

Choice Structures 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

November 2000

Choice Structures and Business Strategy

There is a line among the fragments of the Greek poet Archilochus which says: “The fox knows many things, but the hedgehog knows one big thing.” Scholars have differed about the correct interpretation of these dark words, which may mean no more than that the fox, for all his cunning, is defeated by the hedgehog’s one defense. But, taken figuratively, the words can be made to yield a sense in which they mark one of the deepest differences which divide writers and thinkers, and, it may be, human beings in general. For there exists a great chasm between those, on one side, who relate everything to a single central vision, one system less or more coherent or articulate, in terms of which they understand, think and feel...and on the other side, those who pursue many ends...seizing upon the essence of a vast variety of experiences and objects for what they are in themselves.

—Isaiah Berlin, *The Hedgehog and the Fox*

1. Introduction

The fox and the hedgehog seem to us to capture an important part of one of the central tensions within the strategy field: the cross-sectional vs. longitudinal tug-of-war concerning the best perspective for thinking about linkages among choices. The cross-sectional perspective has increasingly come to emphasize the complexity of coordinating across a large number of choices when the cross-sectional linkages among them are significant (e.g., Porter, 1996). The longitudinal perspective actually subsumes a number of quite different views about how to model intertemporal linkages in choices but all of these views accord much more importance to path-dependence or historical constraint (e.g., Rumelt, 1984; Ghemawat, 1991; Teece and Pisano, 1994). In that sense, the longitudinal perspective has more of an affinity with the hedgehog’s perspective, with a commitment to acting in a particular way, whereas the cross-sectional perspective has more of an affinity with the fox’s.

This paper takes as its point of departure the presumption that both perspectives offer important insights and the strategy field would presumably benefit if we could synthesize them and, in particular, probe the interrelationships between temporal and cross-sectional linkages. The paper starts by exploring the dual role of cross-sectional and temporal linkages in the context of a reappraisal of case studies of the Vanguard mutual company and Southwest Airlines. These two cases have been highlighted as examples of cross-sectional linkages (Porter, 1996; Saloner, Shepard and Podolny, 2000), but we suggest that these cases also suggest that the

temporal ordering of strategic choices can play a critical role. More broadly, the case studies suggest that not all choices can be treated as being equally important: some are more “strategic” than others.

With these cases as backdrops, the paper sets forth a simple, 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 generalization of the framework of NK fitness landscapes introduced by Kauffman (1993). The paper examines not just the random interaction structure explored by Kauffman (1993) and others (Levinthal, 1997; Rivkin, 1997b), 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 commitment. 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.

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 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 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?

In addition to addressing these specific questions, the simulation exercises suggest a useful, if rough, way of partitioning choices, into autonomous choices that are disconnected from others, influential choices that have significant effects on others because of hierarchy, centrality or irreversibility, and contingent choices that are more influenced than influential.

Section 2 provides some historical background on the debate between the cross-sectional and longitudinal perspectives. Section 3 uses the case histories of Southwest Airlines and Vanguard to shed light on the debate, and provide some insights regarding the interrelationships between the two analytical perspectives. Section 4 considers various ways of integrating and analyzing the choice structures implied by the two perspectives, an exercise that leads us to propose a graph-theoretic generalization of the NK model of fitness landscapes. This generalized structure forms the basis for the analysis in section 5. Section 6 concludes.

2. Cross-Sectional and Longitudinal Linkages

The centrality of cross-sectional perspectives in business strategy is quite longstanding. Consider the following description of the Business Policy course offered in 1917 at the Harvard Business School:¹

An analysis of any business problem shows not only its relation to other problems in the same group, but also the intimate connection of groups. For example, not only is any problem of factory management related to other problems in the factory, and any problem of selling related to other problems in the sales department, but also the groups of problems are interdependent. Few problems in business are purely intra-departmental.

More generally, Business Policy courses came to focus on general management rather than functional management and to treat the interdepartmental or interfunctional coordination of choices as one of the most distinctive aspects of the general manager’s job.

Concern with longitudinal linkages has rather different antecedents, including Marshall’s (1920) distinction between short-run and long-run cost curves and Penrose’s (1959) interest in

¹ Harvard University, *Official Register, Graduate School of Business Administration* (March 1917), pp. 42-3. We are indebted to Jan Rivkin for this citation.

the constraints on firm growth. Perhaps the most direct bridge to the strategy literature was supplied by Selznick's (1957) work on the institutionalization process by which organizational practices become infused with value beyond their immediate technical role. Selznick studied inter-services competition in the U.S. military after World War II, when a debate raged about whether to unify the Army, Navy, Marines and Air Force into a single organization. Selznick noted that the Navy Department emerged as the defender of subtle institutional values. In essence,

Navy spokesmen attempted to distinguish between the Army as a 'manpower' organization and the Navy as a finely adjusted system of technical and engineering skills—a 'machine-centered' organization. Faced with what it perceived as a mortal threat, the Navy became highly self-conscious about its distinctive competence.²

Largely on the basis of this case study, Selznick ventured that "commitments to ways of acting and responding are built into the organization."³

Selznick's ideas, particularly his concept of distinctive competence, were quickly seized on by business policy scholars. As an illustration, Andrews' (1971) textbook, which dominated its market segment for much of the 1960s and 1970s, argued that the strategist had to decide which aspects of the firm would be "enduring and unchanging over relatively long periods of time" and directly incorporated Selznick's concept of distinctive competence, and the longitudinal or intertemporal linkages underlying it, into its core analytical framework. As a result, a firm's strengths and weaknesses were supposed to be fixed in the operational short run (Caves, 1980). Cross-sectional linkages were less evident in the Andrews framework, but reportedly continued to account for a significant fraction of the classroom discussion of individual cases.

