their irreversibility? To what extent does the impact of these existing commitments vary with how “strategic” the policy choice is in a cross-sectional sense of hierarchy or centrality?

Section 5 extends the basic analysis in ways that indicate that it is generally robust, and section 6 concludes.

2. The Case of Vanguard

In order to flush out some key modeling criteria, it is useful to ground our discussion in a (somewhat) detailed case history. The one that we will consider in this section concerns Vanguard. The Vanguard Group is a complex of mutual funds that rated first in *Barron's* 1996 rankings of 5-year and 10-year performance by mutual fund families, grew the assets under its management at an annual rate of 31% between 1980 and 1996, compared to 23% for the U.S. mutual fund industry as a whole, and had achieved an expense ratio of 0.3%, compared to 1.0% for a conventional mutual fund. Our discussion of this case is informed almost entirely by a careful longitudinal study of the company by Siggelkow (1998 and 2002a), although the interpretations offered are our own responsibility. To preview our conclusions, we read the Vanguard case as suggesting that it is important to recognize both the multiplicity of choices and the fact that some of them matter more than others because they condition the value of other choices or because of their temporal ordering.

It is simplest to start at the end of the period spanned by Siggelkow's study of Vanguard, with his characterization of Vanguard's activity system in 1997 (see **Figure 1**). To begin by pointing out the obvious, Figure 1 is rather complex. There are thirty-seven light circles and six dark ones, with the dark shading indicating what Porter (1996), who initiated such mapping exercises, called “higher-order strategic themes” that set the context for decisions concerning the light circles, and the circles' labels generally describing departures from industry norms. Many of the circles fit with each other rather than misfitting (as indicated by the solid rather than dashed lines connecting them). There are also six squares denoting distinct product categories, some of which fit with the circles (solid lines), and some of which don't (dashed lines).

By highlighting the cross-sectional complexity of the choices embedded in the Vanguard way of doing things, Figure 1 inevitably underscores the importance of such interactions among choices. As a result, the Vanguard case, among others, has been cited in support of the position

that the cross-sectional interactions among choices are, in some sense, more fundamental than other kinds of interactions and that the highest strategic priority is, therefore, to build tightly-coupled activity systems that exploit such interactions (Porter, 1996).

But while the complexity depicted in Figure 1 can be, and has been, interpreted as an instantiation of the importance of cross-sectional linkages, a longitudinal look at Vanguard is also in order. Vanguard's predecessor, the Wellington Fund, was set up in 1928 and survived the financial crash of 1929 because of its conservatism, but is not reported to have otherwise distinguished itself: assets under its management stagnated after crossing the \$2 billion mark in 1965 and, in the early 1970s, declined substantially (reflecting, in part, general bearishness in U.S. equity markets). Siggelkow's map of Wellington in 1973 contained only four circles, and the single connection between them was a dashed line, indicating misfit. But in 1974, Wellington's successor, the Vanguard Group was incorporated as a mutual holding company in which the shareholders of the underlying funds would own the managing fund complex.² After this unusual—and still unique³—choice of organizational form, progress in articulating the other characteristic features (defined as fundamental or distinguishing aspects) of the Vanguard system was rapid, although they continued to be filled in through the 1980s and 1990s (see Figure 2).⁴

The reasons why the choice of organizational form in 1974 mattered so much are fairly obvious. First of all, administrative services shifted from being a source of profits for the fund manager to being a “cost center” shared by the underlying mutual funds. Not only did this shift eliminate the substantial 40% mark-up on the provision of these services that Wellington Management Company had enjoyed but, perhaps more importantly, it provided motivation for further efforts at cost reduction. As a “true mutual,” all reduction in expenses, whether fund administration or investment management fees, would accrue to fund shareholders and, in turn, enhance the financial return of Vanguard's funds.

² The term mutual fund refers to the joint holding of investment assets. However, with the exception of Vanguard, all “mutual funds” are structured such that shareholders in the fund have no ownership of the entity that manages and administers the investment assets. As Siggelkow (1998) notes, the choice of this organizational form by Vanguard's CEO, John Bogle, may have been motivated by a desire to become autonomous from Wellington Management Company (the investment management company for the funds formerly associated with Vanguard and the company from which Bogle had been fired as president at the beginning of 1974) as much as by a sense of appropriate strategy choice.

³ Siggelkow notes that one other fund had a similar structure when Vanguard adopted this organizational form, but subsequently reverted to the more common structure, involving the management and administration of investment assets by an entity independent of the fund family.

⁴ Here, our characterization differs from Siggelkow's: “Vanguard slowly adopted a very consistent set of choices.”

A low-cost configuration therefore became a “natural attractor” for Vanguard. In 1976, Vanguard introduced the first indexed equity fund, a product category that it subsequently came to dominate. Index funds, as well as bond and money market funds, were natural investment vehicles for a fund complex that was choosing to compete on the basis of low costs since variations in the performance of such funds were largely driven by differences in the various fees charged to fund shareholders.

Vanguard’s choice of configuration affected not only its product offerings, but also how it distributed its products and managed its investments. In 1977, Vanguard shifted from relying on broker-dealers to distribute its funds and began to distribute its funds directly to consumers on a no-load basis. Having shifted to the in-house distribution of funds, Vanguard achieved greater autonomy from its primary investment management advisor, Wellington Management Company.⁵ Vanguard could now both bargain more effectively with its advisors and, ultimately, in 1981, bring some advisory activities in-house.⁶ Finally, in 1978 Vanguard internalized transfer agency activity. Its motivations included the fact that its own activities had expanded to an extent that allowed it to achieve scale economies and the desire to control the quality of customer relations in absence of a system of broker-dealers.

To summarize, by the early to mid-1980s, Vanguard had developed a set of policies with “tight fit” that incorporated many of the characteristics that it was known for in the late 1990s, particularly the key ones (the dark circles in Figure 1). This outcome does not, however, appear to have been the product of either *ex ante* design or a purely emergent process of discovery. A plausible alternative deconstruction is that a radical choice of organizational structure—a mutual form of organization—was embedded in the way Vanguard was set up, and apparently directed management to subsequent choices that “fleshed out” the activity map revealed by current cross-sectional analysis.

