Journal of Economic Behavior & Organization
Vol. 43 (2000) 141–166

JOURNAL OF
Economic Behavior
& Organization

www.elsevier.com/locate/econbase

Optimal search on a technology landscape

Stuart Kauffman [a], José Lobo [b,1], William G. Macready [a,*]

[a] *Bios Group LP, 317 Paseo de Peralta, Santa Fe, NM 87501, USA*
[b] *Graduate Field of Regional Science, Cornell University, 108 West Sibley Hall, Ithaca, NY 14853, USA*

Received 16 September 1998; received in revised form 31 August 1999; accepted 12 October 2000

Abstract

We address the question of how a firm's current location in the space of technological possibilities constrain its search for technological improvements. We formalize a quantitative notion of distance between technologies — encompassing the distinction between evolutionary changes (small distance) versus revolutionary change (large distance) — and introduce a *technology landscape* into an otherwise standard dynamic programming setting where the optimal strategy is to assign a reservation price to each possible technology. Technological search is modeled as movement, constrained by the cost of search, on a technology landscape. Simulations are presented on a stylized technology landscape while analytic results are derived using landscapes that are similar to Markov random fields. We find that early in the search for technological improvements, if the initial position is poor or average, it is optimal to search far away on the technology landscape; but as the firm succeeds in finding technological improvements it is optimal to confine search to a local region of the landscape. © 2000 Elsevier Science B.V. All rights reserved.

JEL classification: C61; C63; L20; O3

Keywords: Combinatorial optimization; Optimal search; Production recipes; Search distance; Technology landscape

1. Introduction

We address the question of how a firm's current production practices and its location in the space of technological possibilities constrain its search for technological improvements. We formalize a quantitative notion of technological distance which encompasses the distinction

* Corresponding author. Tel.: +1-505-992-6721.
E-mail addresses: stu@biosgroup.com (S. Kauffman), jl25@cornell.edu (J. Lobo), wgm@biosgroup.com (W.G. Macready).
[1] Tel.: +1-607-255-5385.

PII: S0167-2681(00)00114-1

between evolutionary change (small distance) versus revolutionary change (large distance). We are particularly interested in the relationship between the firm's current location in the space of technological possibilities and the *distance* at which the firm should search for technological improvements. Our formalization results in a *Landscape Search Model* (LSM) search based upon a *technology landscape* which extends traditional search theory. In the present discussion we focus on a detailed application of the LSM to optimal firm search.[2] Sufficient detail is provided so that readers unfamiliar with the LSM will find a self-contained treatment.

The starting point for our discussion is the representation of technology first presented in Auerswald and Lobo (1996) and Auerswald et al. (2000). In this framework a firm's production plan is more than a point in input–output space; it also includes the *production recipe* used in the process of production. A *configuration* denotes a specific assignment of states for every operation in the production recipe. A production recipe is comprised of N distinct *operations*, each of which can occupy one of S discrete states. The productivity of labor employed by a firm is a summation over the labor efficiency associated with each of the N production operations. The labor efficiency of any given operation is dependent on the state that it occupies, as well as the states of e other operations. The parameter e represents the magnitude of production externalities among the N operations comprising a production recipe, what we refer to as "intranalities". In the course of production during any given time period, the state of one or more operation is changed as a result either of spontaneous experimentation or strategic behavior. This change in the state of one or more operations of the firm's production recipe alters the firm's labor efficiency. The firm improves its labor efficiency — i.e., to say, the firm finds technological improvements — by searching over the space of possible configurations for its production recipe. When a firm finds a more efficient production recipe, it adopts that recipe in the next production period with certainty.

In order to explicitly consider the ways in which the firm's technological search is constrained by the firm's location in the search space, as well as the features of the space, we go beyond the standard search model (based on dynamic programming) and specify a *technology landscape.* The *distance metric* on the technology landscape is defined by the number of operations whose states need to be changed in order to turn one configuration into another. The firm's search for more efficient, production recipes is studied here as a "walk" on a *technology landscape*. The *cost of search* paid by the firm when sampling a new configuration is a nondecreasing function of the number of operations in the newly sampled configuration whose states differ from those in the currently utilized production recipe.

The literature on technology management and organizational behavior emphasize that although firms employ a wide range of search strategies, firms tend to engage in *local search* — i.e., search that enables firms to build upon their established technology (see, e.g., Lee and Allen, 1982; Sahal, 1985; Tushman and Anderson, 1986; Boeker, 1989; Henderson and Clark, 1990; Shan, 1990; Barney, 1991; Helfat, 1994). As discussed in March (1991) and Stuart and Podolny (1996), the prevalence of *local search* stems from the significant effort required for firms to achieve a certain level of technological competence, as well as from the greater risks and uncertainty faced by firms when they search for innovations far away from their current location in the space of technological possibilities. Using both numerical

[2] Further details and other applications of the LSM can be found in Lobo and Macready (1999).

and analytical results we relate the optimal search distance to the firm's initial productivity, the cost of search, and the correlation structure of the technology landscape. As a preview of our main result, we find that early in the search for technological improvements, if the firm's initial technological position is poor or average, it is optimal to search far away on the technology landscape. As the firm succeeds in finding technological improvements, however, it is indeed optimal to confine search to a local region of the technology landscape. We also obtain the familiar result that there are diminishing returns to search but without having to assume that the firm's repeated draws from the space of possible technologies are independent and identically distributed.