This classical synthesis or, perhaps more accurately, even-handedness, seemed to unravel in the 1980s and the 1990s. In particular, there were a number of attempts to elevate the attention paid to the constraints implied by longitudinal linkages as opposed to cross-sectional ones. Within the strategy literature, the development of the resource-based view of the firm (Wernerfelt, 1984; Rumelt, 1984) highlighted the short-run fixity of factor inputs. This interest in the imperfections of factor markets has led to considerable discussion of the development of

² Philip Selznick, *Leadership in Administration* (Evanston, IL: Row Peterson, 1957), pp. 49-50.

firm capabilities (Dierickx and Cool, 1989; Teece and Pisano, 1994). Finally, in contrast to the “smoothly distributed” choice structures emphasized by the literature on capabilities, contributions rooted in mainstream microeconomics, ranging from Schelling (1960) to Ghemawat (1991) emphasized longitudinal linkages derived from discrete choice “commitments” that are, at least to a degree, irreversible. Despite some obvious differences in terms of how these approaches conceived of longitudinal linkages, their common characteristic of path-dependence ensured a focus on historical constraint or prior conditioning of current initiatives.

Perhaps inevitably, the last few years have seen a cross-sectional backlash to this heightened emphasis on longitudinal linkages. In particular, Porter’s recent work on activity systems (Porter, 1996) has revived the early emphasis on interdepartmental and interfunctional coordination while spurring analytical advances in the cross-sectional view of strategy. We formalize this argument by allowing for asymmetries among choices (a generalization of the canonical NK model) and, perhaps more importantly, connect this cross-sectional perspective with important ideas of temporal linkages that the strategy literature has previously identified. In linking these two types of linkages among choices, it is useful to begin by considering the historical evolution of activity systems in two of the salient case examples of such systems—Southwest Airlines and the Vanguard mutual fund group.

3. Two Case Studies

Given the complexity of simultaneously examining both cross-sectional linkages as well as longitudinal ones, it is useful to ground our discussion in detailed case histories. The two that we will consider in this section have both been cited as strong examples of the importance of the cross-sectional perspective on linkages among choices. The cases concern two above-average performers in their respective industries (in the U.S.): Southwest Airlines in the airline industry and the Vanguard Group in the mutual fund industry. Our discussion of Southwest is informed primarily by the treatment in Porter (1996) plus some additional secondary research. Our discussion of Vanguard is informed almost entirely by an in-depth case study of the company by

³ *Ibid*, p. 47.

Siggelkow (1998).⁴ To preview our conclusions, we read the two cases as suggesting that it is important to recognize both the multiplicity of choices (in terms of the epigraph to this paper, the fox's perspective) and the fact that some of them matter more than others (the hedgehog's perspective) for reasons that include temporal ordering.

Southwest Airlines

Southwest Airlines is the only U.S. airline to have been consistently profitable in the last 25 years, has grown revenues at 20 to 30 percent annually over the last five years, maintains the youngest fleet among major carriers as well as the lowest debt levels, and leads the industry in terms of customer service ratings. This example has been deconstructed in several different ways.⁵ Here we focus on Porter's characterization of Southwest's activity system, which is reproduced in **Figure 1**.⁶

Porter explains that many of the choices embedded in Southwest's activity system are exceptions to rather than normal practices within the airline industry. He also provides an indication of the strength of the cross-sectional linkages between them by describing Continental's unsuccessful attempt to imitate Southwest on a number of point-to-point routes by setting up Continental Lite, a carrier that eliminated first-class service and meals, tried to shorten turnaround times at gates, increased departure frequency and lowered fares. Continental Lite continued, however, to provide baggage checking, seat assignments and frequent flyer awards, as well as to use travel agents, since Continental remained a full-service airline on other routes. This hybrid business system quickly proved to be unviable. Delays due to congestion at hub cities and baggage transfers caused numerous late flights and cancellations. Customers were also irked by the fact that, unable to offer the same frequent-flier benefits on the much lower Lite fares, Continental reduced the awards on its entire frequent-flier program. Similarly, Continental could not afford to pay standard travel agent commissions on Lite fares but could not afford to

⁴ Siggelkow, Nicolaj. 1998. "Benefits of Focus, Evolution of Fit, and Agency Issues in the Mutual Fund Industry," unpublished Ph.D. dissertation, Harvard University (1998).

⁵ Also see, for example, the discussion of Southwest in James L. Heskett, W. Earl Sasser, Jr., and Leonard A. Schlesinger in *The Service Profit Chain* (New York: Free Press, 1997). Like many other writers focused purely on execution, however, Heskett *et al.* fail to address the theory of barriers to imitation—or any other path-dependent linkages in the firm's profit function over time.

⁶ One of the circles in Porter's original activity system, "high compensation of employees," has been dropped because other sources indicated that compensation at Southwest is not significantly above average. Another circle,

do without travel agents for its full-service business, so it compromised by cutting commissions across the board. The new operation accumulated hundreds of millions of dollars of losses and had to be grounded.

Based on these and other examples, Porter argues that the interplay of complementarities *and* trade-offs across multiple activities is critical to the possibility of “many best ways to compete.” He emphasizes that firms ought to build tightly-coupled activity systems that recognize the cross-sectional linkages among choices.

Reading Southwest as a pure example of the importance of cross-sectional linkages is, however, just one possible way of deconstructing its history. Another way of doing so is suggested by the dark circles in **Figure 1**, which denote what Porter characterizes as “higher-order strategic themes” at Southwest. This dichotomization would seem to suggest that the dark circles have a contextuating effect on the light circles. The dark circle of “high aircraft utilization” is particularly potent: all of the other dark circles and most of the light circles could be connected directly to it if one were inclined to redraw **Figure 1** along such lines. Consider.

Southwest manages to fly its planes for an average of 11.5 hours per day, compared to 8.6 for the industry (even though one might expect lower numbers for Southwest given that its flights tend to be relatively short and turnaround times on the ground relatively fixed).⁷ Without this resource-utilization advantage, Southwest would need up to one-third more planes to fly the same number of trips! High aircraft utilization is directly connected to the other dark circles in **Figure 1**: lean and highly productive ground and gate crews, frequent and reliable departures, limited passenger amenities, short-haul point-to-point service between mid-sized cities and secondary airports, and very low ticket prices. It seems to condition the majority of the light circles in **Figure 1** as well: standardization of the fleet around Boeing 737s, no meals, baggage transfers, seat assignments or connections with other airlines, and an emphasis throughout the organization on minimizing turnaround times at airport gates. The imperative of high aircraft utilization also seems to explain a number of policies not mentioned in **Figure 1**: employee compensation based on trips rather than hours whenever possible, and the dedication of an operations agent—a “case manager”—to the turnaround of each flight (unlike competitors who

“limited passenger service,” has been modified to “limited passenger amenities” because Southwest does provide quality service along the dimension of convenience, as Porter himself notes.