In addition to highlighting the possibility that longitudinal as well as cross-sectional interactions among choices may be very important, the Vanguard example calls attention, more generally, to the asymmetry of choices: to the idea that while many choices may impinge on performance, they are usually not all of equal importance. We go on to explore the implications

⁵ Investment management companies typically control the distribution of the mutual funds for which they provide investment advice, making it quite difficult to shift advisors.

⁶ Vanguard brought in-house the relatively “plain vanilla” advisory activity of the management of fixed-income funds and index funds, continuing to rely on outside advisors for the management of actively managed equity funds.

of such choice asymmetries, first in a purely cross-sectional context and second in a context that adds in explicitly longitudinal elements. However, before engaging in such analysis, we need to characterize the nature of interactions among choices.

3. Modeling Interactions

Grappling with interactions among choices poses challenges for both decision making agents as well as for those modeling their behavior—what Bellman (1957), one of the progenitors of dynamic programming, described as “the curse of dimensionality.” The difficulties are twofold. Even within a purely cross-sectional frame, rich interactions among a large number of choices imply the nonexistence of a general, step-by-step algorithm that can locate the best set of choices in a “reasonable” period of time (i.e., a polynomial function of the number of variables) (Lewis, 1985; Rivkin, 1997). And from a longitudinal or dynamic perspective, such situations generally do not lend themselves to “pushing forward” from multi-dimensional historical trajectories to identify an optimal path through these possible histories (Sussman, 1975).

The challenge of modeling interdependent choices has recently received additional attention in the economics and management literatures. One approach has been to focus on a very special choice structure, involving supermodularity, in which choices along any two dimensions are pairwise complementary for all values of the choice variables involved, and for all values of other choice variables. Topkis (1978 and 1995) and Milgrom and Roberts (1990 and 1995) have used the resulting lattice models to show that these are the weakest conditions under which it is possible to obtain monotone comparative static predictions linking shifts in optimal choices concerning sets of variables to changes in underlying parameters. How weak these conditions are in absolute terms is another matter: tradeoffs or substitution effects are ruled out, as are reversals between substitution and complementarity as the values of relevant variables change and, consequently, limitations are placed on the number of “best ways to compete” (local peaks on the fitness landscape, as elaborated below.) If one believes, as some strategists (e.g., Porter, 1996) do, that the interplay between complementarities *and* trade-offs across multiple activities is critical to the possibility of “many best ways to compete”, then allowing only global complementarities seems very constricting.

The other response to the problem of multiple, linked choices that has commanded considerable attention recently has been to build on the NK-simulation approach pioneered by Kauffman (1993) in evolutionary biology (cf., Levinthal, 1997 and Rivkin, 2000). Kauffman, drawing on Wright's (1931) notion of a fitness landscape, developed this framework to explore the emergence of order among biological organisms. The model has two basic parameters, N , the total number of policy choices, and K ($< N$), the number of policy choices that each choice depends upon. More specifically, each of the choices is assumed to be binary, and choice-by-choice contributions to fitness levels are drawn randomly from a uniform distribution over $[0,1]$ for each of the 2^{K+1} distinct payoff-relevant combinations a choice can be part of. Total fitness is just the average of these N choice-by-choice fitness levels. Note that with K equal to its minimum value of 0, the fitness landscape is smooth and single-peaked: changes in the setting of one choice variable do not affect the fitness contributions of the remaining $N-1$ choice variables. At the other extreme, with K equal to $N-1$, a change in a single attribute of the organism or organization changes the fitness contribution of all its attributes, resulting in many local peaks rather than just one, with each peak associated with a set of policy choices that have some internal consistency. No local peak can be improved on by perturbing a single policy choice, but local peaks may vary considerably in their fitness levels.

The choice structure underlying the NK simulation approach generalizes Milgrom and Roberts' lattice-theoretic approach based on "complementarities" in two key respects. First, it avoids imposing a specific structure on the linkages among choices. Second, it allows the richness of such linkages to vary across situations (through the K parameter). It embodies a number of other attractions as well, most of which we will discuss and retain below. But for our present purposes, it also has one glaring defect: all choices are assumed to be equally important. This rules out, for example, asymmetries of the sort evident in the distinction between light and dark circles in Figure 1. To remedy that defect, we need more degrees of freedom than are afforded by a single interactivity parameter, K .

One way of proceeding is suggested by the observation that the activity map in Figure 1 bears more than a passing resemblance to a mathematical graph. A mathematical graph can be summarized in terms of its adjacency matrix, which specifies how different choices (the vertices in the graph) are linked (the lines in the graph). In such a matrix, choice variable j 's effect on other variables is represented by the patterns of 0's and 1's in column j , with a value of 1

indicating that the payoff associated with the variable in the row being considered is contingent or dependent on variable j , and a value of 0 denoting independence. Similarly, reading across row i in such a matrix indicates the variables on which the payoff to choice variable i is itself contingent.⁷ The principal diagonal of an adjacency matrix always consists of 1's, but the matrix itself need not be symmetric around that diagonal.

Replacement of the interactivity parameter, K , with an adjacency matrix is meant, most broadly, to allow some choices to vary in the degree to which they influence other choices or are themselves contingent on other choices. A choice is influential to the extent that the column under that policy is populated with 1's, indicating that the value of other policies depends on this choice. Conversely, a policy is dependent upon other choices to the extent that the row corresponding to that policy is populated with 1's in the adjacency matrix. A policy is relatively autonomous to the extent that neither the column nor row associated with this policy is populated with 1's.

