Laboratory of Economics and Management

Sant'Anna School of Advanced Studies
Piazza Martiri della Libertà, 33 - 56127 PISA (Italy)
Tel. +39-050-883-343 Fax +39-050-883-344
Email: lem@sssup.it Web Page: http://www.lem.sssup.it/

LEM

Working Paper Series

The Value of Moderate Obsession: Insights from a New Model of Organizational Search

Gino Cattani *
Alex Dorsch **
Sidney G. Winter ***

* The Stern School, New York University, USA
** Purdue University, USA
*** The Warthon School, University of Pennsylvania, USA

2007/03 January 2007

ISSN (online) 2284-0400

The Value of Moderate Obsession:
Insights from a New Model of Organizational Search

Sidney G. Winter

Deloitte and Touche Professor of Management

The Wharton School, University of Pennsylvania

winter@wharton.upenn.edu

Gino Cattani

Assistant Professor of Management

The Stern School, New York University

gcattani@stern.nyu.edu

Alex Dorsch

Purdue University

dorschd@gmail.com

The Value of Moderate Obsession:
Insights from a New Model of Organizational Search

ABSTRACT

This study presents a new model of search on a "rugged landscape," which employs modeling techniques from fractal geometry rather than the now-familiar NK modeling technique. In our simulations, firms search locally in a two-dimensional fitness landscape, choosing moves in a way that responds both to local payoff considerations and to a more global sense of opportunity represented by a firm-specific "preferred direction." The latter concept provides a very simple device for introducing cognitive or motivational considerations into the formal account of search behavior, alongside payoff considerations. After describing the objectives and the structure of the model, we report a first experiment which explores how the ruggedness of the landscape affects the interplay of local payoff and cognitive considerations (preferred direction) in search. We show that an intermediate search strategy, combining the guidance of local search with a moderate level of non-local "obsession," is distinctly advantageous in searching a rugged landscape. We also explore the effects of other considerations, including the objective validity of the preferred direction and the degree of dispersion of firm strategies. We conclude by noting available features of the model that are not exercised in this experiment. Given the inherent flexibility of the model, the range of questions that might potentially be explored is extremely large.

Key words: Rugged Landscapes; Local Search; Cognition; Obsession; Fractal Geometry

The Value of Moderate Obsession: Insights from a New Model of Organizational Search

1. Introduction

Recent years have seen the emergence of a substantial body of theoretical work that makes powerful use of the metaphor of "search on a rugged landscape" as a way to think about organizational decisions (Levinthal 1997; McKelvey 1999; Gavetti and Levinthal 2000; Fleming and Sorenson 2001; Rivkin 2001; Siggelkow 2001; Rivkin and Siggelkow 2003; Zenger and Nickerson 2004). The modeling approach in this literature builds directly on the "NK model" developed by biologist Stuart Kauffman (1993). When Kauffman created NK as an alternative formalization of Sewall Wright's concept of a fitness landscape (Wright 1932), his principal purpose was to explore how interactions among genes at different loci on the DNA molecule affect the course of biological evolution. Recent applications of NK techniques in organization theory have, obviously, addressed very different issues. This research has been grounded in the work of the Carnegie School, and in particular on the ideas about search put forward by March and Simon (1958) and in the classic volume by Cyert and March (1963), *A Behavioral Theory of the Firm.* By providing a new formal apparatus for the expression and investigation of some basic ideas from the Carnegie tradition, the NK model has re-invigorated that tradition and opened new paths for its fruitful development.

In this paper, our purpose is to build on and extend the significant advances made with the NK body of technique. We describe and exercise a piece of modeling apparatus that we have developed, which uses techniques from the field of fractal geometry to address the task of generating "landscapes." At the level of mathematical description of a landscape, this alternative apparatus is very different from the NK model. In its application to organization theory, however, it is a close cousin to the NK work, sharing in particular its line of descent from Carnegie. Our methods of inquiry also parallel those used in NK-based research. Specifically, we report statistical results for simulated searches over large families of randomly generated landscapes, and we relate search outcomes not to the specifics of individual landscapes but to a small number of theoretically significant *parameters* underlying the random constructions. (At the same time, our approach does offer some distinctive advantages in terms of the possibility of understanding what happens in individual searches.)

Beyond the explication of our new research tool, our principal contribution here is to show how the ruggedness of a landscape affects the relative power of search strategies that rely on two types of guidance. The first source of guidance involves the sort of information that typically guides "local search," the ability to assess the advantages of a move to a position neighboring the current one. The second source is a perception that superior alternatives are to be found in a specific direction from the current alternative, the "preferred direction." We conceive of "preferred direction" as a stand-in for a variety of non-local influences on search that we broadly label "cognitive." In our present model, the preferred direction is a fixed attribute of the individual searcher (though varying across searchers); it is independent of experience and of the current position. The model permits us to vary the relative weight given to the two sources of search guidance, and thereby to investigate the differing consequences of cognition vs. local feedback as determinants of search. In the experiment reported here, we show that an *intermediate* weight on non-local considerations can be particularly advantageous in searching a rugged landscape. Our discussion explains how this effect – the "value of moderate obsession" – arises, and explores its implications for organizational behavior.

As just suggested, the preferred direction can be given a variety of interpretations. It could represent prior information, foresight, causal understanding, heuristic principle, identity or even irrational obsession – or simply, "opinion" based on some combination of the foregoing sources. Although it would be quite possible to include in the model some representation of the distinctions among these, we make no such attempt. What is really to the point, we believe, is that such distinctions are not all that easy to make in practice – as is repeatedly illustrated in accounts of organizations searching for better technologies, strategies or organizational arrangements.[1] What Alf regards as a treasured insight may often appear to Betty as irrational obsession. (For a penetrating analysis of one example of this, see the Tripsas and Gavetti (2000) account of Polaroid's efforts in digital imaging.) Across a wide variety of causes or specific interpretations of "preferred direction," our model affords insight into the implications of this general phenomenon for organizational outcomes. It is important to note that "outcomes" can be assessed either at the level of the individual searcher or at the level of the population of searchers, and that the conclusion can easily be quite different at the two levels. For example, some wider interest (e.g., that of society at large) may be particularly dependent on the *best* of the alternatives found by the population of searchers. That consideration may point, under some circumstances, to the merit of having a diverse population of "obsessives," i.e., actors who are relatively indifferent to the local feedback – proverbially, letting the "hundred flowers bloom."

A number of different examples can be invoked as real-world illustrations of the idea of search guided by a preferred direction. In the realm of technology, our formulation can be viewed as a cousin to the idea of a "technological trajectory." The miniaturization trajectory in semiconductors is the leading example of recent years; improvement is consistently sought in the direction of making circuit line-widths small, hence allowing more circuitry on a chip, hence ultimately supporting the creation of more powerful and cheaper computers. In organizations, the quest for quality via process reliability, as in the "Six Sigma" toolkit, involves an overriding commitment to the proposition that the right trend direction for all defect rates is down. Consistent with our general perspective, it is certainly not the case that the powerful guidance provided by these cognitive frameworks is something that can be convincingly rationalized by a comprehensive cost-benefit analysis. It is easy to speculate on possible "downsides" of such commitments, and in some cases the speculations have systematic support. Such criticism does not restrain the searchers much, and indeed it is scarcely plausible to think that the tortoise of cool analysis will often win the race against the hare of passionate determination.

We hope that this paper is convincing on the point that our apparatus has some advantages that the NK approach does not offer. We acknowledge that the reverse claim can be made, without contradiction, on behalf of NK. In short, we conceive of the present contribution as complementary to the NK modeling approach. In general, we are great enthusiasts of this general approach to formalizing Carnegie insights, and we would not be surprised if it could be fruitfully pursued in a number of related but different modeling schemes. Progress along this line might eventually develop the "landscape search" viewpoint to a level where it could be considered as broad, credible and flexible alternative to other basic ways of understanding behavior, particularly the approach that views behavior as consistently optimizing. But our present undertaking represents only another small step in that preferred direction.

In the following section, we review the basic ideas involved in the representation of behavior as search of a landscape. We then explain the main elements of our model, after which we describe and analyze our experiment.

The concluding section discusses the experimental results in a broader perspective and also sketches features of the computer model that we have chosen not to exercise in this first experiment.

2. Search: Localness and Landscapes

The concept of search has long been central in behavioral accounts of decision making, especially those footed on the giant shoulders of the Carnegie School. A number of considerations account for its prominence. Undoubtedly, the key consideration is that the concept has such a high level of face plausibility. From both direct and vicarious experience, and in a wide variety of specific contexts, people know what it is like to be searching for an answer or solution to a problem. They know in particular that while such searches sometimes do take the form of surveys of well-defined sets of alternatives, as in making a selection from a restaurant menu, this is far from being the general case. In many very consequential cases, the decision alternatives have to be *discovered or invented* before a decision among them can be made – consider the search for a house, a spouse or a dissertation topic. This process of discovery/invention typically consumes significant time and resources, and the need to balance these costs against the benefit of improved search outcomes is an important part of the decision maker's problem. These realistic aspects of search are familiar and recognizable across diverse contexts. Thus, the "search" idea has appealing realism and generality. It naturally follows, however, that the word's meaning and status as part of natural language is taken for granted. It may on that account appear in theoretical writing with little reference to whatever specific meaning or status it should have as a *theoretical* concept, and this can pose (or may historically have created) some hazards of under-specification – an issue we pursue in the discussion here.

All of this continues to apply when "search" is further specified as "local search." It is natural – again, across a wide variety of contexts -- to think of a search in progress as having reached a certain stage or position, the nature of which will influence the future course. First, the current position will shape further progress because it is a key influence on the identification of alternatives that might be examined next. Second, the position is a powerful determinant of the differential costs of these alternatives; the idea of "starting over from scratch" is a type of alternative that may be easy to recognize and acknowledge in principle but is generally unappealing in practice because it entails large up-front investments (or on the converse view, abandoning the investments incurred on the way to the present position, the costs of which are sunk). Third, and most important, the alternatives that lie nearby the current position tend to be ones whose outcome consequences are relatively easy to predict because they involve only small changes from the present position.

In sum, from any given position, "nearby" alternatives are visible, generally inexpensive to implement and easy to assess, and thus in the latter sense are low-risk. These three associated features are commonly encountered in the context of search on a geographic landscape, among other contexts, and this fact helps to qualify landscape search as a good metaphor for local search in general.