⁷ Freiberg & Freiberg, *Nuts!: Southwest Airlines’ Crazy Recipe for Business and Personal Success*, p. 51.

typically assign an operations agent to handle 10 to 15 flights at a time).⁸ All this could be read to suggest that an overarching choice of configuration set the context for most of the other exceptional rather than normal choices embedded in Southwest's activity system.

Since Porter provides only a snapshot of Southwest's activity system, it is impossible to distinguish between these cross-sectional vs. longitudinal readings of the Southwest example on that basis. More discriminatory power is afforded by Siggelkow's (1998) longitudinal case study of the Vanguard Group of mutual funds, to which we turn next.

Vanguard Group

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. Like Southwest, Vanguard has been cited as an example of the power of cross-sectional linkages among choices (Porter, 1996). For expository purposes, it is useful to start with Siggelkow's characterization of Vanguard's activity system in 1997 (see **Figure 2**). To begin by pointing out the obvious, **Figure 2** is rather complex. Thirty-seven light circles and six dark ones "fit" with each other (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).

This complexity can be, and has been, interpreted as an instantiation of the importance of cross-sectional linkages. But a longitudinal look 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). A 1973 map of Wellington 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

⁸ Jody H. Gittel, "Coordinating Service across Functional Boundaries: The Departure Process at Southwest

incorporated as a mutual holding company in which the shareholders of the underlying funds would own the managing fund complex.⁹ After this unusual—and still unique¹⁰—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 3**).¹¹

The reasons why the choice of organizational form in 1974 mattered so much are relatively 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.

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

Airlines,” Harvard Business School Working Paper no. 98-050.

⁹ 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.

¹⁰ 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.

¹¹ Here, our characterization differs from Siggelkow’s: “Vanguard slowly adopted a very consistent set of choices.”

autonomy from its primary investment management advisor, Wellington Management Company.¹² Vanguard could now both bargain more effectively with its advisors and, ultimately, in 1981, bring some advisory activities in-house.¹³ 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 2**). This outcome does not, however, appear to have been the product of either *ex ante* design or an emergent process of discovery. Rather, Vanguard was founded on the basis of a radical choice of organizational structure—a mutual form of organization—which apparently directed management to subsequent choices that “fleshed out” the activity map that current cross-sectional analysis reveals.

The Vanguard example, or at least this reading of it, is probably somewhat extreme: one can easily imagine situations in which there *isn't* one key choice that influences all the rest. Having acknowledged as much, we should add that both the Southwest and Vanguard examples can be read more generally as calling attention 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 of such choice asymmetries, first in a purely cross-sectional context and second in a context that adds in explicitly longitudinal elements. But first, we must assess various ways of analyzing interactions among choices in both kinds of contexts.

4. Modeling of Linkages

Grappling with interdependent 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

¹² Investment management companies typically control the distribution of the mutual funds for which they provide investment advice, making it quite difficult to shift advisors.

¹³ Vanguard brought in-house the relatively “plain vanilla” advisory activity of management fixed-income funds and index funds, relying on outside advisors for the management of actively managed equity funds.

in a “reasonable” period of time (i.e., a polynomial function of time) (Lewis, 1985; Rivkin, 1997a). And from a longitudinal or dynamic perspective, such situations generally do not lend themselves to “pushing forward” in time on the basis of a smaller set of measures (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 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.)

The other response to the problem of multiple, linked choices has been to build on the NK-simulation approach pioneered by Kauffman (1993) in evolutionary biology (cf., Levinthal, 1997 and Rivkin, 1997b). 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 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 **Figures 1** and **2**. 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 systems in **Figures 1** and **2** bear more than a passing resemblance to mathematical graphs. A mathematical graph can, of course, 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 0s and 1s in column j , with a value of 1 indicating that the payoff to the variable in the row being considered is contingent on variable j , and a value of 0 denoting independence. Similarly, reading across row i in such a matrix indicates the variables the payoff to choice variable i is itself contingent upon.¹⁴ The principal diagonal of an adjacency matrix always consists of 1s, 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 be more important than others. Of the many possible adjacency matrices that embody such asymmetries, we focus on simulating the effects of two choice structures that are suggested by the case studies presented earlier.

Reconsider, first of all, the plausible conclusion that Vanguard was founded on the basis of a radical choice of organizational structure—a mutual form of organization—which apparently directed management to subsequent choices that "fleshed out" the activity map that current cross-sectional analysis reveals. One might think of the corresponding adjacency matrix

¹⁴ In addition to such direct effects on value contributions, variables may, of course, be indirectly related through other variables.

as containing 1s in one column (corresponding to the choice of organizing as a “true” mutual) as well as in the principal diagonal, with 0s 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 which higher-order policy choices influence lower-order choices but the reverse is not true. In graph-theoretic terms, hierarchies are best thought of as directed (or at least rooted) trees, with the action taking place (i.e., the 1s cropping up) to the left of the principal diagonal. The particular form of hierarchy we explore in this paper has 1s as all the entries to the left of the principal diagonal (see **Figure 4a**). Choice 1 is hierarchically the most important, choice 2 the second most important, and so on, which lets us take a finer-grained look at the effects of variations in the degree of hierarchical importance than a star structure would permit.