Figures 3a and 3b present two extreme adjacency matrices. The former represents a hierarchical structure in which higher-order policy choices influence lower-order choices but the reverse is not true. The characterization of Vanguard as having been founded on the basis of a radical choice of organizational structure from which other choices naturally flowed has, roughly speaking, this quality. One might think of the corresponding adjacency matrix as containing 1's in one column (corresponding to the choice of organizing as a "true" mutual) as well as in the principal diagonal, with 0's elsewhere—in graph-theoretic terms, a star. A star graph is an extreme example of the much more general class of hierarchical choice structures. In graph-theoretic terms, hierarchies are best thought of as directed (or at least rooted) trees, with interdependencies (i.e., the 1's) populating one side of the principal diagonal. Figure 3a depicts a pure hierarchical form with 1's as all the entries to the left of the principal diagonal. Choice 1 is hierarchically the most important, choice 2 the second most important, and so on; such a structure lets us take a finer-grained look at the effects of variations in the degree of hierarchical importance than a star structure would permit.

Figure 3b, in contrast, represents an activity system with a structure in which policy choices vary in their degree of centrality, i.e., the degree to which they are mutually dependent

⁷ In addition to such direct effects on value contributions, variables may, of course, be indirectly related through other variables.

on other choices. In addition to its intuitive appeal, such a notion of centrality is responsive to the potential inferential problem that all we might be able to do is to observe linkages between choices, not the direction of influence, i.e., that in observational terms, we might have to work with undirected graphs—or in adjacency matrix terms, with matrices that are symmetric around the principal diagonal. One systematic procedure for ranking choices in terms of their importance in such a context would be to array them in terms of the number of other choices with which they interact. Application of this procedure to Figure 1's purely cross-sectional map of Vanguard, for example, would imply that “focus on low costs,” with 17 links, is more central to Vanguard's strategy than, say, the “Vanguard award for excellence,” with one such link. The particular form of centrality depicted in Figure 3b embodies a structure and a labeling scheme that has 1's as all the entries to the left of the inferior diagonal (but distributed symmetrically to the left and the right of the principal diagonal). Thus, choice 1, with links to 9 other choices, is the most central, choice 2 the second most central, and so on.

To allow us to explore systematically the range of possible adjacency matrices, we specify the following stochastic process to generate matrices. For each policy choice we specify a probability p_i^H that policy i influences other policy choices and a probability p_i^C that the payoff to this policy is in turn contingent on other policies. Thus, the likelihood of a linkage such that choice i influences policy choice j is $p_i^H p_j^C$. Or to reparametrize these variables, p_i^H and p_i^C , in a useful way, let $r_i = p_i^H / (p_i^H + p_i^C)$ represent the relative tendency towards influence as opposed to contingency and let $p_i = (p_i^H + p_i^C)$ represent the likelihood of some form of interaction as opposed to autonomy.⁸ Thus, by varying r_i from zero to one we specify the relative degree to which a policy is contingent or influential and by varying p_i from zero to one we vary the policy's degree of autonomy. Specifically, we set p_i equal to 0.5 for all i values and vary r_i from zero to one in order to explore structures in which policy choices vary in the degree to which they are influential or contingent and we explore the role of centrality by setting r_i equal to 0.5 for all i values and varying p_i from zero to one.

For all interaction structures, an organization's policy choices are represented by a vector of length N where each element of the vector can take on a value of 0 or 1 (not to be confused with the 0's and 1's that are used to denote the absence or presence of linkages between every pair of policy elements). The overall fitness landscape will then consist of 2^N possible policy

choices, with the overall behavior of the organization characterized by a vector $\{x_1, x_2, \dots, x_N\}$, where each x_i takes on the value of 0 or 1.⁹ If the contribution of a given element, x_i , of the policy vector to the overall payoff is influenced by K_i other elements—in ways that vary across the three structures we will analyze—then it can be represented as $f(x_i | x_{i1}, x_{i2}, \dots, x_{iK_i})$.

Therefore, each element's payoff contribution can take on 2^{K_i+1} different values, depending on the value of the attribute itself (either 0 or 1) the value of the K_i other elements by which it is influenced (each of these K_i values also taking on a value of 0 or 1). Specifically, we follow prior researchers and assign a random number drawn from the uniform distribution from zero to one to each possible $f(x_i | x_{i1}, x_{i2}, \dots, x_{iK_i})$ combination with the overall fitness value then being defined as $\sum_{i=1 \text{ to } N} f(x_i | x_{i1}, x_{i2}, \dots, x_{iK_i}) / N$.

A number of additional assumptions, based on prior applications, that are built into this specification should also be mentioned. First of all, there is the emphasis on choice under uncertainty. In addition to its arguable descriptive realism, initial uncertainty helps explain why an organization launched over a fitness landscape may not instantly alight on the globally optimal policy vector. Second, there is the assumption that randomness takes the form of a uniform distribution. While some might argue that this distribution is too diffuse, we retain this assumption to provide at least some basis for numerical comparability with prior work, which suggests, among other things, that the structure of the fitness landscape is not sensitive to the particular probability distribution employed (Weinberger, 1991). Third, there is the equal weighting of different choices in terms of their direct contribution (potential) to overall fitness. Solow, Bunetas, Roeder, and Greenspan (1999) explore the implications of differentially weighting the contribution of different policy variables to overall performance.¹⁰ While asymmetries in weights are clearly important, our focus here is on asymmetries in the structure of interactions and the implications for effective strategy formation of exploiting them. Finally, we should add the caveat that while the analysis highlights the effects of linkages among the

⁸ Using this parameterization, $p_i^H = p_i r_i$ and $p_i^C = p_i (1 - r_i)$.

⁹ The model can be extended to an arbitrary finite number of possible values of an attribute, but the qualitative properties of the model are robust to such a generalization (Kauffman, 1989).

¹⁰ The focus of their work is to demonstrate that sufficiently extreme weighting differences, in particular weighting the contribution of one policy by $1-e$ and the other $N-1$ variables by e for sufficiently small values of e , can allow a process of local search to reach the global optimum even under conditions of high interaction (K) across policy choices.

organization's policy choices, it does not address linkages across firms. In particular, one could imagine spatial competition (or cooperation) among firms so that the fact that one or more firms occupy a particular point on the policy landscape changes the payoff to other firms' occupying that region. Clearly, such effects exist and are important. But, for simplicity, we do not explore them in the present analysis.