Note, however, that the identified aspects of "localness" are logically independent. Each, in fact, suggests somewhat different criteria for whether search in a particular situation should be considered local or not. Additional complexity enters the picture in the form of a key refinement of the "easy to assess" point. Does that mean "easy to assess" *without actually implementing* the alternative, or easy only in the context of actual implementation? This is the important distinction between "off-line" and "on-line" search or learning processes (Lippman and McCall 1976; Levitt and March 1988; Gavetti and Levinthal 2000). Within each of these categories, there is a rich variety of possible situations and interpretations. "Off-line" could mean thought experiments, theoretical calculations, simulations, laboratory experiments, pilot plants – all of which are ways of "trying" something different without actually making a significant change in the prevailing "position," the principal ongoing operation that permits the organization to pay its bills. Within the "on-line" category there are ranges of cases distinguished by whether the assessment of the consequences of an actual change is not only plainly feasible, but also prompt, relatively automatic and cheap -- or whether on the contrary it demands time, attention and other resources. There is also the question of whether an on-line experiment is fully reversible or leaves a legacy of some sort, the latter implying that it is not strictly possible to return to the previous position if the experiment (directed to improvement) should have a negative outcome.

Even much of the verbal discussion of local search generally abstracts from some of these diverse complexities, and formal modeling efforts naturally go farther in that direction. The models typically address a set of possibilities covering a rather small subset of the range just sketched. In particular, it is commonly assumed that "nearby" alternatives – which usually means those reachable by the minimum possible step in some discrete space –

are visible and perfectly assessed in the "off-line" mode. If the assessment indicates a superior outcome is available at a nearby position, that step is made, and the outcome is the anticipated one. Further, it is *only* those nearby positions that are available as alternatives in a given period. The "try it and see" aspect of on-line search is often missing, as are the complications of delayed and noisy feedback. All of these very significant restrictions are accepted in the present work. If we simply proceed along that line, however, the great *a priori* appeal of the "local search" idea is considerably diminished. The image of behavior presented is too ant-like – too groping, too uninformed by broader views of the "landscape".

3. Introducing "Cognition"

There is an emerging consensus that this will not do, that it is important to come to grips with the fact that human beings and their organizations (while certainly not the confident global optimizers of economic theory) are (nevertheless) not ants (either). Detailed studies of the ways in which individuals and organizations address problems generally underscore the profound importance of *cognition*, "… a forward-looking form of intelligence that is premised on an actor's beliefs about the linkage between the choice of actions and the subsequent impact of those actions on outcomes" (Gavetti and Levinthal 2000). For those of us who believe that formal modeling often makes a valuable contribution to understanding, the challenge of coming to grips with cognition is a formidable one. Where case studies consistently, and unsurprisingly, reveal great complexity, the role of formal modeling is to provide structures that offer a relatively transparent analysis of key causal determinants. Such a service can hardly be provided by a model that gives full play to the underlying complexity; yet, it is important to avoid simplifications that might create the impression that there are "easy answers" that would always permit organizations to respond, in a reliable and fine-grained way, to the details of complex real-world problems.

Like researchers in the NK tradition, we believe that the rugged landscape metaphor (coupled with random sampling of parametrically characterized landscapes) is an effective theoretical device for representing realistic complexity without getting bogged down in it. It serves to isolate the multiple contingencies affecting individual searches from the large-scale, systematic tendencies that theory might be able to characterize. The metaphor also provides a direct representation of learning from local feedback, and explains its limitations. The question that then arises is the nature of the "cognitive overlay," the representation of considerations that compete with or supplement the guidance provided by local feedback. The exploration of this difficult question has only recently begun.

Gavetti and Levinthal (2000) assume that a cognitive representation corresponds to the specification of a subset N1 of the N dimensions of the policy problem faced by the searcher, i.e., the cognitive representation is a low-dimensional counterpart of the actual landscape. Further, the searcher's impression of the payoff implications of the N1 dimensions is in a sense objectively correct: each cognitive representation is assigned a score that is the average score of all landscapes consistent with it. In "cognitive" search mode, the searcher responds to the simplified picture of the global landscape provided by the cognitive representation, effectively choosing a region of the landscape in which to explore further according to the usual hill-climbing principles of local search. Thus, in this set-up, the cognitive representation may be highly inaccurate, but it does permit the searcher to survey the landscape both globally and objectively. Further, with respect to the N1 dimensions included in the cognitive representation, there are no inertial forces impeding "long jumps" – searchers can accomplish whatever organizational transformation they can envisage.

In unpublished work, Rivkin (1998) has proposed a different "overlay" on the local search process in the NK context. In his scheme, the firm may invest in "integrative capacity." When it has integrative capacity at level M the firm can perceive and assess alternatives that involve simultaneous changes in M of its N policies. Thus, search becomes less local as M increases from 1, and becomes global at M = N. Integrative capacity has a cost that increases with M, so the firm must make a higher level decision about what level of M to invest in. As Rivkin explains, a changing environment renews the need to search for a better solution and hence enhances the benefit that comes with a higher value of M – by increasing the speed and effectiveness of search. This parallels the Gavetti and Levinthal treatment in the sense that the representation of the departure from local search is one that yields an unambiguous improvement in the expected efficacy of search – but since the improvement is costly, the implications for net payoff reflect context-dependent tradeoffs.

Recent work by Nelson (2006) presents a challenge to the whole strategy of treating cognition as an overlay on a landscape search scheme, regardless of the details of the landscape and the search process.[2] In his view, the landscape itself is a cognitive phenomenon, fundamentally reflecting the actor's understanding of what might matter and is worth attending to. Further, the feedback that guides search in Nelson's conception is vastly richer than the local payoff or "fitness" changes contemplated by the NK models (and the present one). Feedback can modify understanding of the present position at different levels, including those that are fundamental to the image of the landscape. While these proposals are not expressed in a formal model, they are offered at a level of abstraction that clearly points to the possibility of at least partial incorporation in a model. This is arguably more urgent in the modeling of off-line search, since the on-line search context is typically one in which there are multiple obstacles to the actual incorporation of the new insights that are obtained. In any case, extending the depth to which cognitive considerations extend, and the realism of their treatment, is a worthy objective for the modeling community. But, as noted at the start of this section, the work of doing this in a way that yields real insight is only beginning.

This brief review of some related work identifies an important aspect that distinguishes our formulation from these others. When changing position, our searchers always *move* locally regardless of the existence of more global cognitive influences on the step taken. In this sense, we maintain our commitment to those aspects of the appeal of "local" search that derive from "inertia" or "inexpensive to implement" – i.e., the fact that localness is often imposed by *non-cognitive* costs or constraints that limit "long jumps" in policy space. Consider, for example, NCR's transition from mechanical to electronic cash registers, as recounted by Rosenbloom (2000). One could argue that the company acknowledged (at a cognitive level) that the electronic age was coming when it acquired an electronics company Computer Research Corporation (CRC) in 1953. At the operational level, however, the commitments to the mechanical technology were profound, and the actual transition occurred many years later, in a context of extreme competitive crisis. This point is closely related to the "on line" vs. "off line" distinction; indeed, it is close to the heart of it. Our model is fundamentally one of on-line search; only the surveying of the local alternatives is an off-line process. We would also refer to this same point if asked to defend the presence of the word "organizational" in our title (considering that our model searchers have no significant internal structure). Organizations have deep structures of commitment to existing ways of doing things. Whatever constraints an individual faces in moving from having a "better idea" to actually implementing it, the constraints in a substantial organization are vastly larger. Thus the dialectic between cognition and feedback is more deeply rooted in the organizational case, and that dialectic is what we want to explore.

Of course, our model is also distinguished by the fact that the "preferred direction" idea is such a simple approach to capturing some of the complex influences of cognition. By the standard of the rich case studies it is indeed very simple – but such a comment could still be made about more complex formal models, at least if the model logic remains understandable. We believe that it is instructive by virtue of, as well as in spite of, its simplicity -- a useful complement to the rugged landscape metaphor, and an understandable platform for more complex formulations. In particular, we can illustrate in an understandable, reproducible context the wisdom of much scholarship and experience, which seems to say that it is possible to go wrong by having too little regard for the feedback – and also by having too much regard. By experimentally varying the weights on the different sources of search guidance, and considering a variety of contexts, we can throw some light on the structure of this basic existential puzzle.

4. Technical Description of the Simulation Model

In this section, we first explain how fractal fitness landscapes are generated and highlight their properties. The significant contrasts between the fractal geometry approach and the NK modeling approach are identified. We then explain how the search of an individual searcher (henceforth, a "firm") is guided.

4.1. Landscape Generation via Fourier Synthesis.

The landscapes in our model are surfaces over a plane, and they bear a reasonable resemblance to actual physical landscapes. The point of the search is to attain as high a point as possible on this landscape; this vertical dimension is called "technical fitness" (or "fitness," for short). In the NK model, by contrast, the alternatives searched are strings of N bits, corresponding geometrically to the vertices of the unit hypercube in N space. Thus, in NK modeling the term "landscape" is metaphorical, not merely in relation to the space of real world alternatives, but also to the model's representation of that space. Our model implements the landscape metaphor in a more literal way. Given that this metaphor is being drawn upon for heuristic guidance in thinking about the challenges of search, we think it is advantageous to have this more direct connection between the formalism and the heuristic imagery.

That it is possible to use computer algorithms to generate realistic landscapes – and random families of such landscapes – is the key fact that we exploit in our approach.[3] This striking fact is an aspect of a broader (and even more striking) discovery about the characteristic geometric features of natural phenomena, due to Mandelbrot (1983). The basic principle is that of self-similarity across scales, as illustrated by the fact that the geometry of drainage basins remains recognizably similar as we go from a great river to its tributaries to the tributaries of the tributaries, etc., until we reach the small gullies on the hillsides. In the computer algorithms, simulated versions of such phenomena are created by a logic in which smaller scale features are superimposed on larger ones in a layered fashion, following a stochastic specification that is the same at all layers, except that it is scaled down in magnitude by a certain fraction at each layer. A variety of algorithms in the toolkit of fractal geometry share this basic feature. (See (Peitgen and Saupe 1988) for discussion of these techniques, including the one we use.)

In the specific algorithm that we employ, the different layers are essentially sine curves of different frequencies. The procedure is a cousin to the mathematical technique of Fourier analysis. In the simplest case of Fourier analysis, a function of a single variable is represented as a sum of sine waves of different frequencies. A wide class of functions can be represented or closely approximated in this manner. Here, however, the functions are of two variables and our concern is not with analysis but with synthesis – i.e., with constructing a function as a sum of trigonometric functions that differ in frequency and phase. We interpret the constructed functions as yielding landscape heights/fitness values. When random elements are included in the construction procedure, the result is a technique for constructing random families of related functions of two variables. The "relatedness" within a family involves properties that matter to the efficacy of search procedures, particularly the degree of ruggedness. Rugged landscapes result when the weight on the higher frequency (short wavelength) components of the sum is relatively high – as is further explained in the subsequent discussion of Figure A 1 in the appendix.