And then there is the case of Southwest, less conclusive because we have only a snapshot of Southwest’s activity system, not a longitudinal view. Consider the more general problem that given informational barriers, all that we might be able to do is observe linkages between choices, not the direction of influence, i.e., that in observational terms, we might have to work with undirected graphs. Then a systematic procedure for ranking choices in terms of their importance would be to array them in terms of the number of other choices with which they are linked. For strictly illustrative purposes, application of this procedure to **Figure 1**’s map of Southwest might imply that “lean, highly productive ground and gate crews” and “limited passenger amenities” are the most central choices and the small circles on the periphery on the left hand side of the figure the most peripheral.¹⁵ More generally, we can think of ranking choices in terms of their importance with a centrality measure that tallies up the number of other choices (up to $N - 1$) with which they are linked. Since directionality cannot—by assumption—be observed, there is symmetry in a specific sense: the 1s to the left of the principal diagonal are mirrored by 1s to its right. Whether the 1s cluster centrally in the adjacency matrix, however, depends on the order in which choice variables are labeled. The particular form of centrality we explore in this paper embodies a structure and a labeling scheme that has 1s as all the entries to the left of the inferior diagonal (but distributed symmetrically to the left and the right of the principal diagonal)—see

¹⁵ Note however, that while labor and customers get their (limited) due, capital efficiency, in the form of “high aircraft utilization” seems underemphasized given our earlier conclusion that all of the other dark circles and most of the light circles could be connected directly to it if one were inclined to redraw **Figure 1** along such lines. Evidently, the results from the procedure of counting the number of links are only going to be as appealing as the activity map with which one starts out.

Figure 4b, which transposes **Figure 4a** along these lines. Thus, choice 1, with links to 9 other choices, is most central, choice 2 second most central, and so on.

We benchmark what happens in these hierarchical and central structures against the symmetric but random interactivity built into the canonical NK structure. In adjacency matrix terms, what the canonical NK structure embodies is a symmetric matrix in which, apart from the elements of the principal diagonal, there are K 1s in each row and each column, but they are randomly distributed within these constraints. Note, in particular, that with K equal to $N-1$, the graph of choices is completely connected and the adjacency matrix is full of 1s. In the simulations that follow, we assume for this random benchmark that K equals 6, because that generates roughly the same number of local peaks as our other two structures.

For all three 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 0s and 1s assigned, respectively, to denoting 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 1.¹⁶ 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) and the luck of the draw. Specifically, it is common to 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 $\sum_{i=1 \text{ to } N} f(x_i | x_{i1}, x_{i2}, \dots, x_{iK_i}) / N$. We follow this procedure as well.

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

¹⁶ 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).

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. Again, we retain this prior assumption even though we intend to focus on asymmetries among choices. Putting different weights on the direct contributions of choice elements does not seem to us to be the best way of gaining insight into the indirect contributions that choice elements can make to overall performance by virtue of the linkages among them. (And again, there are also considerations of maintaining comparability with prior work.)

We also assume that N equals 10—a level of multidimensionality that is sufficient to generate more than 1 million distinct graphs. The results that we report are averaged over 1,000 independent landscapes that share the same structure. That is, the landscapes are hierarchical, central, or random with a particular adjacency matrix structure that is held fixed, but a distinct seeding (drawn from the uniform $[0,1]$ distribution) is carried out for the fitness values of the policy variables, the $f(x_i | x_{i1}, x_{i2}, \dots, x_{iK_i})$ terms, for each of the 1,000 runs. The repetition is meant to allow for the averaging out of purely random effects.

Finally, we should add the caveat that while the analysis highlights the effects of linkages among the 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.

5. Simulation Results

We apply simulation techniques to the choice structures described above to answer several different types of questions. First, what are the effects of presetting a certain number of policy choices equal to their values at the global optimum with the remaining choices determined by a process of local search? If a few higher-order choices make subsequent lower-level choices

self-evident, the possibility of achieving a coherent configuration of policy choices through grand strategizing at a high level appears much more promising. Second, what happens when one of the N values of the policy variable is preset to a value inconsistent with the value of that variable for the global peak? This is a simple way of incorporating historical constraints on firms' movement in policy space into the analysis. Under both kinds of presets, we look at differences in choices as well as fitness levels: how does the hamming distance—the number of policy choices that differ across two policy vectors—between a preset (i.e., constrained) optimum and the global optimum vary across the three choice structures, and are the differences largely differences in low-level policy choices or are they invariant to a choice variable's hierarchy or centrality? The answers to the last question also lead us to consider, in the next section, a different sort of choice structure that highlights the differences among policy choices in terms of patterns of influence, contingency and autonomy.

Our results concerning fitness levels are normalized 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 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, due to the seeding of fitness levels by drawing on 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. The landscapes examined here typically have 40 to 50 local peaks. The average value of these local peaks is one benchmark of the level of fitness that cognitively constrained choices might generate.¹⁷

To normalize both with respect to 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

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

¹⁷ Actually, local search would generate a somewhat higher level of fitness than this. Higher peaks have broader basin of attraction. Thus, if you started a single organization on each point in the landscape, more organizations would end up at higher local peaks than lower local peaks.

where Avg is the average value of the local peaks, and Global is the average value of the global peak across the 1,000 landscapes. These normalized fitness levels are what we actually report in the subsections that follow.

Strategic Choices and Local Search

Our exploration of preset matches 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 **Figures 1 and 2**), 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 **Figures 1 and 2**, followed by a process of local search), then the requirements for strategic planning remain relatively modest.

Figure 5 addresses this issue in the following manner. A certain number of policy choices (“degree of match”), selected in decreasing order of “strategic” importance (in the hierarchical and centrality structures), are set to equal their value at the global optimum, and the remaining policy choices are then determined 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. 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 preset partially seed the organization in the fitness landscape, with the remaining non-preset variables randomly assigned. From this starting point, the organization then identifies the local peak within whose “basin of attraction” it has fallen.

With a 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 5**. However, the global optimum is not approached

until nearly all policy variables are specified to equal their settings at the global optimum. Similarly, hamming distances—the number of policy choices that differ between two policy vectors—tend to be quite large (see **Figure 6**).¹⁸

The gap between the curve depicting performance under the random network structure and the other two curves indicates 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 value of 1, indicates the loss from not fully articulating the optimal policy array. 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 40 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 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 7**, tracing the value of partially articulated activity maps in a hierarchical context in which preset choices cannot be varied (à la **Figure 5**) 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 consciously specified as well as those initially set at random. 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

¹⁸ Note that a randomly assigned string of length 10 would, in expectation, have a hamming distance of 5 from the global optimum.

convergence as the degree of match hits 10. Similarly, the unconstrained approach fails to generate smaller hamming distances than the constrained approach. In other words, strategic discipline *can* be useful, even in a world that is “discovery-driven” because of cognitive constraints as well as true uncertainty.

Constraints of History

We now turn our attention to the constraints imposed by history, which we model by presetting variables one at a time to values inconsistent with their values at the global optimum.