We also assume that N equals 15—a level of multidimensionality that, based on a standard result in graph theory, is sufficient to generate more than 10^{19} distinct graphs.. The results that we report are averaged over 10,000 landscapes. The repetition is meant to allow for the averaging out of two kinds of purely random effects. The first reflects the set of possible adjacency matrices that may result for a given set of values of p and r ; the second results from the seeding of an organization on a given performance landscape. To address the former source of randomness, we generate 100 adjacency matrices for each vector of p and r values. Each of these 100 landscapes will have an independently drawn adjacency matrix, although based on the same p and r values. In addition, given the realized adjacency matrix, the landscapes will have a distinct seeding of fitness values. To address the latter sort of random effect, we generate 100 distinct fitness landscapes for each of these 100 adjacency matrices. In analyzing our results, we also normalize fitness levels to control for two sorts of effects. First, the magnitude of the global peak will vary from landscape to landscape, even if the landscapes share the same structural properties. As a result, the highest possible performance is specific to a particular fitness landscape. Second, it is also important to normalize with respect to what might constitute poor performance. A random point of the fitness landscape has an expected fitness level of 0.5 given the seeding of fitness levels by draws from a uniform distribution ranging from zero to one. In addition, local search processes will, by themselves, suffice to take an organization to a local peak in the landscape, of which there are typically a dozen or more in the landscapes examined here. The average value of these local peaks is one benchmark of the level of fitness that cognitively constrained choices might generate.¹¹ Thus, the value of an average local peak can

¹¹ Technically, the average value of local peaks in the landscape is not quite what would correspond to the expected outcome of local search from a random starting position. The reason for this divergence is that higher local peaks have somewhat broader basins of attraction (Kauffman, 1993). Therefore, an alternative baseline would be to locate at random on the landscape a large population of firms and let them engage in local search until each firm identified a local peak in the landscape. The problem with this approach is that in a large space of possibilities (2^N), there are a large number of possible starting positions. In addition, two organizations from the same starting position may, as a result of the stochastic nature of the search process, end up at two distinct peaks. Thus, for simplicity, we have

be thought of as a floor in the sense of corresponding to the level of performance to be expected given incremental search from a random starting point in the fitness landscape.

So in order to normalize with respect to both the global optimum and the level of performance associated with the value of the average local peak, we transform the raw fitness level obtained from the simulations, f_i , into

$$(f_i - \text{Avg}_i) / (\text{Global}_i - \text{Avg}_i),$$

where Avg_i is the average value of the local peaks on a particular landscape i , and Global_i is the value of the global peak in the given landscape. These normalized fitness levels, averaged over the 10,000 runs, are what we actually report in the subsections that follow.

4. Simulation Results

We explore the emergence of strategic positions from two perspectives, both of which involve strategic choices followed by local search over what might be described as tactical choices. We first look at the possibility, or demands, of *a priori* specification of strategies when policy choices are highly interdependent: can higher-order or strategic guidance along a few dimensions followed by tactical adjustment and alignment of the remaining dimensions through local search be expected to lead to high levels of performance? Second, we consider the impact of mistakes that are strategic, in the sense of irreversible: what are the residual costs of different types of initial misspecified choices, after local search and tactical adjustment aimed at mitigating these “mistakes”?

Strategic Guidance and Tactical Alignment

How might a complex policy system such as the cross-sectional one depicted in Figure 1 for Vanguard arise? Broadly speaking, there are two possibilities. One is through ex-ante design of a coherent and fully articulated activity system. Another possibility is via a process of search and adjustment on the fitness landscape defined by the payoff associated with different constellations of policy choices. In particular, a process of local search will identify an internally

simply calculated the average value of local peaks in each landscape as the basis for our calculation of a baseline level of performance from which to evaluate the performance returns to deliberate strategy making.

coherent set of policy choices; that is, a set of choices from which any incremental one-policy-at-a-time change would be dysfunctional, or what has been called a local peak in the fitness landscape (Kauffman, 1993). However, local peaks come without warranties as to their global, or absolute desirability, so there is no assurance that local search processes will, on their own, lead to satisfactory performance.

The actual evolution of successful strategies probably involves elements of both *ex ante* design and *ex post* adjustment. It seems implausible that a strategic position with dimensionality as high as in Figure 1 could be fully articulated *a priori*; at the same time, it seems unlikely to be purely emergent. A more plausible picture of managerial processes seems to be that while there is some top-down prespecification of both some broad principles and some particular policy choices, these represent starting points in processes aimed at improving firms' positions over time (Siggelkow, 2002a; Gavetti and Levinthal, 2000). This representation also has the attractive feature of embodying elements of both the conscious choice of strategies, in the spirit of the "content" style of strategy research, and of the emergence of strategic positions that is central to "process" discussions of strategy formulation (Mintzberg, 1978; Burgelman, 1994).

Our use of this representation is motivated by the idea that the effectiveness of strategic planning may be inversely related to the dimensionality required of a strategy to ensure the achievement of a reasonably consistent set of policies. If strategy must be defined at a detailed operational level to achieve consistency (e.g., if it must spell out the choices corresponding to all the circles in Figure 1), then the requirements for strategic planning escalate dramatically. In contrast, if a few higher-level choices make subsequent lower-level choices self-evident (e.g., if it suffices to spell out the choices corresponding to just the dark circles in Figure 1, followed by a process of local search), then the requirements for strategic planning remain relatively modest.

Figure 4 explores this issue in the following manner. A certain number of policy choices ("degree of match"), selected in decreasing order of "strategic" importance (with reference to the hierarchical and centrality structures), are set to equal their value at the global optimum, and the initial values of the remaining policy choices are specified at random. These remaining policies are then modified by a process of local search. Local search (March and Simon, 1958; Cyert and March, 1963) involves the comparison of an existing policy choice with adjacent, or neighboring choices. This process is operationalized here as involving the comparison of the current policy vector with all the other policy vectors that differ from the current vector in terms of just one

choice element.¹² If a superior alternative is identified in the immediate neighborhood of the existing policy array, it is adopted. In subsequent periods, more local search follows until no further replacement that immediately enhances fitness values can be found. This dynamic leads, inexorably, to local peaks in the fitness landscape (Levinthal, 1997). Thus, the choice variables that are correctly preset influence the initial seeding of the organization in the fitness landscape. From this starting point, the organization then identifies a local peak within whose “basin of attraction” it has fallen.