The outline of the paper is as follows. Section 2.1 presents a simple model of firm-level technology; production recipes are introduced in Section 2.1, production "intranalities" are defined in Section 2.2, and firm-level technological change is discussed in Section 2.3. The material in these sections draws heavily from Auerswald and Lobo (1996) and Auerswald et al. (2000) and further details can be found there. Section 3 develops the notion of a technology landscape, which is defined in Section 3.1. The correlation structure of the technology landscape is introduced in Section 3.2 as an important characteristic defining the landscape. Section 4 treats the firm's search for improved production recipes as movement on its technology landscape. The cost of this search is considered in Section 4.1. Section 4.2 then presents simulation results of search for the *Ne* technology landscape model defined in Section 2.2. We then go on to develop an analytically tractable model of technology landscapes in Section 5. We also describe in this section how a landscape can be represented by a probability distribution under an annealed approximation. Section 6 considers search under this formal model. The firm's search problem is formally defined in Section 6.1 and the important role of reservation prices is considered in Section 6.2. Section 6.3 determines the reservation price which determines optimal search and results are presented in Section 6.4. We conclude in Section 7 with a summary of results and some suggestions for further work.

2. Technology

2.1. Production recipes

A recent body of work, both empirical and theoretical, emphasizes the importance of firm specific characteristics for explaining technological change (for empirical contributions to this literature, see, e.g., Dunne, 1988, 1989; Audretsch, 1991, 1994; Davis and Haltiwanger, 1992; Bailey et al., 1994; Dwyer, 1995; Dunne et al., 1996; for theoretical contributions, see, e.g., Jovanovic, 1982; Herriott et al., 1985; Hopenhayn, 1992; Kennedy, 1994; Ericson and Pakes, 1995). Our representation of firm-level technology incorporates this perspective.

A firm using production recipe ω and labor input l_t produces q_t units of output during time period t,

$$q_t = F[\theta_t, l_t]. \tag{1}$$

The parameter θ represents a cardinal measure of the level of organizational capital associated with production recipe ω. The firm's level of organizational capital determines the firm's labor productivity (i.e., how much output is produced by a fixed amount of labor). Firm-level output is thus an increasing function of organizational capital, θ. A firm's level of organizational capital is a function of the *production recipe* utilized by the firm. The firm's production recipe encompasses all of the deliberate organizational and technical practices which, when performed together, result in the production of a specific good. (Our concept of organizational capital is very similar to that found in Presscott and Visscher (1980) and Hall (1993).) We assume, however, that production recipes as we define them are not fully known even to the firms which use them, much less to outsiders looking in. In order to allow for a possibly high-level of heterogeneity among production recipes utilized by different firms, we posit the existence of a set of all possible production recipes, Ω. We will refer to a single element $\omega_i \in \Omega$ as a production recipe. The efficiency mapping

$$\theta : \omega_i \in \Omega \to \Re^{++} \tag{2}$$

associates each production recipe with a unique labor efficiency.

Production recipes are assumed to involve a number of distinct and well-defined *operations*. Denote by N the number of operations in the firm's production recipe, which is determined by engineering considerations. The ith recipe ω_i can then be represented by

$$\omega_i = \{\omega_i^1, \dots, \omega_i^j, \dots, \omega_i^N\}, \tag{3}$$

where ω_i^j is the description of operation j for $j = 1, \dots, N$. We assume that the operations comprising a production recipe can be characterized by a set of *discrete* choices. These discrete choices may represent either qualitative choices (e.g., whether to use a conveyor belt or a forklift for internal transport), quantitative choices (e.g., the setting of a knob on a machine), or a mixture of both. In particular we assume that

$$\omega_i^j \in \{1, \dots, S\} \tag{4}$$

for each $i \in \{1, \dots, N\}$ and where S is a positive integer. Each operation ω_i^j of the production recipe ω_i can thus occupy one of S states.

We denote a specific assignment of states to each operation in a production recipe as a *configuration*. Making the simplifying assumption that the number of possible states is the same for all operations that comprise a given production recipe, the number of all possible and distinct configurations for a given production recipe associated with a specific good is equal to

$$|\Omega| = S^N. \tag{5}$$

New production processes are created by altering the states of the operations which comprise a production recipe. Technological change in this framework takes the form of finding production recipes which maximize labor efficiency per unit of output (i.e., technological progress is Harrod-neutral).