Figure 8 summarizes the normalized fitness level achievable when one of the 10 policy variables is preset to a value inconsistent with its value at the global optimum. Under a hierarchical pattern of interactions, fitness improves markedly as the preset mismatch shifts from one of the higher-order variables to a lower-level policy choice. Not surprisingly, under the random interaction pattern, fitness is independent of which policy choice is mismatched. What is surprising is that the improvement in fitness is relatively modest under the centrality interaction structure once one goes beyond the most central policy variables. Less central variables not only do not constrain, or substantially 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 compensatory shifts in policy variables other than the one that is preset. As a result of the absence of such contingencies, the presetting of lower-order policy choices is more damaging to fitness levels under the centrality structure.

Another striking feature of this set of simulations concerns how few of the optima with preset mismatches constitute local peaks of the fitness landscape. Given the importance of configurational effects, one might reasonably conjecture that constraining one variable to differ from the global optimum would lead to the selection of a different, non-global, peak in the fitness landscape. However, **Figure 9** indicates that this is relatively uncommon except as one turns to presetting the least important variables under the hierarchy and centrality structures.¹⁹

Additionally, the optima with preset mismatches tend to be quite distant from the global optimum. **Figure 10** summarizes the hamming distance between the global optimum and the constrained optima as different policy variables are preset, in mismatched ways, under the three

¹⁹ We do find, however, that the optimum with a preset mismatch would be a local peak about 90% of the time if the preset constraint were relaxed with the other N-1 variables held fixed.

interaction structures. Since N equals 10, a random choice of policy values would, in expectation, result in a hamming distance of 5 from the global optimum. **Figure 10** reveals that the constrained optima do not differ much from this random baseline except when the least strategic variables are preset to mismatch. In other words, getting even one moderately important choice wrong can lead to significant distancing from the global optimum.

Linkages among Choices: Influence, Contingency and Autonomy

While the prior analysis has indicated in an aggregate manner the extent to which strategy choices will fall short of the optimum under various settings, it does not indicate which variables are being misspecified. Alternatively put, while we know that the hamming distance from the global optimum is apt to be substantial, is that distance largely driven by “secondary” policy choices or is it relatively invariant to a variable’s level in terms of hierarchy or centrality?

Figure 11 focuses on preset mismatches (similar results obtain for preset matches) and indicate that there is more of a tendency to set higher order variables to their values at the global optimum.²⁰ However, particularly under the centrality structure, we also observe a somewhat greater tendency for the least “strategic” variables to match their values at the global optimum. Under the centrality structure, lower-level policy variables have the weakest linkages with central policy variables. As a result, lower level variables may be set to match their values at the global optimum even if a more central variable is condemned to a mismatch.

The broader suggestion is that the “natural” adjacency matrices we have looked at so far mix up at least three very different types of effects: influence, contingency and autonomy. Variables may be more or less influential to the extent that they affect the payoffs to other variables. In an adjacency matrix, this is represented by the prevalence of 1s in the relevant column. Independent of influence, the payoffs from specific variables may be more or less contingent on other choices, as reflected in the number of 1s in the relevant row of the adjacency matrix. Finally autonomy is characterized by variables that are neither influential nor contingent: variables that correspond, in graph-theoretic terms, to disconnected vertices. In this subsection, we look at a choice structure—distinct from the three that we have already examined—that separates out these three effects.

²⁰ Note that given the policy choice is represented by a binary variable, a random baseline yields an expected value of 0.5 in these figures.

Consider the stylized interaction pattern depicted in **Figure 12**.²¹ The first three variables are purely influential. That is, the direct payoff to these variables is only dependent on their own setting, but other variables are contingent on their value. Variables 4 to 6 comprise these “contingent” variables. While contingent on the value of variables 1 to 3, they themselves do not influence the payoff of other variables. Finally, variables 7 to 9 are independent of all other variables. Their setting influences the overall fitness level, but their value neither influences the payoff of other variables nor is their fitness contribution contingent on other variables.

This stylized interaction allows us to tease out the underlying forces in the results we observe with the hierarchical and centrality interaction patterns. **Figure 13** indicates 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.²² 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, the equifinality that high levels of interaction engender also allows for a certain robustness. In contrast, when an autonomous variable is misspecified, that doesn’t create negative ramifications elsewhere in the system of policy choices; 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 some greater optimism for the power of high-level strategy making. **Figure 14** 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. Our results tend to suggest that it is sufficient to specify the influential variables correctly: given such a specification, a process of local search is sufficient to yield the global optimum. The contingent variables will “automatically” 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

²¹ It is convenient to reduce N from 10 to 9 in this context in order that there are an equal number of each of the three types of variables.

²² Note that due to the limited number of interactions in this interaction matrix (and secondarily due to the reduction in N from 10 to 9), there are relatively few peaks in this landscape: on the order of 4 or so, in contrast to 40-plus in

at their optimum value. In 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 to prove sufficient in the case of the hierarchical and centrality structure.

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, as well as engage in empirical work that is fine-grained enough to permit exploration of the nuances of choice structures. The analysis above was 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. They should therefore be separated out in practical terms and probably excluded from the domain of strategy in academic terms.

The same is sometimes true of groups of choices, as in the (nearly) decomposable systems originally highlighted by Simon 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.

the other matrixes. As a result, the magnitude of the fitness differences across local peaks is not enormous, but our normalization structure does yield a large spread.

Choices that aren't autonomous or decomposable, in contrast, are lumped together under the random interaction or canonical NK model as having equal potential to be influential. The mini-cases on Southwest and Vanguard suggested, however, that it is important to recognize both the multiplicity of choices (or themes) and the fact that some of them matter more than others. The operational question is, as always, how the influential or "strategic" choices should be identified.

The modeling effort that followed the mini-cases focused, for the most part, on setting up two cross-sectional alternatives to the random interaction model: hierarchy and centrality. The simulation results confirmed that hierarchy and centrality are, given the respective underlying choice structures, useful indicators of "influential" choices that it is particularly important to get right. 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. Lastly, even though we did not think it analytically interesting to allow the weights on (the direct effects of) choices to vary, it is clear that that is another aspect of the differences among choices observed in the real world, and often an important one. So weight might be yet another indicator of influence, either individually or in interaction with irreversibility.