The landscape is represented as a surface over a point grid which we typically choose to be $2^6 = 64$ points on a side, or $2^{12} = 4096$ points total. The latter is the total number of alternatives available to a searcher, and is comparable to the upper end of the range of values commonly encountered in the NK literature. (A commonly used bit-string length is N = 10, which implies $2^{10} = 1024$ alternatives.) While specific interpretations can be offered for the two locational dimensions below the surface, there is no straightforward interpretation of these comparable to the interpretation that NK modeling offers for its N dimensions (each is a "policy"). A feature of our program called "contour control" provides the option of choosing the point at which the longest wavelength component reaches its maximum. For a smooth landscape the direct influence of this control is strong, and the generated landscape's actual maximum point is invariably quite close to the designated location. As we consider increasingly rugged landscapes, the control is manifested in a progressively weaker statistical tendency for the global maximum to lie in the general region of the designated location – a reflection of the higher weights applied at the upper layers of the construction, where fluctuations of shorter period (higher frequency) are represented. Contour control is a feature that expands the options of experimental design relative to what NK modeling provides; we can (but need not) specify roughly where the right answer is.

After an initial landscape has been generated by the method just described, it can be modified in three ways that are significant in relation to the experiments we perform. First we have a simple parametric way to control the objective validity of a preferred direction. The landscape can be tilted in the x or y direction, which means essentially that the height of a planar landscape is added to the initial one: the initial algorithmically-generated fitness value at grid point (x, y) is modified by adding the quantity

$$\Delta z = a * x + b * y$$

where a and b are parameters. Second, we can set a new zero point on the vertical scale, re-setting all initial landscape heights that were below that height to zero. We interpret this zero point – which we call "sea level" -- as corresponding to a level of (technical) fitness that is just at the threshold of "ecological" fitness. For an economic model of evolutionary competition among firms, a simple interpretation of technical fitness is the amount of customer utility produced per unit of cost incurred; sea level then corresponds to a technical fitness value just low enough so that even a monopolist with that fitness could not find a single buyer. (More generally, the ecological fitness yielded by a particular positive value of technical fitness depends on the competition and what it has to offer.) Finally, a rectangular region in the landscape can be directly re-set to zero height, which has the effect of assuring that local feedback plays no role in guiding search in that area.

4.2. Guidance of Local Search

A firm located at a given point in the grid evaluates the eight neighboring grid points and moves to the one with the highest score. The score of a grid point is computed as a weighted average of two components, one based on its fitness value and the other based on the degree to which the point lies in the preferred direction from the starting point. As discussed previously, "preferred direction" is our device for incorporating the role of considerations affecting the search that do not involve the direct assessment of fitness value. A parameter v, $0 < v < 1$ is the weight on preferred direction, while (1-v) weights the fitness value. When $v = 1$ the firm is only concerned with its direction. By contrast, when $v = 0$ the firm is only concerned with immediate payoff: the search is "hill climbing" and consequently terminates when a local peak is reached.

Because the two components of the score are not naturally commensurable, some care must be taken with the scaling of the components if outcomes are to be sensitive to v values over a convenient range.[4] The fitness component is defined as the fitness increment from the given point to the neighboring one, divided by a scaling factor. The default value of the latter is the maximum height on a typical landscape (namely, 100) divided by half the length of the grid side, N/2 (but see the following paragraph). For the direction component, the score is $1 – \theta/90$, where θ is the angle in degrees between the direction of the neighboring point and a perturbed value of the preferred direction. (Note that negative scores are implied when the angle exceeds 90°.) The role of the perturbation is to mediate between the concept of preferred direction, which in principle can vary continuously, and the reality that the grid offers only eight directions of possible movement. The perturbation is a zero-mean normal deviate with a standard deviation which we normally take to be 22.5 degrees. In effect, this device expresses the preferred direction as a "mixed strategy" over the eight directions offered by the grid. It also introduces a random element into the search path, the influence of which can be amplified by increasing the standard deviation of the perturbation, or by increasing v.

When doing experiments like those reported below, which involve comparing the average performance of search strategies across different landscape families; it is important to attend carefully to the vertical scaling of the landscapes. If, for example, we might want to conclude that a given search strategy (characterized, e.g., by the v value) is more effective in landscape family A than in family B, we must take care to assure that this is an interesting comparison – as it would not be if, for example, the peaks of B are typically lower than the low foothills of A. To address this issue, we attempt to scale landscapes in such a way as to make the median value of maximum height a constant across families. Because the distribution of maximum heights depends in a complex way on the parameter values defining the family, there is no completely straightforward way to realize this ambition. Through experimentation, we developed an empirical relation relating the median value of the maximum height to the defining parameters of the family, specifically those controlling ruggedness, sea level and tilt. (The R^2 of the regression underlying this relationship is .998.) We then adjust the scaling factor in each family so that the median value of the maximum height (above sea level) is 100.

4.3. Initial Conditions

At the beginning of a simulation run, the searching firms are assigned their preferred directions at random, by drawing from a normal distribution with specified mean and standard deviation. They are also assigned random initial positions in the point grid, distributed over a specified rectangular region, but in this case the distribution is uniform. We typically interpret initial position differences in terms of capability differences derived from previous experience of the same organizations; experience leaves some better prepared for the new contest than others. We could similarly interpret differences in preferred direction as reflecting different experiences in the past. Thus, the firms are initially heterogeneous in two key respects that may affect search outcomes.

5. An Illustrative Experiment

We now describe an experiment that illustrates the effects of guiding local search not just by the locally assessed advantages of different moves, but also by the higher-level cognitive considerations represented by preferred direction. More specifically, we show how the consequences of giving weight to 'preferred direction' depend on the ruggedness of the landscape, exploring three values of the weight on preferred direction and two of ruggedness. Our experimental manipulation involves three considerations in addition to these two: (i) the mean preferred direction of the population of firms (3 conditions), (ii) the dispersion of firm direction preferences around that mean (2), (iii) the objective validity of the preference, as determined by the presence or absence of a corresponding tilt in the landscape (2). We determine results for a total of 72 different conditions in all, as shown in Table 1.

Across all conditions, we posit a situation in which firms initially have zero ecological fitness; they are at sea level. Thus, the search processes we consider can be interpreted as representing the efforts of firms to find viable ways of participating in a new industry – an industry that exists only as a latent potential when the story starts.[5] As a firm begins its search, it cannot be guided by local payoff because there is none; it is necessarily guided only by the cognitive considerations summarized by the preferred direction. Firms take different approaches in this uncertain situation, and have different positional advantages, because they (or their founders) differ in their prior experience. The intimations of opportunity that motivate these searches are broadly correct; there are valuable things to be found in the general direction in which firms are searching. Further, the most advantageous position of all is likely to be found toward the middle of the general field of search. However, the differences among firms and the combined contingencies of specific landscapes and searches make the outcomes highly uncertain at the individual firm level. While there are systematic features in patterns of firm success and failure, they would be hard to discern without the aid of the large samples that simulation makes possible. Simulation thus provides what reality can only suggest, an image of the reliable causal links among situations, strategies and outcomes that in individual cases are deeply obscured by the random noise.

In the experiments, the parameters of the model are set to values that capture the themes in the preceding paragraph regarding the context of the search. Given an initially generated landscape, the origin of the vertical scale (sea level) is re-set to a value 20% of the way from the min to the max of fitness on that original landscape ($s = 0.2$). Contour control is set to locate the peak of the long wavelength component of the fractal structure to occur at grid coordinates (32, 32). Under these conditions, the neighborhood of the origin is at sea level in the great proportion of the landscapes generated. In particular, the box with corners at grid points (0, 0) and (5, 5), wherein we randomly distribute the initial positions of firms, is typically at sea level. However, to assure that this holds in all cases, we also specify directly that this box is at sea level.

5.1. Parameter Settings for Experimental Conditions

Here we characterize our experimental conditions in terms of specific parameter settings in the model. The two values of the parameter determining ruggedness are $h = 1.25$ (rugged) and $h = 1.75$ (smooth). In the rugged case, landscapes typically have a substantial number of local maxima. In the "smooth" case local maxima are fewer, but cases with a unique local maximum are uncommon. Illustrative examples of landscapes and search paths are shown in Figures 1 and 2.

The second of our focal variables represents the extent to which firms are guided by local feedback or by the broader cognitive considerations summarized in the preferred direction. We vary the value of v from 0.2 to 0.5 to 0.8. While a high value of v (e.g., 0.8) corresponds to the case where firms disregard feedback and adhere to their preferred direction; a low value of v (e.g., 0.2) implies that firms make local moves primarily on the basis of a local quest for higher ground on the fitness landscape.

Among the other three experimental variables, the most important is the tilt of the landscape, which affects the objective validity of the firm cognitions summarized in preferred direction. In the tilt condition, the high ground of fitness tends to lie toward the North-Northeast. The parameter values are a=0.5, b=1. At these values, the tilt effects are substantial relative to the randomness in the landscape generation process, but not so large as to reliably overwhelm that randomness. There is no obvious way to determine the average effect of the tilt on the height of local peaks on a landscape, since the tilt affects both the number and location of such peaks, the locations are partly random, and the specific increment the tilt adds to a peak depends on its location.

To gauge the general effect of different preferred directions on outcomes, we use three different values for the mean direction vector relative to the origin – i.e., 22.5 degrees or "East Northeast," 45 (Northeast), and 67.5 (North Northeast). The latter value is broadly consistent with the "reality" of the landscape in the "tilt" condition. We also set two conditions for the degree of individual firm variation around the mean preferred direction, using two different values for the direction standard deviation of the controlling normal distribution – namely, 7.5 (low variation) and 22.5 (high variation).

5.2. Hypotheses and Conjectures

One function of this first experiment is to verify that the model does the "obvious" things that its motivating logic would suggest. We expect, for example, that it pays to search in an objectively superior direction. This implies that in the absence of tilt the payoffs should be larger on the average when the mean preferred direction is 45 and, under the tilt condition, they should be larger when the mean preferred direction is 67.5.

On rugged landscapes (low values of h), we expect lower payoffs than on smooth ones, because the search process will tend to "hang up" on local peaks. We expect a moderate v to be more favorable than low v on rugged landscapes, since this should provide some resistance to the tendency to stop the search at a local peak encountered early on. On the other hand, the high v setting may lead the searchers to charge over the high ground and wind up on lower ground toward the eastern and northern edges of the grid – a symptom of this would be a pattern of higher payoff at intermediate stages of the search than at the end, and it is particularly likely when the landscape is smooth and not tilted. Finally, under the no-tilt condition, greater variation should improve results when the mean preferred direction angle is low (22.5) or high (67.5). In these cases, greater variety should mean that a few firms that deviate toward the center from the mean preferred direction should tend to do better, and probably, as a result, pull up the average (since sea level limits the damage that misguided firms can do to the average). Under the tilt condition, greater variation should improve average results when the mean preferred direction is at the unfavorable value (22.5).