With a preset degree of match of 1, only the first, most strategic, variable is set equal to the global optimum. As more variables are matched with their settings at the global optimum, fitness rises steadily according to Figure 4. However, the global optimum is not approached until nearly all policy variables are preset to their values at the global optimum. As a further test of the importance of identifying relatively strategic policies for strategy making, we also consider a random baseline in which the matrix of interactions is the same (hierarchical or centrality) but the policies that are correctly prespecified are randomly chosen. The gap between the curve depicting performance under such a random choice of policies to prespecify correctly and the other two curves (for hierarchy and centrality) indicate the power of presetting more strategic variables to their values at the global optimum. In contrast, the gap between the realized fitness level and the (normalized) value of 1, indicates the loss from not fully articulating the optimal policy array. This analysis implies, most broadly, that a priori strategy-making matters. The more policies that can be specified correctly a priori the higher the level of fitness the organization is able to obtain subsequent to its process of local search. However, there appears to be very little leverage to correctly prespecifying more or less strategic (influential or central) policies.

To make more sense of these patterns, it is useful to note that the fitness landscapes we are analyzing are quite complex, typically comprising over a dozen local peaks. In such worlds, the powers of local search are relatively limited. Local search rapidly leads to the identification of a local peak but conveys no assurance about the local peak’s global properties (i.e., its fitness value relative to the global optimum). Presetting the most strategic variables to their values at

¹² An important consideration, highlighted in Figure 5, is whether in this process of local search, the initial conscious choice of policies is adhered to and local search occurs only over the set of policies not prespecified or whether all policies are free to change in the process of local search. Figure 4 indicates the result when these initial policies are adhered to, while Figure 5 indicates the result of the lack of this strategic discipline.

the global optimum does lead to the identification of a better-than-average local peak (recall that the normalized fitness value would have a value of zero if the average realized fitness level equaled the average value of local peaks in the fitness landscape). However, a high level of specificity is necessary to obtain the highest possible fitness levels or configurations close to the global optimum: in rugged landscapes, there are just too many positive-gradient paths that lead to local peaks other than the global one.

Also note that while the articulation of and insistence on adherence to a single (or low-dimensional) strategic choice may not be sufficient to lead to the identification of a high-performing set of choices, a lack of such strategic discipline can lead to even less attractive results. Compare the top line in **Figure 5**, tracing the value of partially articulated activity maps in a hierarchical context in which preset choices cannot be varied (à la Figure 4) with the bottom line, which looks at a hierarchical context in which the preset policy choices *can* be revised in the process of local search. Thus, in both settings organizations share the same initial seeding, but in the unconstrained case the organization is free to vary all policy choices including those that had been preset to match their values at the global optimum. It turns out that with the degree of match of 1, the latter, “unconstrained” approach underperforms the “constrained” approach, and the gap between the two widens for intermediate degrees of match prior to convergence as the degree of match approaches the number of policy variables (i.e., 15). In other words, strategic discipline of the sort that constrains search processes can be quite important, even in a world that is “discovery-driven.”

The reason that respecting preset matches helps performance is that it places the organization in a more favorable basin of attraction than an a priori strategy with fewer matches would. As a result, strategic discipline in the sense of adherence to preset matches prevents the process of local search from leading the organization to a different basin of attraction that, on average, will yield an inferior local peak.¹³ The loss from a lack of discipline is greatest for a moderate degree of match in the initial specification of policy variables. With nearly all of the policies correctly specified, local search, even if unconstrained, is unlikely to lead the firm astray. Alternatively, with only one or two policies correctly specified, even if the firm adheres

¹³ Indeed, with only one policy preset correctly but without strategic discipline, this choice ends up being misspecified at the end of the search process 20% of the time under the hierarchy structure and 30% of the time under the centrality structure.

to these policy values, it is likely to find itself stuck at an only modest local peak consistent with those policies. At moderate levels of prespecification, the policies provide valuable guidance, but such a starting point does not preclude the possibility that unconstrained local search may lead to inferior local peaks inconsistent with the initial strategy choices. Indeed, the payoff to ex ante matching is concave in the context of strategic discipline, rising quite rapidly with greater levels of precision, and convex in the absence of such discipline, rising slowly with greater precision and then increasing rapidly as the policy is more completely specified. Without strategic discipline, the policy needs to be quite fully specified in order for the wisdom of the initial seeding not to be lost; with strategic discipline, in contrast, the beneficial effects of a good initial seeding are more strongly cumulative.

Strategic Mistakes and Tactical Mitigation

Success is not the only possible outcome to strategic prespecifications: they may also turn out to be mistakes. Alternatively, even if a policy choice made sense at one point in time, it may no longer be suited to an environment that has shifted and yet, if commitment-intensive, will be hard to reverse. The analysis in this subsection focuses on the downside rather than the upside of the effect of initial positioning in policy space. Specifically, it models the commitment-intensity or irreversibility of choices—perhaps their most basic temporal quality—by focusing on totally irreversible “mistakes” in the sense of policy variables whose values are preset to mismatch rather than match their values at the global optimum. The objective is to explore how the underlying structure of interactions among choices affects the residual costs of such mistakes, after local search aimed at tactical adjustment through both mitigation of these mistakes and efforts to align the full system of policy choices.