While it is particularly important to get influential choices right, it is also worth considering the case where they hurt rather than help—most plausibly, perhaps, because of the dead hand of the past. This is where contingent choices—choices that are more influenced than influential—come into the picture. 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 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 local peaks in the fitness landscape. Rather, they comprise, what from a purely cross-sectional perspective, appear to be inconsistent activity systems. These apparently inconsistent patterns of choices constitute a "second-best" compromise among the conflicting demands of cross-sectional interdependence. A surprising finding, fleshed out in our final analysis, is that cross-sectional interdependence of a particular form—when choices are

contingent upon other choices but don't themselves influences other choices—provides a certain robustness to the choice problem.

Such robustness notwithstanding, it remains true that high levels of specificity or (presumably) more sophisticated adaptation strategies are 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. In other words, selection matters, not just *ex ante* design.

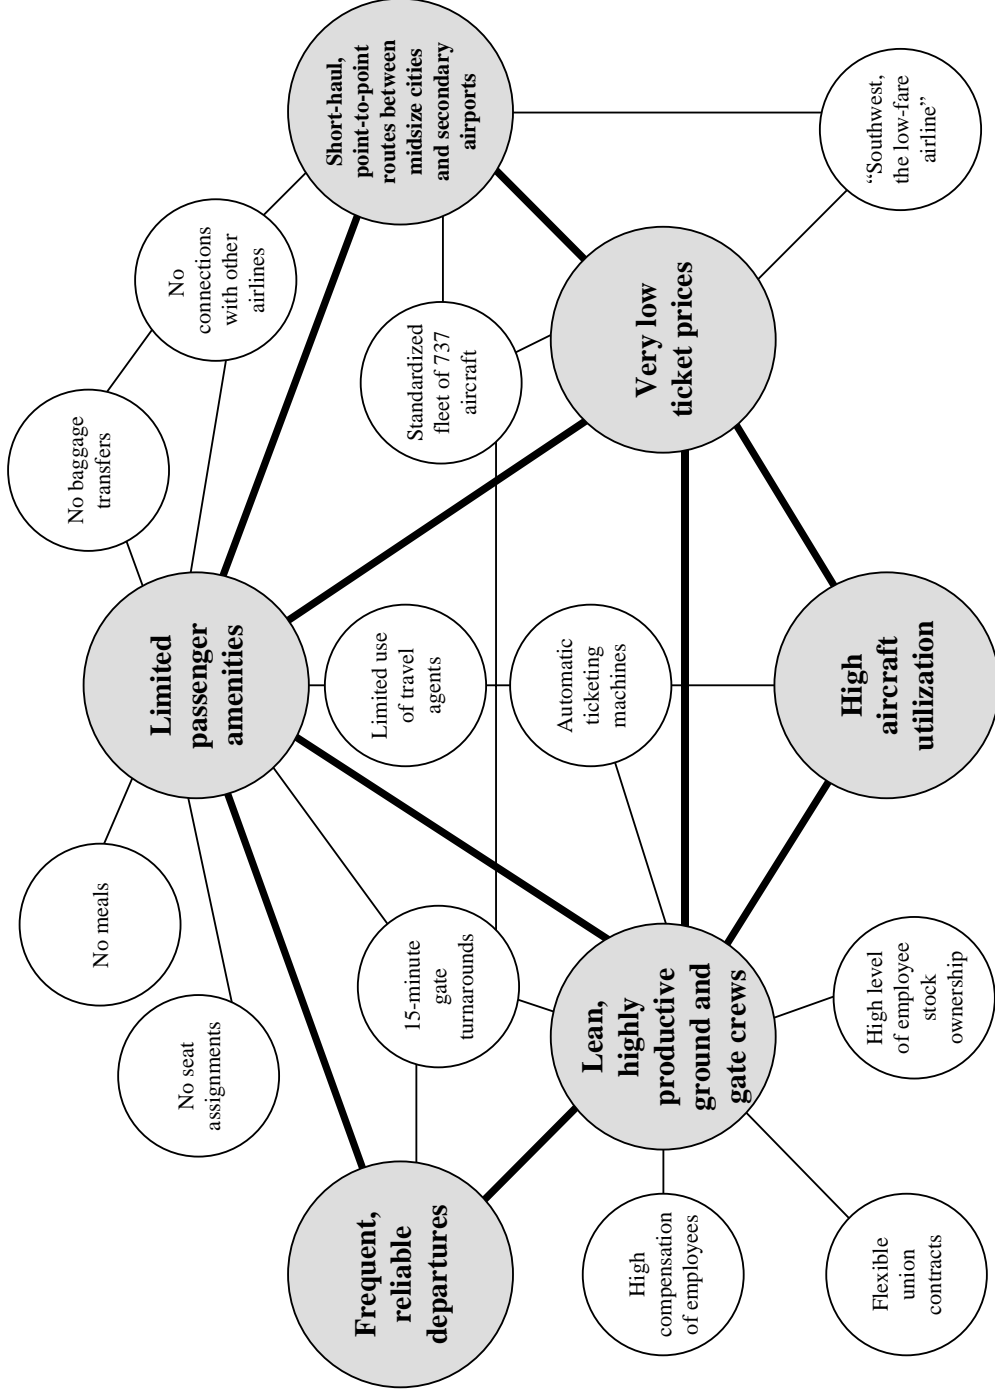
Finally, to refer back to the epigraph with which this paper began, this paper can be read as a first attempt to meld the two very different perspectives on choice structures and strategy that were exemplified by the fox and the hedgehog. Interesting opportunities for further work include looking at the properties of other choice structures, undertaking additional empirical investigation—preferably longitudinal as well as cross-sectional—of actual choice structures and patterns, examining temporal linkages in more detail, and allowing for interactions across competitors. Still, we hope that the analysis in this paper has made a start at suggesting how to shift the discourse within the strategy field from a dialectical perspective on the fox vs. the hedgehog to a more synthetic one.