5.3. Experiment Results

We generated the results for each family of landscapes as follows. First, for a given value of v – with v ranging from 0.2 to 0.5 to 0.8 – we recorded fitness values for 500 individual firms' over 250 and 500 periods across 100 randomly generated landscapes. Our model gives us control of the probability that an individual firm searches in a particular period, and it is here set to 0.2. Thus, the two durations correspond to expectations of expected searches equal to 50 and 100 (with actual values independent across firms). We computed averages and standard deviations for the (average) fitness values. We recorded these values at different search durations in recognition of the point that search outcomes can depend on search duration – and, as we shall see, longer is not always better. We further tested the robustness of our findings by running the experiment over 500 (not just 100) randomly generated landscapes, again for a sample of 500 firms observed over 500 periods. The results, which are available from the authors upon request, turned out to be quite similar to those reported here.

Figures 1 and 2 display the 500-period search path of a single firm under each of six conditions – for rugged and smooth landscapes and each of the three values of v (no tilt). Prominent features of these individual paths correspond directly to prominent patterns in our overall quantitative results. In particular, under the high v condition the firm reaches a high point on the landscape but then abandons it to proceed further in the preferred direction. This behavior accounts for the principal contrast between average results at period 250 and period 500, as we further discuss below. (The landscape in the bottom panel of Figure 2 is the same as in the top two panels, but shown in a different orientation to make the path visible.)

The results of the simulation for the full 500 periods are summarized in Table 2, which reports average fitness values under all experimental conditions. The main tendencies of the results are displayed graphically in Figures 3a-3b, which show one pair of bar charts for the high dispersion case at periods 250 and 500, when the mean preferred direction is 45 degrees. As expected, the intermediate value of v (0.5) is particularly advantageous when the landscape is rugged, and the low value is advantageous when the landscape is smooth. The high value of v (0.8) is disastrous when the landscape is not tilted, and gets worse the longer it is pursued – these searchers are so committed to their *a priori* views about direction that they are effectively divorced from reality. On the other hand, when the *a priori* belief is modestly aligned with the reality (as it is in the tilt condition), its consequences are not nearly so bad, although still inferior to the results achieved at lower v values. It is also noteworthy in Figures 3a-3b that the intermediate value of v produces results that are highly robust to the conditions.

The average fitness values support a finer-grained analysis of the relative effectiveness of feedback-based and oriented search under different scenarios, as indicated in Tables 3 and 4. These tables report means and standard deviations (across landscapes) for average fitness values (across firms), under all of the different conditions. All of the pair-wise differences in average fitness across different v values, under otherwise identical conditions, are statistically significant at the .01 level or better. As the tables show, the large sample size permits us to affirm the significance of differences that are quite modest in absolute terms, including ones that are virtually invisible on visual inspection of Figures 3a-3b. While the differences that are *plainly* visible in these figures are overwhelmingly significant, they remain only tendencies, contradicted with substantial frequency in individual cases. The lesson taught clearly favors moderation in balancing the influence of feedback and *a priori* belief, but it provides no interesting reassurance about the likelihood of avoiding regret.

The results show that the relative weight given to adherence to the preferred direction (v) vs. the local feedback does shape the effectiveness of a firm's search behavior. The results vary significantly with the ruggedness of the landscape. In rugged landscapes (h=1.25), firms tend to get stuck on a relatively low local peak for low values of v (as in Figure 1, top). By contrast, in smoother landscapes, the likelihood of settling on a higher peak, perhaps even the global maximum, is highest at the lowest value of v (Figure 2). Overall, the results suggest that search characterized by a moderate degree of 'obsession' or 'overconfidence' (v=0.5) is beneficial when searching a rugged landscape. It provides protection against an early stop on a low local peak, without necessarily precluding a later stop at a much higher level. On a smoother landscape, the hazard of the early stop is much lower while the hazard of a dramatic overshoot of the optimum is much enhanced – thus a lower value of v will give a higher level of performance.

Tables 5a and 5b report significance tests statistics on the difference between mean fitness values computed after 250 and 500 periods – again for different values of h, v, preferred direction, standard deviation of direction, with and without tilt. In the absence of tilt (see 5a), a high degree of "obsession" consistently leads to a decline in average fitness values with the number of periods. In other words, if firms disregard feedback in favor of their *a priori* notions, they will wind up searching in a region of the landscape where available payoffs tend to be lower. At lower levels of v, the extension of search on the untilted landscapes generally produces small and typically insignificant increases in fitness. In the presence of tilt (see 5b), the continuation of search when v is high still diminishes the achieved fitness in 9 of 12 cases. The other 3 cases all occur in the condition where "obsession" is most aligned with reality; the mean preferred direction is 67.5. In two of those the extension of search increases average fitness modestly but significantly, while in the other case the difference is insignificant. On landscapes that are smooth and tilted (5b, right-hand panel), continuing search tends to reduce payoffs even when v = .5, except when that moderate obsession is itself aligned with reality.

Contrary to our original expectation, "variety" (variation around preferred direction) does not seem to be beneficial. Although not reported, t-tests on the difference between mean fitness values indicate that on the average firms with low (7.5) standard deviation around their preferred direction do significantly better than firms with high standard deviation (22.5). Our hypotheses on variety were based on the supposition that the favorable effects of a difference from the mean direction would tend to outweigh the negative effects. This analysis may have given too much emphasis to the role of the downside limit imposed by sea level, and too little emphasis to the effect of the underlying concave shape imposed by the centering of the average max at (32, 32). We intend to explore variety further in experiments that introduce the evolutionary selection features of the model, which were not involved in the current experiment. When search is costly and the realized fitness of a firm determines its ability to search further, the average fitness of surviving firms should be expected to be strongly influenced by the most fit among them. This mechanism corresponds, of course, to the role of variation in biological evolution.

6. Discussion

We now step back from the details of the experiment and return, informed by the experiment, to the substantive theme of how the model relates to the phenomena of organizational search in the real world. We consider first the interpretation of the "ruggedness" metaphor, then consider the basis for the experimentally revealed fact that an intermediate v value is favored in rugged landscapes, and then discuss further implications and directions for future work.

6.1. Interpreting "Ruggedness"

The obvious implication of a rugged landscape is that it involves a multiplicity of local peaks, and therefore is an unpromising environment for a search process that is conducted solely by local hill climbing. A strongpoint of the NK research tradition in organization theory is that it goes beyond the obvious. It links the model phenomenon of landscape ruggedness to underlying substantive considerations, and therefore provides at least some broad guidance on what markers might characterize a rugged landscape in an organizational situation. In the NK scheme, ruggedness reflects complexity, and complexity arises from interaction among the N parameter choices, as parameterized by K. A problem can be large (N) without necessarily being on that account complex, since in the absence of interaction its solution is easily decentralized. It is complexity that stands in the way of an organization trying to find the best answer to its problem, the true optimum.[6] The linkage from interaction (K) to complexity and ruggedness is provided by the theorems characterizing the NK model, the proofs of which depend on the special structure of the model's random procedure for generating the payoff (fitness) function.

To interpret ruggedness in the fractal landscape model, one alternative would be simply to import the NK interpretation – we have ruggedness, hence complexity, hence there must in some sense be interaction "underneath." There is, however, a more fruitful line of interpretation reachable by exploring the deeper meaning of fractal structure. Organizational capabilities can be viewed as nested structures generated by iterated "how?"

questions. Specification of broad objectives linked to the organization's survival and performance gives rise to a set of questions about how those objectives might be achieved, and then tentative answers to those questions generate another layer of implementation questions, and then that layer another layer, and so on, until all details of the performance are settled by one mechanism or another. In some cases, there is a lot of "equifinality" among choices at a given level, i.e., the part of the outcome called "performance" is not highly sensitive to the details, even if there is significant outcome variety in other respects. In other cases, this is not true at all—"for want of a nail, the shoe was lost, for want of a shoe, the horse was lost …" etc. Situations where "details matter" are represented by rugged landscapes; a short step can change the performance quality drastically. Situations where adherence to a few basic principles guarantees strong performance correspond to smooth landscapes. The successive layers by which the landscape is built correspond, on this interpretation, to the succession of increasingly narrow "how?" questions that are answered in a specific course of organizational action.

This interpretation inspires a question: is there any reason to think that organizational knowledge actually has a *fractal* structure as distinguished from a more general *layered* structure? The answer is not obvious, and certainly there is no obvious way by which one might attempt to resolve the question empirically. Speculatively, however, it seems reasonable to suggest that (for example) a performance that is "all art and no science" might be one in which the details are not merely important, but in fact more important than any base-level principles. The relationships of amplitude weights at different layers might actually have a pattern opposite to that posited in fractal structure.

6.2. The Value of Moderate Obsession

The most interesting of our experimental results is the superiority of the intermediate v value when the search is on a rugged landscape. That this sort of result might emerge was anticipated as we began our effort to characterize a search process blending local feedback and global cognition – but it was not anticipated with high confidence. It seemed quite possible that a v value large enough to allow escape from local minima would also be strong enough to send the searcher charging across the high ground and onward to whatever outcome might lie at the edge of the landscape, in the preferred direction. In other words, it seemed *a priori* that there might be little or no intermediate range of v values separating those that generate the types of outcomes we have illustrated by v = 0.2 from those that generate the types illustrated by v = 0.8. That the experiment generated a strong signal at the obvious (but arbitrary) intermediate value of v = 0.5 came as something of a surprise. The sources of this result deserve careful probing, for the sake of understanding the substantive message that the model is conveying.

In the context of the specific search process described above (sec. 4.2), what is it that stops a searcher from continuing to move locally in the preferred direction? There are two factors in the answer, (i) the landscape slopes downward in the preferred direction, so the local feedback is not encouraging, (ii) the weight (1-v) on that feedback is strong enough so that it overcomes the pull of the preferred direction. The firm may stop permanently when this is the case, or it may move in another direction if a sufficiently favorable fitness increment can be found at one of the other neighboring points, without too much sacrifice of direction. (The random element in the "mixed strategy" representation of preferred direction also plays a role, so that we sometimes see searches terminating in a random "dither" among a small number of approximately equal-fitness grid points.)

What is key here is this. At a given value of v (less than one), a local peak will be a stopping point for the search if the descent from that peak to all the neighboring points is steep enough. Similarly, a "hillside" position facing a down-slope in the preferred direction may be a stopping point if the down-slope is steep enough, and if the fitness gains in other directions are small. The latter possibility is easily illustrated in the model by creating a very smooth hill and allowing searchers with diverse preferred directions to locate the top of it (e.g., by setting v = 0.001). We then intervene and increase v to, for example, v = 0.1. The result is that the searchers move a short distance from the top, each in its preferred direction, and then come to a stop – where the hill slope becomes steep enough to deter a further move. Increase v a bit further, and they will move a bit further from the global peak, and stop – sketching a larger ring around the maximum point. But there comes a point at which the hill does not present a steep enough slope to offset the prevailing v value, and as v is increased beyond that point all searchers proceed recklessly to the edge of the landscape, each in its preferred direction.