Figure 6 summarizes the normalized fitness level achievable when one of the 15 policy variables is preset to a value inconsistent with its value at the global optimum.¹⁴ Under a hierarchical pattern of interactions, fitness improves markedly as the preset mismatch shifts from one of the higher-order variables to lower-level policy choices. What is surprising is the modest improvement in fitness under the centrality interaction structure once one goes beyond the most central policy variables. Less central variables not only do not constrain, or substantially

¹⁴ It does not make sense to explore a random specification of the misspecified policy as the analysis explores the impact, exhaustively, of all policies being misspecified.

influence the payoff of many other choices, but they themselves are not greatly contingent upon other policy choices. Being contingent on other policy choices facilitates mitigating shifts in policy variables other than the one that is preset. As a result of the absence of such contingencies, the preset mismatch of lower-order policy choices is comparatively more damaging to fitness levels under the centrality structure.

Consistent with this effect on normalized performance of constraining a more peripheral policy variable, we see in **Figures 7** and **8** that firms operating under the constraints of history often identify a policy configuration that is distinct from the global optimum but internally consistent. In particular, **Figure 7** indicates the substantial proportion of configurations that constitute local peaks (i.e., are internally consistent) while **Figure 8** indicates that these configurations are not particularly close to the global optimum. Thus, while the firm could always choose the constrained optimum such that the hamming distance between the constrained optimum and the global optimum diverges by just one policy, the firm does not, as a second-best, specify a policy array that minimizes the distance between the constrained optimum and the global optimum.

The role of relatively peripheral policy variables in this regard bears repeating. To the extent that a focal policy that is misspecified is dependent on or influences other policies, compensating changes in these other policies can be made that facilitate a distinct, but nevertheless reasonably effective constellation of policies. In contrast, when a relatively peripheral policy is misspecified under the centrality structure, the constrained optimum tends not to correspond to a local peak (i.e., a consistent set of policy choices). Rather, the firm in some sense accepts this misspecification (the number of local peaks declines in **Figure 7**), but the hamming distance between the global peak and the constrained optimum shrinks (see **Figure 8**).

Modes of Interaction: Influence, Contingency and Autonomy

Figures 6, 7, and 8 taken together, suggest that the misspecification of a highly contingent policy does not impose the same performance costs as the misspecification of other variables. Indeed, there appears to be a certain robustness associated with contingent variables (see

Siggelkow(2002b) for a similar argument). The prior analysis of interaction structures in Figure 4 and elsewhere conflates the role of influence and contingency in that policies that are relatively less contingent are also less influential.¹⁵ In Figure 9, we consider an extreme adjacency matrix that disentangles these effects. We specify the first five policy variables such that with probability 1 they are influential and with probability 1 they are not contingent (i.e., $r = 1$ and $p = 1$). Similarly, we specify policies 6 to 10 as being with probability 1 contingent and with probability 0 influential (i.e., $r = 0$ and $p = 1$). The remaining five policies (policies 11 to 15) are treated as being autonomous (i.e., $p = 0$).

This stylized interaction allows us to tease out the underlying forces in the results we observe with the hierarchical and centrality interaction patterns. Figure 9 confirms that constraining one of the “influential” variables to differ from the global maximum has a profound effect on the relative fitness level of the constrained optimum. Somewhat more surprisingly, constraining the independent variables to differ from the global optimum has a larger impact than constraining the seemingly more important “contingent” variables. The reason for this is that the presence of contingency allows for the possibility of substituting or compensating changes in policy variables. While tightly linked interaction patterns have generally been viewed as fragile, they also allow, through equifinality, for a certain robustness. In contrast, when an autonomous variable is misspecified, that has no negative implications for other choice variables; at the same time, however, there is no opportunity to compensate for the misspecification.

The parsing out of effects in this stylized adjacency matrix also offers greater optimism for the power of high-level strategy making. Figure 10 tracks normalized fitness levels as an increasing number of variables are preset to match their values at the global maximum, with the remaining variables identified through a process of local search. The results suggest that it is sufficient to specify the purely influential variables correctly and then to follow up with a process of local search. The contingent variables are likely to be correctly specified if the influential variables are set to the global optimum, and the autonomous variables, as non-contextualized choices, can readily be set at their optimum value via a process of local search. In

¹⁵ Specifying the interaction structures solely by varying p^H and p^C would not eliminate such confounding effects. Variation in these parameters not only affects influence and contingency, but also the level of autonomy or interdependence. Thus, the analysis in this section is an important supplement to the prior analysis, but not a substitute.

that sense, at least, the intuition of the sufficiency of “grand strategy making” and the presumption that operating details can safely be left unspecified are validated. It is the intertwining of influence and contingency that prevents such top-level strategy making from proving sufficient in the case of the hierarchical and centrality structures.

5. Robustness

It is important to comment on the robustness of the analyses presented in the previous section. The characterizations are based on the averaged results from 100 independent runs for each of 100 distinct adjacency matrices. That is, an adjacency matrix is drawn given the specification of p and r and 100 distinct randomly seeded fitness landscapes are examined for each of these adjacency matrices. This process of examining 100 distinct fitness surfaces is replicated for each of the 100 adjacency matrices randomly drawn given a particular p and r values.

While this should allay concerns about robustness related to random effects, questions of robustness may remain with respect to the parameters of the model. In order to address these, it is useful to start out by noting that there are essentially 3 structural parameters: r , p , and N .¹⁶ The analysis above explored landscape structures in which either r or p was held fixed at 0.5 and the other parameter varied from 0 to 1 among the N policy variables. As part of a structural robustness analysis check, it is important to examine how these results might change if the conditioning parameter assumes values other than 0.5. These results are available from the authors in a Technical Appendix.¹⁷ The qualitative effects of the hierarchical position or centrality of a policy variable on the results of searching from a partially specified optimum or a constrained suboptimum turn out to be quite robust in general.

¹⁶ One could also explore alternatives to the uniform distribution for seeding the landscape; however, prior analysis (Kauffman, 1993; Rivkin 1997) indicate that the analysis of landscapes tends to be quite robust to the specification of alternative distributions for the random variate.