References

- Andrews K. 1971. *The Concept of Corporate Strategy*. Richard D. Irwin: Homewood IL
- Bellman R. 1957. *Dynamic Programming*. Princeton University Press: Princeton NJ
- 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
- Freiberg K, Freiberg J. 1996. *Nuts!: Southwest Airlines' Crazy Recipe for Business and Personal Success*. Bard Press: Austin
- 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
1998. Competition and Business Strategy in Historical Perspective. Harvard Business School Press: Boston MA. HBSP #798-010
1999. *Strategy and the Business Landscape*. Addison-Wesley: Reading
- Gittell JH. 1998. Coordinating Service across Functional Boundaries: The Departure Process at Southwest Airlines. Harvard Business School Working Paper no. 98-050
- Heskett JL, Sasser WE, Schlesinger LA. 1997. *The Service Profit Chain*. Free Press: New York
- 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
- Learned EP, Christensen CR, and Andrews K. 1961. *Problems of General Management* Richard D. Irwin: Homewood IL
- 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
- Marshall A. 1920. *Principles of Economics* (8th edition). Macmillan: London
- 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
- 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 1996
- Rivkin JW. 1997a. Consequences of Fit. Unpublished Ph.D. dissertation, Harvard University
- 1997b. Imitation of Complex Strategies. Working paper, Harvard Business School.
- Rumelt RP. 1984. Towards a Strategic Theory of the Firm. In RB Lamb, ed., *Competitive Strategic Management*. Prentice Hall: Englewood Cliffs, NJ
- Saloner, G, Shepard A, Podolny, J. 2000. *Strategic Management*. John Wiley: New York
- 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
- 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. Southwest Airlines' Activity System



Source: Michael E. Porter "What is Strategy" *Harvard Business Review*, Nov-Dec 1966

Figure 5. Value of Partially Articulated Activity Maps




Figure 6. Partially Specified Activity Maps and Proximity to Global Optimum


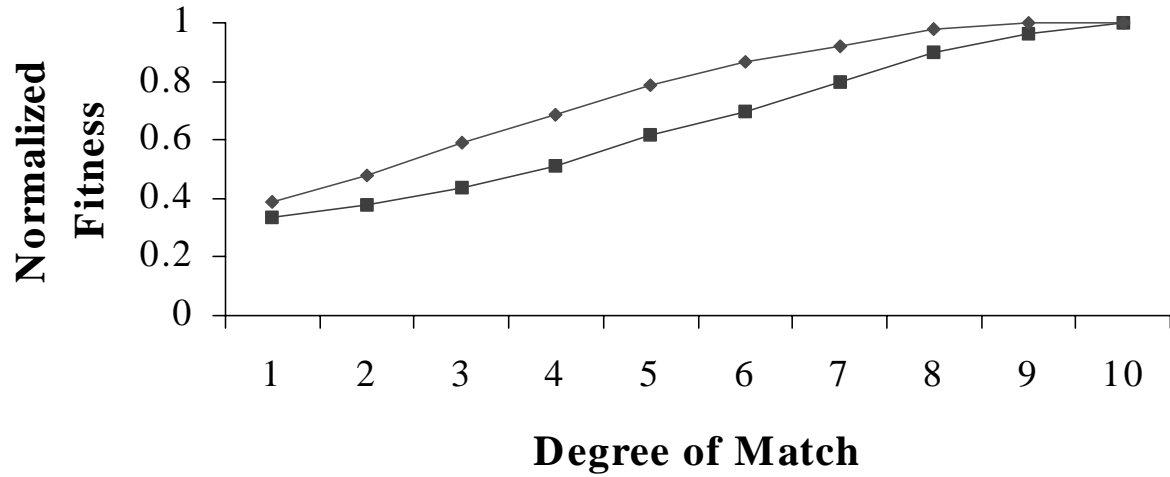
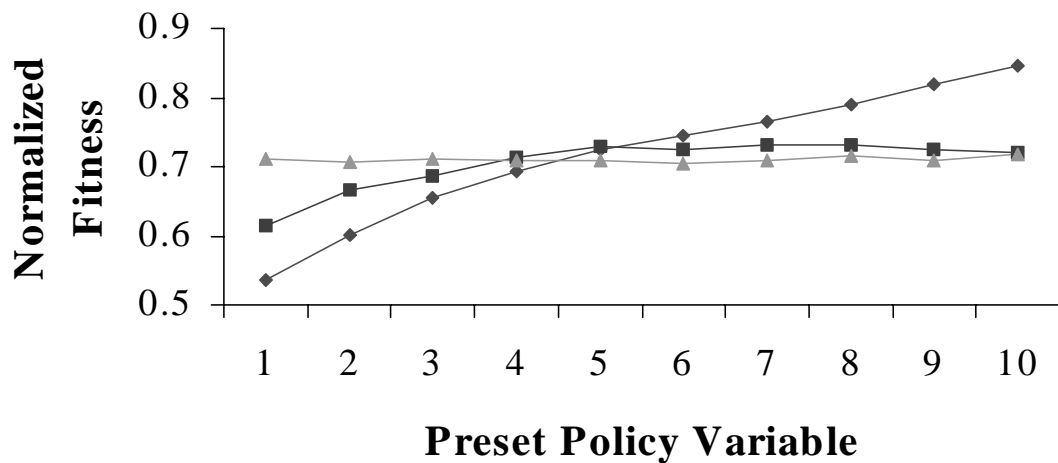