In short, it is steep slopes that stop movement in the preferred direction, and therefore make it possible for a moderate v value to be consistent with a full stop at or near the global fitness peak. This will tend to happen if the global peak is itself a pointed local peak, or if there are such pointed local peaks in its vicinity.[7] The value of moderate obsession is realized not simply because the landscape is rugged, but because the landscape's fine structure permits movement from local maxima at low fitness levels, and then inhibits movement at the local maxima found at high fitness levels. Interestingly, the required pattern is characteristic of many natural landscapes: flat plains give way to rolling low hills, then steeper hills, then rugged mountain terrain, then pointed peaks that only well-equipped professional climbers should attempt. Mandelbrot's appreciation of "the fractal geometry of nature" apparently contains hidden lessons for organizational search.

6.3. Smooth Hill Tops and Tight Coupling

The logic by which moderate obsession works well when the high ground is occupied by sharp peaks has its contrapositive: it works poorly when the high ground is a smooth hill. Although the search for the global maximum is greatly facilitated when the landscape is smooth and has a single maximum point, the problem of *holding on* to the optimal solution becomes very difficult under that same condition. If we assume that "fitness" represents the true interest of the organization, and that the cognition represented by preferred direction is supposed to be in the service of that interest, we see that a smooth hill top presents a problem. Even a small positive v is enough to produce a substantial displacement from the location of the maximum. Larger displacements can arise as this logic applies one step at a time, until the policy is not only wrong, but the fitness outcome is substantially degraded.

Stepping outside of the strict logic of the model, we can see that this problem is more serious than the model makes it appear. It is not only constructive attempts at cognitive understanding that can create such distortion, but also direction preferences grounded only in fantasy, true obsession, or parochial interests that are opposed to the fitness interests of the organization. Also, the search process in the model is based on the perfect effectiveness of off-line evaluation of local moves, which also entails the assumption that feedback is prompt and non-stochastic. These model commitments are clearly optimistic. Without them, our model searchers would not stop when the smooth hill slope gave way to a cliff edge, they would proceed onward to disaster – and it appears that real world searchers sometimes do precisely that (see, e.g., (Starbuck and Milliken 1988). Absent clear warnings from local feedback, the logic of "this has worked well up to now" is hard to resist at the individual level, and very hard indeed to resist in an organization.

This line of thinking supports the view that it is sometimes desirable to replace a smooth landscape by a rugged one, to the extent that discretionary features of organizational design can accomplish that, and if the context is one where the problem faced is more about holding on to an achieved solution than locating a solution for the first time (Levinthal and March 1993). While it might seem perverse to intentionally create systems that exact substantial penalties for minor departures from nominal procedures, the image of the increasingly hazardous journey across a smooth hilltop supports the wisdom of creating artificial alarm systems where natural ones do not exist. Thus, "just-in-time" inventory policy, or "lean manufacturing" more generally, involves a willingness to reduce the buffering of the current process and allowing remediable deficiencies to be revealed, thereby enhancing learning. The impact of such approaches depends crucially on what actually happens when the alarms go off. Unfortunately, experience shows that "worked up to now" is powerful enough to induce complacency even as the alarms continue to sound in the background, or are simply disconnected (Perrow 1984). Awareness that the alarms are a feature of a deliberate design, which is strengthened when the short-term consequences of alarms are minimal, contributes to this result.

6.4. Future Work

A number of possibilities for future work have been alluded to along the way. One could, for example, parse the "preferred direction" concept into one or more of its constituent interpretations and modify the model accordingly. The narrow informational interpretation would be an easy target: one could establish the preferred directions in an initial stage in which each actor observes a sample of fitness values from random locations in the landscape, deduces an impression of where the high ground lies, and chooses preferred direction accordingly (presumably, as a consistent orientation to that high ground rather than a globally constant direction).[8] In that scheme, a larger sample size makes the preferred direction more "objective," and less variable across searchers. Strong information suppresses variety, and one can readily imagine that this could have interesting adverse consequences at the level of the population as a whole, and especially so when the landscape is rugged.

Expanding the canvas even more, we could eliminate the effective off-line evaluation of local alternatives, and instead allow the searchers to retrace their steps – or try to, with some chance of error. We could admit that firms might try to learn the best value of v, or respond to the fact that competitive pressure (decreasing ecological fitness) requires an adventurous departure from a local peak of technical fitnesss.

At the top of our own agenda are questions that can be investigated with the present version of the model. We can exploit existing features that involve costly searches and costly moves, initial financing of firms, accumulation of income from operations based on realized fitness, and firm viability requirements. These features, collectively known as the "search economics" of the model, expand our story of organizational search into a story of industry dynamics -- a model that is a cousin to many predecessors in the evolutionary economics tradition, as well as to those in the NK tradition of organization theory. Among other things, we can use the model to examine some of the questions that motivated it in the first place – questions involved in the task of understanding the logic of the early stages of a new industry in which participating firms are diverse in their relevant backgrounds and capabilities (Levinthal 1998; Cattani 2005; Cattani 2006).

We conclude with one reminder to our readers. Remember, the model firms on our model landscapes cannot see the landscapes. *You* can see the landscapes because we have made that view available to you, and we hope that the view will not stop you from empathizing with the firms in their existential situation. Their "global view" is simply their cognition about where the high ground lies. Imperfect as it is, it is all they have. One step at a time, they proceed by trying to balance the guidance of that cognition against the guidance of local feedback – or alternatively, by trying to balance passion and discipline.

Acknowledgments

Financial support for this research from the Mack Center at Wharton is gratefully acknowledged. For helpful guidance and suggestions, we are indebted to Witold Henisz Constance Helfat, Daniel Levinthal, and Johannes Pennings. Useful comments on an earlier draft were received from two anonymous referees, as well as from participants in the Conference on the Behavioral Theory of the Firm, Carnegie Mellon University, May 2006. We would like to acknowledge, particularly, the valuable comments of Jan Rivkin. The responsibility for our imperfect understanding and/or flawed implementation is our own.

<u>Table 1</u>

Parameter Values Characterizing Experimental Conditions

1. **Ruggedness of the Landscape, 2 values**

   (i) Rugged: h= 1.25, (ii) Smooth: h = 1.75

2. **Weight on Preferred Direction, 3 values**

   (i) Low: v = 0.2, (ii) Medium: v = 0.5, (iii) High: v = 0.8

3. **Mean Preferred Direction, in degrees, 3 values**

   (i) East-Northeast (ENE): 22.5, (ii) NE: 45, (iii) NNE: 67.5

4. **Variation (standard deviation) of firm values around Mean Direction , in degrees, 2 values**

   (i) Std = Low: 7.5, (ii) High: 22.5

5. **Tilt to North-Northeast (NNE), fitness increment per grid step, 2 values**

   (i) None: a=b=0, (ii) Positive: a = 0.5, b = 1.0

In all conditions: Sea level is set 20% of the way from the min to the max of the initially generated landscape (s = 0.2). Vertical scaling implies that the median of the max height across landscapes in the family is 100 (approximately). Initial positions are random over the grid region (0, 0) to (5.5) and that region is set to sea level. Contour control is set to (32, 32), i.e., the global peak is likely to be in the vicinity of that grid point, especially in the high h (smooth) condition.

Table 2

Summary of the Results: Average Fitness Values at Period 500

Std Dev	Tilt		dir = 22.5		
		\ v =	0.2	0.5	0.8
7.5	No tilt	Rugged	62.22	81.37	35.52
		Smooth	93.02	91.70	28.69
	Tilt	Rugged	82.31	90.22	69.87
		Smooth	100.28	88.17	74.47
22.5	No tilt	Rugged	60.14	78.29	26.87
		Smooth	90.92	88.00	17.94
	Tilt	Rugged	79.48	85.68	58.56
		Smooth	99.19	86.34	60.49
			dir = 45		
		\ v =	0.2	0.5	0.8
7.5	No tilt	Rugged	64.38	87.10	20.88
		Smooth	95.61	94.49	8.82
	Tilt	Rugged	84.44	93.93	73.37
		Smooth	101.05	90.50	72.55
22.5	No tilt	Rugged	62.82	83.32	28.08
		Smooth	93.46	92.08	19.68
	Tilt	Rugged	81.86	90.78	71.22
		Smooth	100.52	90.68	74.15
			dir = 67.5		
		\ v =	0.2	0.5	0.8
7.5	No tilt	Rugged	62.43	82.14	35.70
		Smooth	93.43	92.90	29.24
	Tilt	Rugged	80.99	92.03	82.47
		Smooth	100.69	94.09	85.94
22.5	No tilt	Rugged	60.89	78.54	27.54
		Smooth	91.21	88.85	18.96
	Tilt	Rugged	79.49	88.54	72.08
		Smooth	99.69	90.85	72.51

Table 3

Means and Standard Deviations of Average Fitness*
No Tilt, and periods=500

	Rugged (h=1.25) v=0.2	Rugged (h=1.25) v=0.5	Rugged (h=1.25) v=0.8	Smooth (h=1.75) v=0.2	Smooth (h=1.75) v=0.5	Smooth (h=1.75) v=0.8
Mean Direction =22.5						
Std=7.5	**62.22**	**81.37**	**35.52**	**93.02**	**91.70**	**28.69**
	(3.21)	(4.66)	(5.65)	(3.04)	(3.39)	(7.44)
Std=22.5	**60.14**	**78.29**	**26.87**	**90.92**	**88.00**	**17.94**
	(5.51)	(9.66)	(10.19)	(7.10)	(10.52)	(11.96)
Mean Direction =45						
Std=7.5	**64.38**	**87.10**	**20.88**	**95.61**	**94.49**	**8.82**
	(1.35)	(1.89)	(5.26)	(1.16)	(0.87)	(5.90)
Std=22.5	**62.82**	**83.32**	**28.08**	**93.46**	**92.08**	**19.68**
	(3.21)	(5.91)	(8.85)	(3.93)	(5.57)	(11.96)
Mean Direction=67.5						
Std=7.5	**62.43**	**82.14**	**35.70**	**93.43**	**92.90**	**29.24**
	(1.90)	(4.49)	(6.06)	(2.27)	(2.19)	(7.85)
Std=22.5	**60.89**	**78.54**	**27.54**	**91.21**	**88.85**	**18.96**
	(4.26)	(9.97)	(10.51)	(6.26)	(9.58)	(12.45)

Table 4

Means and Standard Deviations of Average Fitness*
Tilt (a=0.5; b=1), and periods=500

	Rugged (h=1.25) v=0.2	Rugged (h=1.25) v=0.5	Rugged (h=1.25) v=0.8	Smooth (h=1.75) v=0.2	Smooth (h=1.75) v=0.5	Smooth (h=1.75) v=0.8
Mean Direction =22.5						
Std=7.5	**82.31**	**90.22**	**69.87**	**100.28**	**88.17**	**74.47**
	(2.35)	(2.79)	(6.12)	(0.67)	(0.63)	(5.80)
Std=22.5	**79.48**	**85.68**	**58.56**	**99.19**	**86.34**	**60.49**
	(7.06)	(11.74)	(22.74)	(3.27)	(7.09)	(24.20)
Mean Direction =45						
Std=7.5	**84.44**	**93.93**	**73.37**	**101.05**	**90.50**	**72.55**
	(1.19)	(0.72)	(2.80)	(0.24)	(1.43)	(4.28)
Std=22.5	**81.86**	**90.78**	**71.22**	**100.52**	**90.68**	**74.15**
	(4.34)	(6.07)	(12.95)	(1.18)	(3.39)	(11.51)
Mean Direction=67.5						
Std=7.5	**80.99**	**92.03**	**82.47**	**100.69**	**94.09**	**85.94**
	(2.76)	(2.68)	(4.41)	(0.69)	(1.11)	(6.45)
Std=22.5	**79.49**	**88.54**	**72.08**	**99.69**	**90.85**	**72.51**
	(6.58)	(9.15)	(13.82)	(2.33)	(6.40)	(16.11)

* For each landscape, average fitness over 500 searching firms is recorded. Means and standard deviations of these values are computed over 100 randomly generated landscapes with identical parameter settings. All pair-wise differences across v values across v values are statistically significant at p =.01 (two-tailed).