¹⁷ We have also examined the robustness of our results as N varies. Examining large values of N is extremely computationally intensive as the set of possible fitness values that need to be considered grows exponentially with N . We have re-analyzed our results for a range of N values and the qualitative effect of r and p (hierarchy and centrality) remain. The one qualitative effect of larger N that is evident is that it tends to reduce performance as the global peak increases relative to realized fitness. However, the comparative effect of hierarchy and centrality remains unchanged. At very low values of N , the results can take on very extreme values due to the limited number of peaks and the fact that the denominator of our normalized fitness level approaches zero as the value of the average local peak approaches the value of the global peak.

While there are some exceptions to this rule, these are generally artifacts of extreme values in the parameter space rather than the statistical fragility of the reported results. In particular, we see that for values of p near 0, the realized normalized fitness is quite low for all policy choices being constrained. The reason is that for very low values of p , there are few interactions and in turn few peaks in the landscape. The normalized fitness indicates the ratio of the difference between the realized fitness level and the average peak to the difference between the global peak and the average peak. As the fitness of the average peak approaches that of the global peak (since there are only a handful of peaks in the landscape), the denominator of this ratio approaches 0 and, as a result, any deviation from the global peak results in a low normalized fitness.

We also observe relatively low normalized fitness levels for extreme settings of r when one of the policy variables is constrained to be misspecified. As r approaches the extreme values of 0 or 1, the number of interactions shrink. In the limit, with r at 0 or 1 there is only a single peak in the landscape. Thus, for the reasons outlined above, in such settings our normalized fitness value tends to accentuate deviations from the global peak. What underlies both special cases is the point that our measure of performance becomes overly sensitive in extremely simple landscapes (even though the substantive results regarding centrality and hierarchy remain quite robust). This should not, on reflection, be too surprising. We specified this performance measure to highlight variations in performance in complex settings with a large number of local peaks among which we wished to distinguish rather than in simple (relatively non-rugged) settings.

Numerous other extensions to the basic analysis could be attempted. Given the simplicity of the temporal dimension of our modeling effort (which assumed the total irreversibility of specific variables), we did not discuss degrees of irreversibility, although differences in this regard might supply an additional useful marker of influence. Similarly, even though we did not allow the weights on (the direct effects of) choices to vary, it is clear that that is another key indicator of influence in the real world, either individually or in interaction with irreversibility (see Solow, et al., 2002 for an examination of this form of heterogeneity). Clearly, however, such extensions would constitute a distinct modeling exercise rather than an examination of the robustness of current specification.

6. Conclusion

Some choices condition other choices. This conditioning may be of a cross-sectional nature, as implied by the activity systems approach, or of longitudinal nature, as in models of path dependence and commitment. If the strategy field is to move beyond rhetorical appeals regarding the relative importance of one set of “linkages” or another, we must develop both more carefully specified theoretical models (cf., Miller, 2002), as well as engage in empirical work that is fine-grained enough to permit exploration of the nuances of choice structures (cf., Siggelkow, 2002a). The current analysis is clearly targeted primarily at the former goal.

One way of summarizing the results from this exercise is that it suggests a useful, if rough, way of partitioning choices: into autonomous choices, influential choices and contingent choices. Autonomous choices are choices that are disconnected from others. In relation to such choices—but not others—the notion of universal best practices makes some sense. Note that while getting these choices wrong does not, by definition, alter the payoffs from other choices, it is also true that these kinds of choices, if wrong, cannot be compensated for by contingent choices. Still, such choices can be made independently of an overarching choice of strategy and therefore have the quality of operational policies (Porter, 1996).

The same is sometimes true of groups of choices, as in the (nearly) decomposable systems originally highlighted by Simon (1962) and recently analyzed in the business context with an NK approach by Gavetti (2000). At the limit, a subsystem of choices that do not interact with any choices outside the subsystem can be treated like an individually autonomous choice: partitioned and made on standalone terms. The implied reduction in the complexity of the overall choice problem tends to be significant.

Choices that aren’t autonomous or decomposable, in contrast, should not be lumped together—as they are by the random interaction or canonical NK models—as having equal potential to be influential. As the mini-case on Vanguard suggested, it is important to recognize both the multiplicity of choices (or themes) and the fact that some of them matter more than others. The modeling effort that followed the mini-case focused, for the most part, on setting up two cross-sectional alternatives to the random interaction model of NK landscapes that encompassed variations in individual choice elements’ interactions with others: hierarchy and centrality. The initial analysis of strategy making revealed, under the assumption that choices are of symmetric weight but asymmetric in their interactions, that the identity of the policies that

are specified correctly is not as critical as the number of correctly specified policies and whether the organization is disciplined in adhering to these initial specifications in its subsequent search efforts. The subsequent analysis of the pure effects of influence, contingency, and autonomy reaffirmed the intuition of the importance of correctly prespecifying more strategic choices. In the initial analysis, as we varied the parameter r , we changed both the likelihood that a policy is influential and, conversely, the degree to which it is contingent. Thus, specifying a variable that is not very influential, is, by implication, specifying a variable that is highly contingent. As a result, the natural looking interaction structures such as depicted in Figures 3a and 3b have a more complex structure that may appear at first. It is necessary to parse out the separate effects of influence, contingency, and autonomy to fully reveal what constitutes more or less strategic choices and what the implications are of such choices being specified correctly. Hierarchical structures have, embedded in them, implications for contingency as well as influence. Similarly, centrality structures have implications for influence, as well as contingency and autonomy.

The simulation results on the constraints of history, or strategic mistakes, confirm that a choice's position within a hierarchical or centrality structure are, given the respective underlying choice structures, useful indicators of "strategic" choices that it is particularly important not to misspecify. This is where contingent choices—choices that are more influenced than influential—come into particularly sharp focus. The modeling effort indicated that such choices can afford two very distinct types of benefits: enabling the more effective pursuit of the direction determined by higher-order choices by aligning with them, and mitigating the effects of higher-order handicaps. In other words, contingent choices can be either advantage-seeking or disadvantage-mitigating, although the first role is the one that is typically stressed in the literature on strategy.