Figure 7: Value of Constrained Local Search



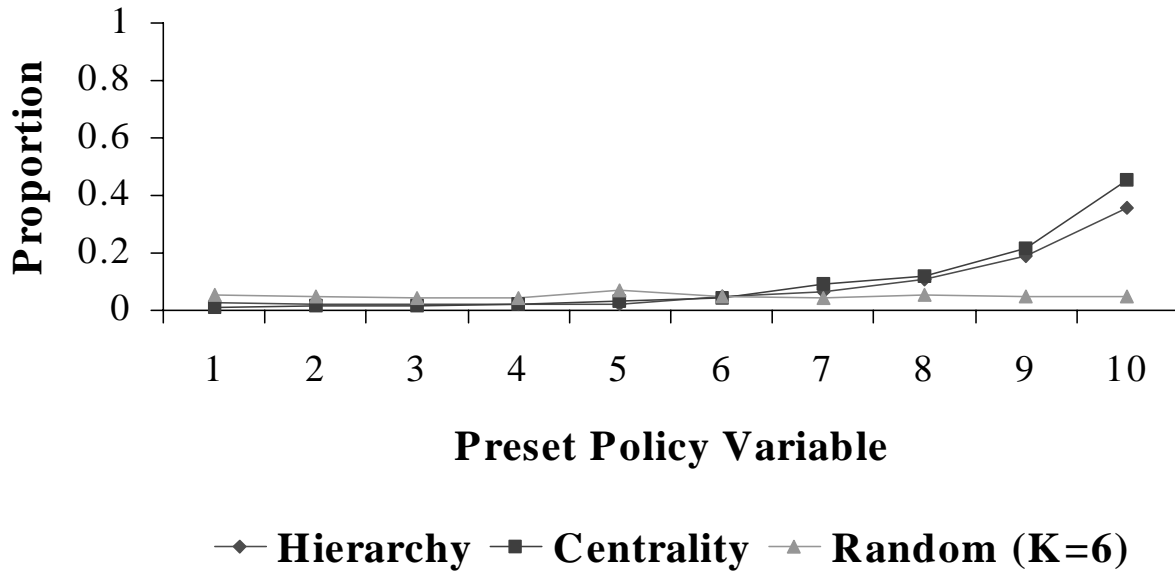
■ Hierarchy (Unconstrained) ◆ Hierarchy (Constrained)

Figure 8. Constraints of History

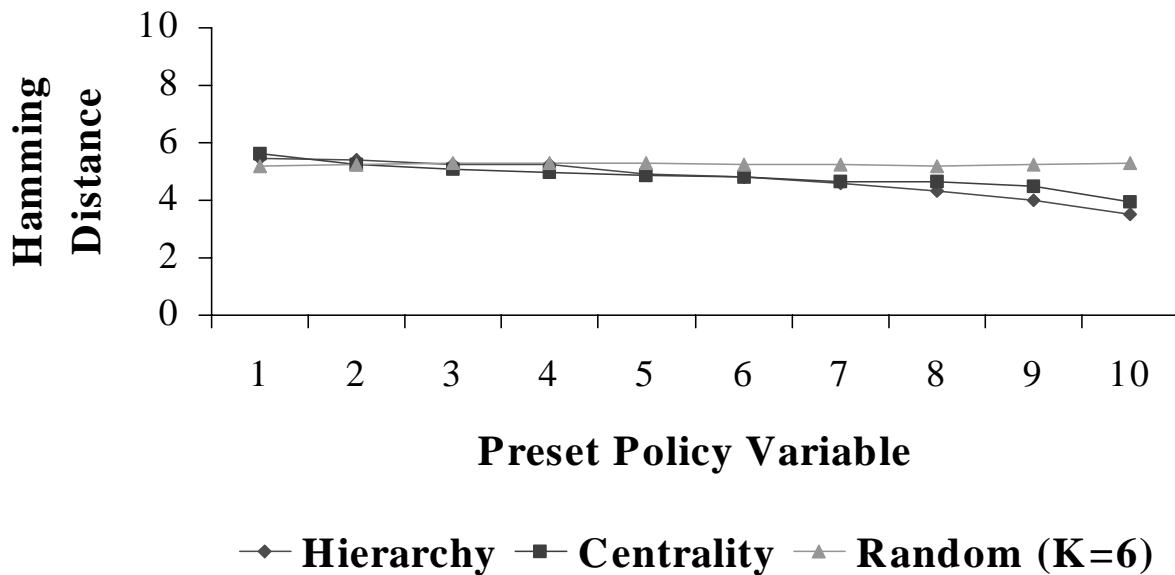


◆ Hierarchy ■ Centrality ▲ Random (K=6)

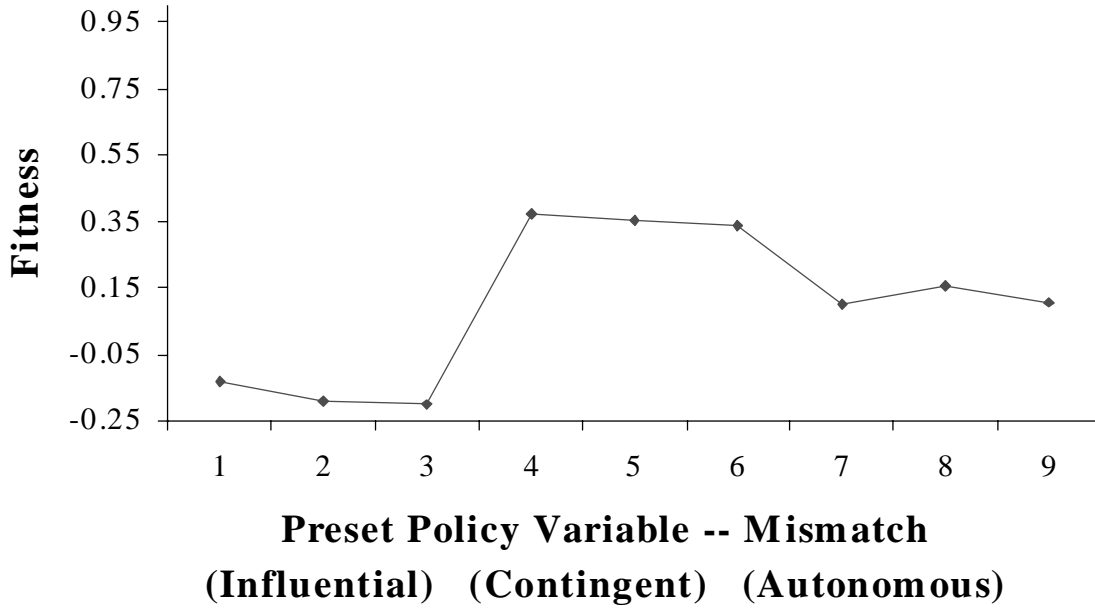
**Figure 9. Constrained Optima:
Proportion that are Local Peaks**



**Figure 10. Constrained Optima:
Hamming Distance**



**Figure 13. Constraints of History:
Normalized Fitness Levels**



**Figure 14. Value of Partially Articulated Activity
Maps**