Table 5a

Test Statistics for Means Differences in Averages Fitness Values after 250 and 500 periods
h=1.25 and h=1.75 - No Tilt

Std	Direction	h=1.25 Periods=250	h=1.25 Periods=500	Test Statistics (500 vs. 250)	h=1.75 Periods=250	h=1.75 Periods=500	Test Statistics (500 vs. 250)
	Mean Direction = 22.5						
7.5	v = 0.2	62.18	62.22	0.18	92.75	93.02	1.39
	v = 0.5	81.16	81.37	0.72	91.49	91.7	0.83
	v= 0.8	57.96	35.52	-52.24**	53.25	28.69	-45.59**
22.5	v = 0.2	59.94	60.14	0.55	89.97	90.92	1.84
	v = 0.5	77.8	78.29	0.78	87	88	1.36
	v= 0.8	51.54	26.87	-30.90**	46.81	17.94	-32.38**
	Mean Direction = 45						
7.5	v = 0.2	64.35	64.38	0.28	95.37	95.61	3.29**
	v = 0.5	86.89	87.1	1.78	94.95	94.49	-7.96**
	v= 0.8	62.94	20.88	-150.77**	57.4	8.82	-144.47**
22.5	v = 0.2	62.76	62.82	0.3	92.99	93.46	1.73
	v = 0.5	83.02	83.32	0.77	91.85	92.08	0.57
	v= 0.8	58.25	28.08	-49.20**	53.32	19.68	-45.13**
	Mean Direction= 67.5						
7.5	v = 0.2	62.38	62.43	0.39	92.97	93.43	2.97**
	v = 0.5	81.82	82.14	1.11	92.68	92.9	1.27
	v= 0.8	57.62	35.7	-49.44**	54.33	29.24	-51.12**
22.5	v = 0.2	60.69	60.89	0.69	89.93	91.21	2.66**
	v = 0.5	77.98	78.54	0.89	87.25	88.85	2.25*
	v= 0.8	51.55	27.54	-29.99**	46.63	18.96	-31.29**

* $p < 0.05$, ** $p < 0.01$ - Two-tailed tests

Table 5b

Test Statistics for Means Differences in Average Fitness Values after 250 and 500 periods
h=1.25 and h=1.75 - Tilt (a=0.5; b=1)

Std	Direction	h=1.25 Periods=250	h=1.25 Periods=500	Test Statistics (500 vs. 250)	h=1.75 Periods=250	h=1.75 Periods=500	Test Statistics (500 vs. 250)
	Mean Direction = 22.5						
7.5	v = 0.2	81.76	82.31	3.47**	99.4	100.28	16.37**
	v = 0.5	89.3	90.22	4.51**	91.84	88.17	-48.31**
	v= 0.8	71.78	69.87	-3.92**	77.2	74.47	-5.73**
22.5	v = 0.2	77.92	79.48	2.88**	96.8	99.19	6.16**
	v = 0.5	82.94	85.68	3.20**	86.72	86.34	-0.59
	v= 0.8	61.77	58.56	-2.19*	65.52	60.49	-3.29**
	Mean Direction = 45						
7.5	v = 0.2	84.19	84.44	3.37**	100.3	101.05	39.41**
	v = 0.5	93.43	93.93	10.60**	93.63	90.5	-44.55**
	v= 0.8	81.7	73.37	-53.44**	85.73	72.55	-61.16**
22.5	v = 0.2	81.28	81.86	1.97	99.36	100.52	9.20**
	v = 0.5	89.1	90.78	3.70**	92.31	90.68	-6.43**
	v= 0.8	74.43	71.22	-3.66**	80.28	74.15	-7.96**
	Mean Direction= 67.5						
7.5	v = 0.2	80.54	80.99	2.52*	99.51	100.69	20.89**
	v = 0.5	90.08	92.03	9.44**	93.77	94.09	3.76**
	v= 0.8	79.49	82.47	8.84**	84.59	85.94	3.09**
22.5	v = 0.2	78.64	79.49	1.84	97.78	99.69	8.04**
	v = 0.5	85.79	88.54	4.08**	89.01	90.85	3.50**
	v= 0.8	71.96	72.08	0.13	74.91	72.51	-2.30*