It is worth adding that the kind of constrained optima associated with disadvantage mitigation typically do not constitute an effort to minimize the (constrained) distance to the global peak. The policy configurations selected typically lie a considerable distance from the global peak but consist of an internally consistent set of choices. Consistency rather than proximity to the global peak appears to constitute a more powerful heuristic for improving performance. The exceptions arise when relatively peripheral policy variables are constrained. In such settings, sacrificing internal consistency is not too costly and the firm optimally gives up on it to some extent in order to get closer to the global peak. .

Even more broadly, the discussion in this paper should have highlighted the importance of thinking longitudinally as well as cross-sectionally in order to understand the interactions among choices. Neither the cross-sectional nor the longitudinal perspective seems likely to have a unique claim on truth. Nor is it the case that we can consider each perspective separately. Instead, it is their conjunction that seems to provide the richest insights about how effective competitive positions might be established.

References

- Bellman R. 1957. *Dynamic Programming*. Princeton University Press: Princeton NJ
- Burgelman, R. 1994. Fading Memories: A Process Theory of Strategic Business Exits. *Administrative Science Quarterly*, 39: 24-56.
- Caves R. 1980. Industrial Organization, Corporate Strategy and Structure. *Journal of Economic Literature* 18: 64-92
- Cyert RM, March JG. 1963. *A Behavioral Theory of the Firm*. Prentice-Hall: Englewood Cliffs NJ
- Dierickx I and Cool K. 1989. Asset Stock Accumulation and Sustainability of Competitive Advantage. *Management Science*, 35(12): 1504-1514
- Gavetti G. 2000. Cognition, Capabilities and Corporate Strategy Making. Unpublished working paper, January 2000 draft.
- Gavetti G, Levinthal D. 2000. Looking forward and looking backward: Cognitive and experiential search. *Administrative Science*. March 2000
- Ghemawat P. 1991. *Commitment*. Free Press: New York
1999. *Strategy and the Business Landscape*. Addison-Wesley: Reading
2002. Competition and Business Strategy in Historical Perspective. *Business History Review* 76 (spring 2002): 37-74.
- Kauffman S. 1989. Adaptation on rugged fitness landscapes. In D. Stein (Ed.), *Lectures in the Sciences of Complexity*. Addison-Wesley: Reading MA
1993. *The Origins of Order: Self-organization and Selection in Evolution*. Oxford University Press: New York
- Levinthal D. 1997. Adaptation on rugged landscapes. *Management Science* 43(7): 934-950
- Lewis A. 1985. On effectively computable realizations of choice functions. *Mathematical Social Sciences* 10: 43-80
- March J, Simon H. 1958. *Organizations*. John Wiley & Co.: New York
- Milgrom P, Roberts J. 1990. The Economics of Modern Manufacturing: Technology, Strategy, and Organization. *American Economic Review* 80(3): 511-28
1995. The Economics of Modern Manufacturing: Reply. *American Economic Review* 85(4): 997
- Miller, K. 2002. Knowledge inventories and managerial myopia. *Strategic Management Journal*, 23: 689-706.
- Mintzberg, H. 1978. Patterns in Strategy Formulation. *Management Science*, 24: 934-948.
- Nelson RR, Winter SG. 1982. *An Evolutionary Theory of Economic Change*. Harvard University Press: Cambridge MA
- Penrose E. 1959. *The Theory of the Growth of the Firm*. Basil Blackwell: Oxford

- Porter ME. 1996. What is Strategy? *Harvard Business Review*. November-December: 61-78
- Rivkin JW. 1997. *Consequences of Fit*. Unpublished Ph.D. dissertation, Harvard University.
2000. Imitation of Complex Strategies. *Management Science*, 46: 824-844.
2001. Reproduction of knowledge: Replication without imitation at moderate levels of complexity. *Organization Science*, 12: 274-293.
- Rumelt RP. 1984. Towards a Strategic Theory of the Firm. In RB Lamb, ed., *Competitive Strategic Management*. Prentice Hall: Englewood Cliffs, NJ
- Schelling TC. 1979. *The Strategy of Conflict*. Harvard University Press: Cambridge MA (First published in 1960)
- Selznick P. 1957. *Leadership in Administration*. Row Peterson: Evanston IL
- Siggelkow N. 1998. Benefits of Focus, Evolution of Fit, and Agency Issues in the Mutual Fund Industry. Unpublished Ph.D. dissertation, Harvard University
- 2002a. Evolution Towards Fit. *Administrative Science Quarterly*, 47: 125-159.
- 2002b. Misperceiving Interactions among Complements and Substitutes: Organizational Consequences. *Management Science*, 48: 900-917.
- Simon, H. 1962. The architecture of complexity. *Proceedings of the American Philosophical Society*, 106: 467-482.
- Stinchcombe A. 1965. Social structure and organizations. in J. March (Ed.), *Handbook of Organizations*. Rand McNally: Chicago IL
- Sussman HJ. 1975. Catastrophe Theory. *Synthese* 31(2) p. 4
- Teece D, Pisano G. 1994. The Dynamic Capabilities of Firms: An Introduction. *Industrial and Corporate Change* 3: 537-556
- Topkis DM. 1978. Minimizing a Submodular Function on a Lattice. *Operations Research* 26(2): 305-331
1995. Comparative Statics of the Firm. *Journal of Economic Theory* 67(2): 370
- Weinberger E. 1991. Local properties of Kauffman's N-K Model: A tunably rugged energy landscape. *Physical Review A* (44): 6399-6413
- Wernerfelt B. 1984. A Resource-based View of the Firm. *Strategic Management Journal* 5: 171-180
- Wright S. 1931. Evolution in Mendelian populations. *Genetics* 16: 97-159

Figure 1. Map of interactions among Vanguard's characteristic features in 1997

Source: Siggelkow, 1998

Figure 4. Value of Partially Articulated Activity Maps

Figure 5: Value of Constrained Local Search

Figure 6. Constraints of History **Figure 7. Constrained Optima:
Proportion that are Local Peaks**

**Figure 8. Constrained Optima:
Hamming Distance**

**Figure 9. Constraints of History:
Normalized Fitness Levels**

Figure 10. Valued of Partially Articulated Activity Maps