* $p < 0.05$, ** $p < 0.01$ - Two-tailed tests

Figure 1: Searching a Rugged Landscape, Differing Values of v

h = 1.25, No Tilt, Number of Periods = 500

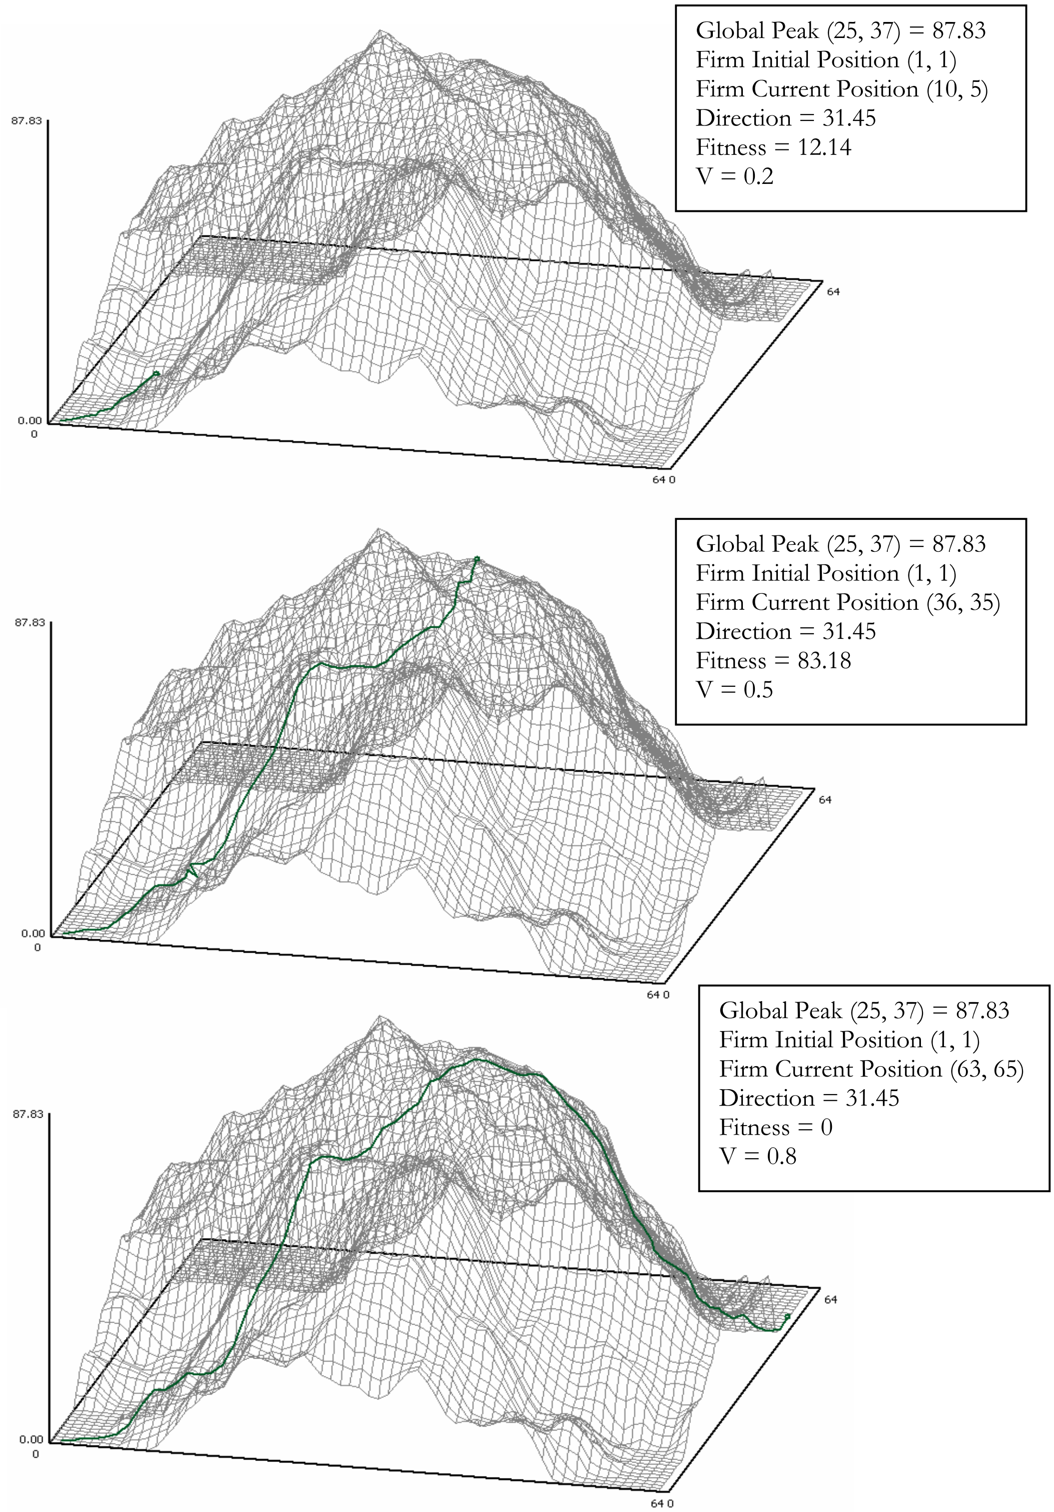

Figure 2: Searching a Smooth Landscape, Differing Values of v

h = 1.75, No Tilt, Number of Periods = 500

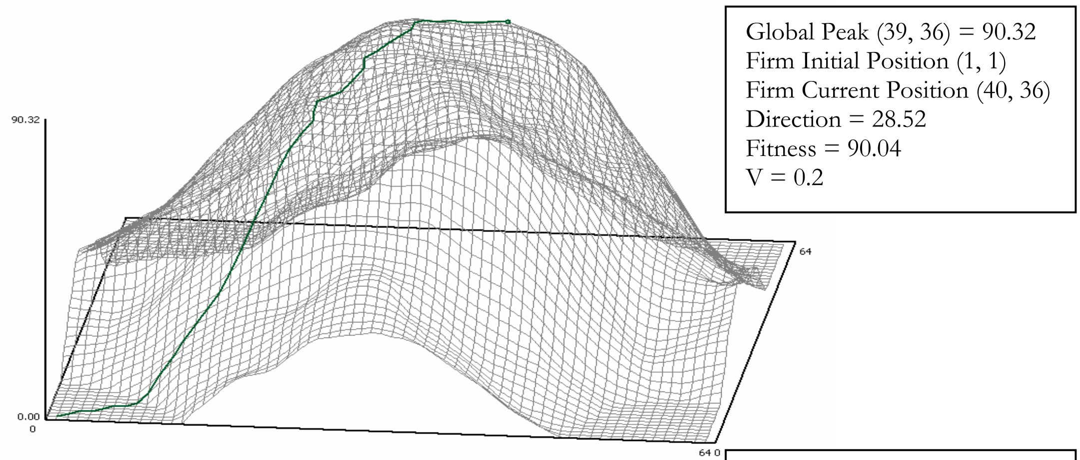

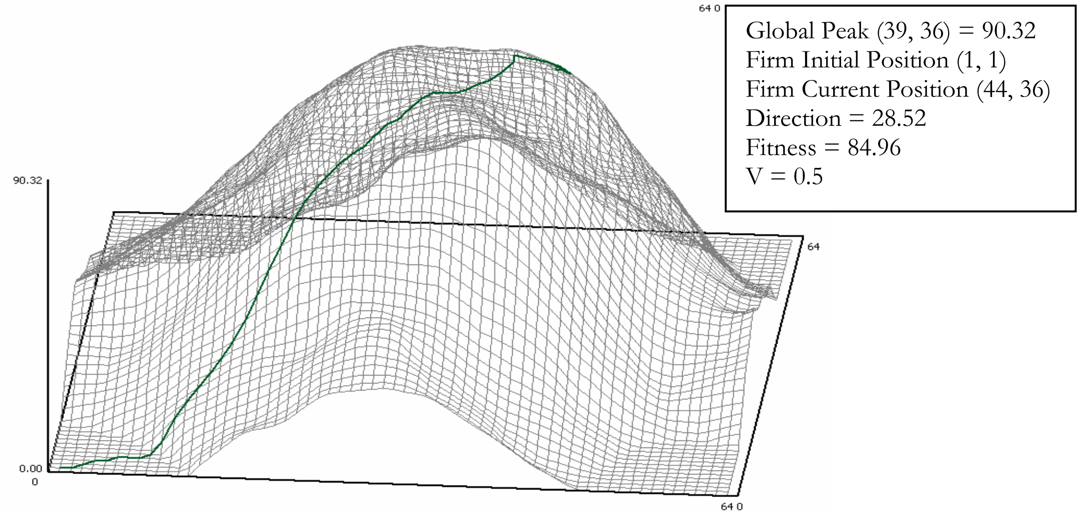

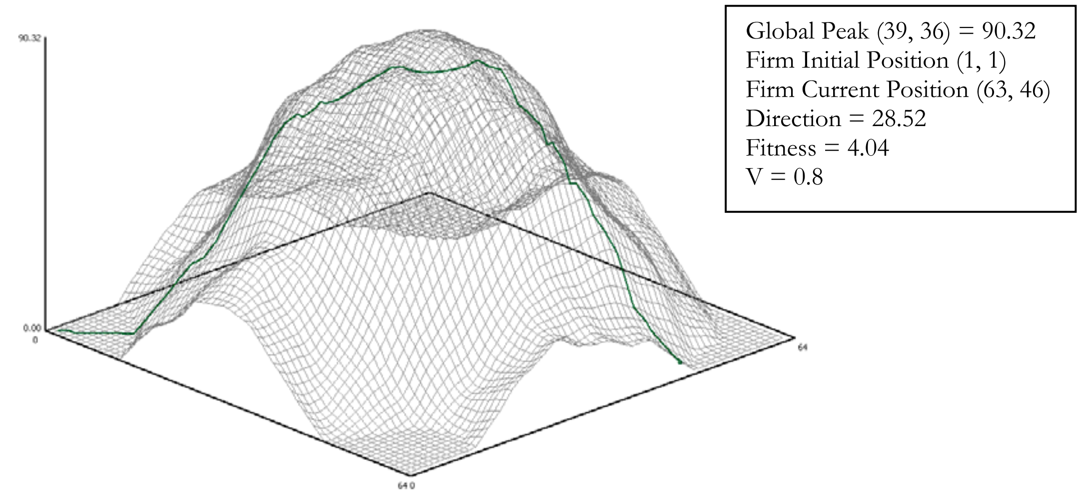

<u>Figure 3a</u>
Average Fitness, Period 250
(mean direction = 45, std dev = 22.5)

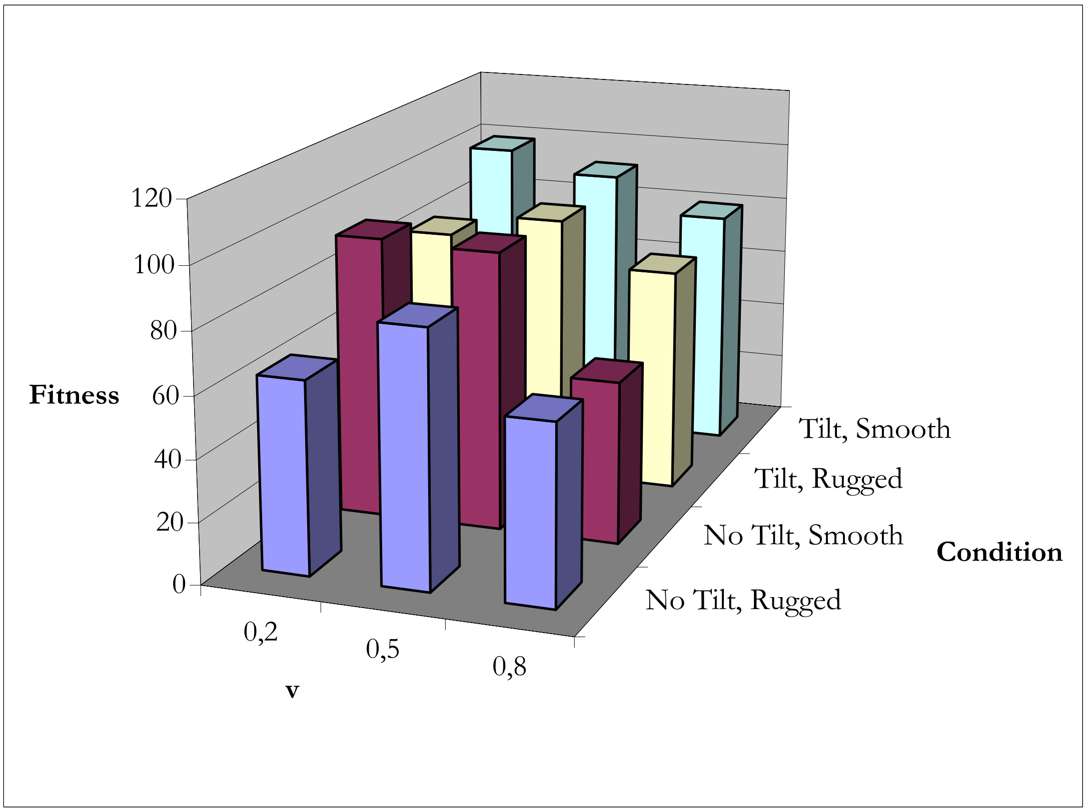

Figure 3b

Average Fitness, Period 500
(mean direction = 45, std dev = 22.5)

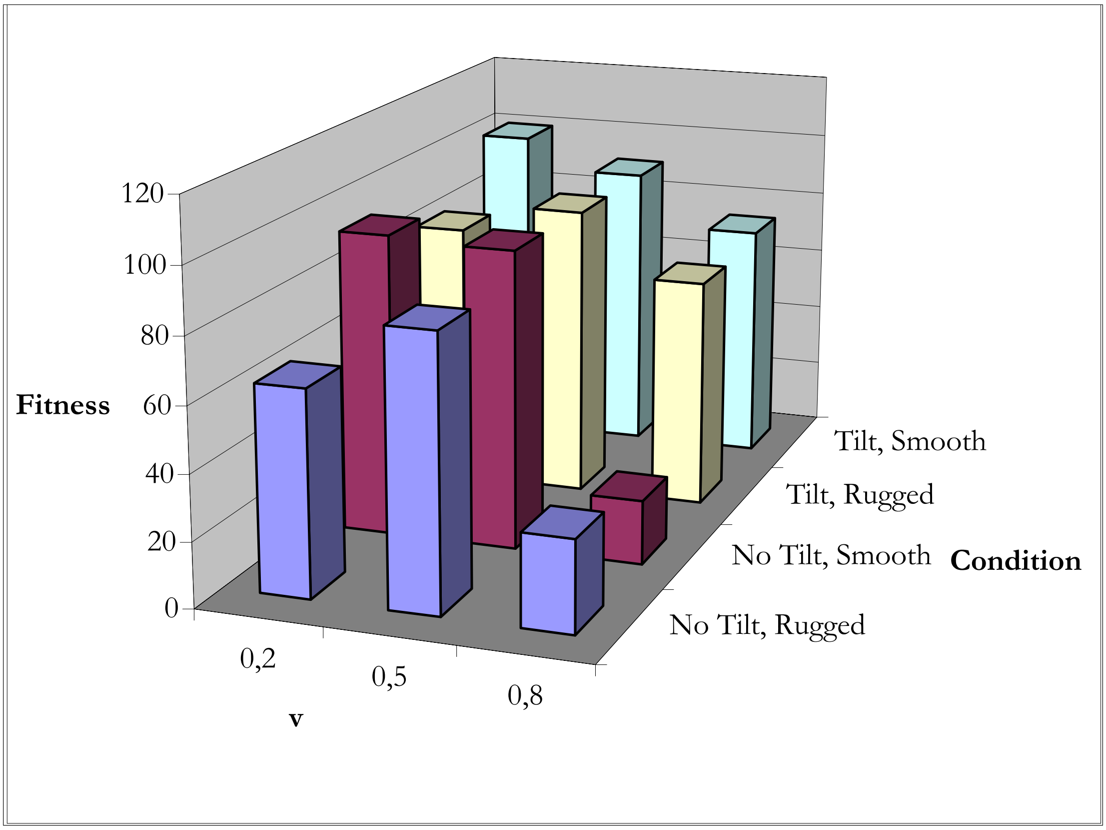

Appendix: Technical Aspects of Landscape Construction

As noted above, the landscape is constructed over a square grid that is 63 steps (64 points) on a side. The lowest frequency in the construction corresponds to a wavelength (1/frequency) equal to the size of the grid, 63 grid steps/cycle. The highest frequency corresponds to a wave length of a single grid step. The random element in the construction selects a weight for the amplitude contributed by a particular frequency from a standard normal distribution, which is then applied in a symmetric fashion to frequencies corresponding to wavelengths (x, y) and (N-x, N-y), where N is the grid size. This weight is then modified, however, to reflect the fractal aspect of the desired construction.[9] Specifically, the weight on frequency pair (u/N, w/N) is multiplied by the factor $r = (u * u + w * w)^{(-(h+1)/2)}$, which has the effect of diminishing the influence of higher frequencies to a degree that increases with the value of the parameter h. Thus, the lower the value of h, the higher the weights on the second and higher layers, and the more rugged the landscape. Layers beyond the first are also subject to a "random phase shift," uniformly distributed over 0 to 2π, which means that the positions of the curves at different frequencies are randomly distributed over the grid. The phase of the lowest frequency component is fixed by the "contour control."

Figure A 1 illustrates the logic above in the one dimensional case, with a two layer fractal structure. Panel 1a shows the fundamental building block, the cosine curve over the interval 0 to 4 π radians. "Contour control" corresponds to picking out the interval from 3 π to 5 π radians for attention, with the designated maximum point at 4 π. In Panel 1b, the two curves shown are the sum of the basic curve and a second curve that is reduced in amplitude and randomly shifted in phase. The amplitudes of the two second layers are at fractions f = 0.4 and f = 0.8 of the amplitude of the basic curve; the phase shifts are chosen to be the same to facilitate comparison. Note that in the f = 0.4 case, adding the second layer does not eliminate the unique maximum displayed in the first layer curve, though the shape is altered and the location of the maximum is shifted to the right. When the second layer is a larger fraction of the first (f = 0.8), there are two local maxima. In Panel 1c, the two curves are rescaled and the vertical origin is shifted, with sea level set a 0.2 of the distance from the minimum to the estimated maximum. Consistent with the procedure described above, we do not rescale the individual curves to achieve a maximum of precisely 100; the rescaling is at the family level and intended to achieve comparability between families.

<u>**Figure A 1**</u>

1a: Cosine Curve

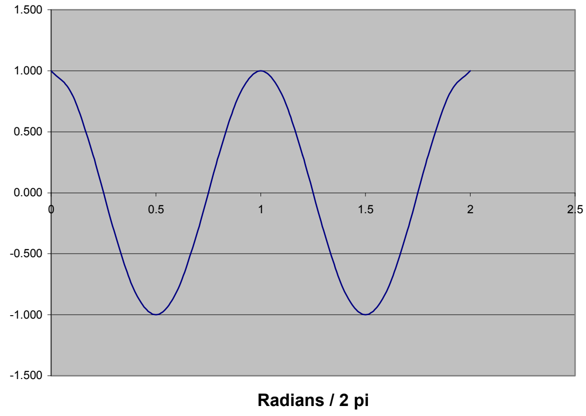

1b: Ilustrating Fractal Structure
f = .4 and f = .8

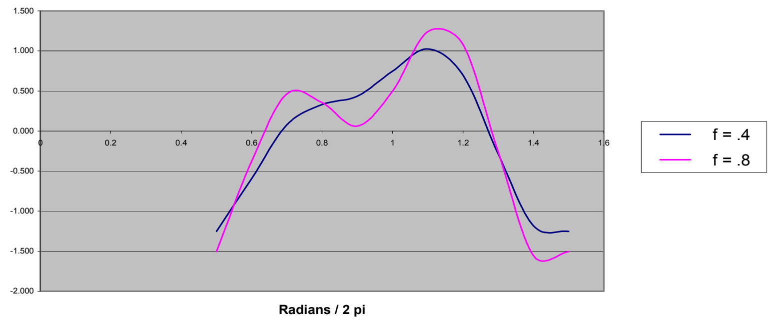

1c: Illustrating Rescaling

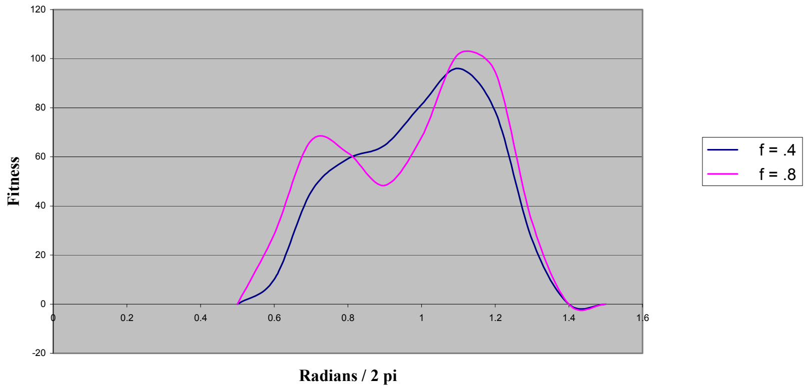

References

Cattani, G. (2005). "Pre-adaptation, firm heterogeneity and technological performance: A study on the evolution of fiber optics, 1970-1995." Organization Science **16**: 563-580.

Cattani, G. (2006). "Technological pre-adaptation, speciation and genesis of new technologies: How Corning invented and developed fiber optics." Industrial and Corporate Change, **15**(2): 285-318.

Cyert, R. M. and J. G. March (1963). A Behavioral Theory of the Firm. Englewood Cliffs, NJ, Prentice-Hall.

Fleming, L. and O. Sorenson (2001). "Technology as a complex adaptive system: Evidence from patent data." Research Policy **30**: 1019-1039.

Gavetti, G. and D. Levinthal (2000). "Looking forward and looking backward: Cognitive and experiential search." Administrative Science Quarterly **45**(1): 113-137.

Kauffman, S. (1993). The Origins of Order. New York, Oxford University Press.

Levinthal, D. (1997). "Adaptation on rugged landscapes." Management Science **43**: 934-950.

Levinthal, D. (1998). "The slow pace of rapid technological change: Gradualism and punctuation in technological change." Industrial and Corporate Change **7**(2): 217-247.

Levinthal, D. A. and J. G. March (1993). "The myopia of learning." Strategic Management Journal **14**(Special issue): 95-112.

Levitt, B. and J. G. March (1988). "Organizational learning." Annual Review of Sociology **14**: 319-340.

Lippman, S. A. and J. McCall (1976). "The economics of job search: A survey." Economic Inquiry **14**: 155-189.

Mandelbrot, B. B. (1983). The Fractal Geometry of Nature. New York, W.H. Freeman and Co.

March, J. G. (2003). Passion and Discipline: Don Quixote's Lessons for Leadership. Palo Alto, CA, Stanford University.

March, J. G. and H. A. Simon (1958). Organizations. New York, Wiley.

McKelvey, B. (1999). "Avoiding complexity catastrophe in co-evolutionary pockets: Strategies for rugged landscapes." Organization Science **10**(3): 294-321.

McKelvey, M. (1996). Evolutionary Innovations: The Business of Biotechnology. Oxford, Oxford University Press.

Nelson, R. R. (1982). "The role of knowledge in R&D efficiency." Quarterly Journal of Economics **97:** 453-470.

Nelson, R. R. (2006). "Bounded rationality, cognitive maps, and trial and error learning." Journal of Economic Behavior and Organization (forthcoming).

Peitgen, H.-O. and D. Saupe, Eds. (1988). The Science of Fractal Images. New York, Springer-Verlag.

Perrow, C. (1984). Normal Accidents: Living with High-Risk Technologies. New York, Basic Books.

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

Rivkin, J. W. (1998). Optimally local search on rugged landscapes. Boston, MA, Working Paper, Harvard Business School.

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

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

Rosenbloom, R. S. (2000). "Leadership, capabilities and technological change: The transformation of NCR in the electronic era." Strategic Management Journal **21**: 1083-1103.

Siggelkow, N. (2001). "Change in the presence of fit: The rise, the fall, and the renaissance of Liz Claiborne." Academy of Management Journal **44**: 838-857.

Starbuck, W. H. and F. J. Milliken (1988). "Challenger: Fine-tuning the odds until something breaks." Journal of Management Studies **25**: 319-340.

Tripsas, M. and G. Gavetti (2000). "Capabilities, cognition, and inertia: Evidence from digital imaging." Strategic Management Journal **21**: 1147-1161.

Winter, S. G. (2000). "The satisficing principle in capability learning." Strategic Management Journal **21**(Oct-Nov (special issue)): 981-996.

Wright, S. (1932). The roles of mutation, inbreeding, crossbreeding and selection in evolution. Proceedings of the Sixth International Congress of Genetics. **I:** 355-366.

Zenger, T. A. and J. A. Nickerson (2004). "A knowledge-based theory of the firm: The problem-solving perspective." Organization Science **15**: 617-632.

[1] The kinship of identity and irrational obsession at the individual level is a major theme of James March's remarkable video production, *Passion and Discipline: Don Quixote's Lessons for Leadership* (March 2003).
[2] In this work, Nelson returns to the development of themes that have concerned him in the past. (Nelson 1982) is a key background reference.
[3] While the quest in computer graphics research has been for techniques that generate realistic landscapes, such naturalness is not *per se* a key consideration in our research. But it is advantageous to the extent that it strengthens the heuristic advantages just referred to, in terms of the visual display and interpretation of the search processes modeled. We return subsequently to the more significant question of whether organizational knowledge might have a fractal structure.
[4] Otherwise it might happen, for example, that even v values as high as .99 still correspond to the "hill climbing" form of guidance.
[5] Thus our model addresses a more exploratory search process than the "problemistic" search featured in Cyert and March (1963). The concept of problemistic search is helpful in understanding how an organization gets back on track when something throws it off. It is less relevant to the process of finding the "track" (stable routines) in the first place, as in capability learning (Winter 2000).
[6] As Rivkin (2000) has shown, the complexity of a search in the NK model is logically linked to its computational tractability. For $K > 2$, the problem is "NP-complete," meaning that search length rises exponentially with N.
[7] Of course, our Fourier synthesis methods would exclude the possibility of a true "point" if our location variables were continuous. A (relatively) sharp local peak appears on our landscape when the amplitude weight on the shortest wavelength component is high.
[8] Such a change would bring the model into much closer alignment with Gavetti and Levinthal (2000), who also model cognition as summarizing weak information about remote points on the landscape.
[9] In the context of the Fourier synthesis method, this symmetry in weighting is required to zero out the imaginary terms in the construction, exploiting the fact that the conjugate pairs have imaginary parts that are equal in absolute value but opposite in sign.