The contribution to overall labor efficiency made by the jth operation depends on the setting or state chosen for that operation, ω_i^j, and possibly on the settings chosen for all other operations, $\omega_i^{-j} \equiv \{\omega_i^1, \dots, \omega_i^{j-1}, \omega_i^{j+1}, \dots, \omega_i^N\}$. Hence the labor efficiency of the jth operation is in general a function ϕ_i^j of ω_i^j and ω_i^{-j}, so that we can write

$$\phi_i^j = \phi^j(\omega_i^j, \omega_i^{-j}). \tag{6}$$

We assume that the N distinct operations that comprise the production recipe contribute additively to the firm's labor efficiency

$$\theta(\omega_i) = \frac{1}{N}\sum_{j=1}^{N}\phi_i^j = \frac{1}{N}\sum_{j=1}^{N}\phi^j(\omega_i^j, \omega_i^{-j}). \tag{7}$$

We can think of $\phi^j(\omega_i^j, \omega_i^{-j})$ as the payoff to the jth operating unit when it is in state ω_i^j and the other operations are in the states encoded by the vector ω_i^{-j}. In our cooperative setting, operations act not to maximize their own labor efficiency, but rather the aggregate labor productivity of the firm (i.e., $\theta(\omega_i)$).

2.2. *Production intranalities*

Working from the view that an important role of the firm is to "internalize" externalities (Coase, 1937; Williamson, 1985), we assume that in the typical case there are significant external economies and diseconomies among the N operations comprising a production recipe — i.e. to say, significant production and management externalities exist *within* the firm. These "intranalities" can be thought of as connections between the operations constituting the production recipe (Reiter and Sherman, 1962). To say that a connection exists between two operations is simply to say that the performance of the two operations affect each other (positively or negatively) either bilaterally or unidirectionally.

For a production recipe $\omega \in \Omega$, we define the *production intranality* scalar

$$e_k^j = e(\omega^k, \omega^j) \tag{8}$$

as follows:

$$e_k^j = \begin{cases} 1, & \text{if the setting of operation } j \text{ affects the labor requirement of operation } k, \\ 0, & \text{otherwise} \end{cases} \tag{9}$$

for $j, k = 1, \dots, N$. Since the choice of the setting for the jth operation always affects the efficiency for the jth operation, we have $e_j^j = 1$ for $j = 1, \dots, N$. We make the strong simplifying assumption of equal number of connections, namely

$$e_k^j = e_j^k \equiv e. \tag{10}$$

We assume throughout that e and N are given by nature.

When $e = 1$, Eq. (8) is additively separable, otherwise

$$\theta(\omega_i) = \frac{1}{N}\sum_{j=1}^{N}\phi_i^j = \frac{1}{N}\sum_{j=1}^{N}\phi^j(\omega_i^j; \omega_i^{j_1}, \dots, \omega_i^{j_e}). \tag{11}$$

The S^e possible contributions to total labor efficiency made by the jth operation (through ϕ_i^j) are treated as i.i.d. random variables drawn from some distribution F. In what follows we assume that the values returned by ϕ are drawn from the uniform distribution, $U(0, 1)$, over the unit interval, although our results are insensitive to this choice.

2.3. Firm-level technological change

We now describe the general features of the technological problem facing the firm in our model. The firm's production recipe determines the firm's level of organizational capital and thus its labor efficiency. The production recipe is comprised of a number of distinct operations which at each moment can be in one of a finite number of possible, and discrete, states. Consequently, improvements in the technology used by the firm entails changes in the state of the operations comprising the production recipe. [3] Firm-level technological improvements result from the firm finding improved configurations for its production recipe. Thus stated, the firm's technological problem is a combinatorial optimization problem (Reiter and Sherman, 1965; Papadimitriou and Steiglitz, 1982; Cameron, 1994). A compelling question to ask in the context of combinatorial optimization is, whether the globally optimal configuration can be reached from any given initial configuration.

We propose to study the firm's technological problem by means of a *technology landscape,* developed in Section 3. The economies or diseconomies resulting from the interaction among the operations constituting the production recipe constraintgreatly affect the firm's search for technological improvements. The intranalities parameter e provides a measure of the conflicting constraints confronting the firm as it seeks to optimize its production recipe. Just as a topographical map is a way of representing height over a two-dimensional physical space, a technology landscape is a means of representing the problem faced by the firm in its search for the optimal configuration for its production recipe.

3. The technology landscape

3.1. Defining the technology landscape

To define a technology landscape we require a measure of distance between two different production recipes, ω_i and ω_j, each drawn from Ω. The distance metric used here is not based on the relative efficiencies of production recipes, but rather on the similarity between the operations constituting the recipes. More precisely, the *distance* $d(\omega_i, \omega_j)$ between the production recipes ω_i and ω_j is the *minimum* number of operations which must be changed in order to convert ω_i to ω_j. Since changing operations is symmetric $d(\omega_i, \omega_j) = d(\omega_j, \omega_i)$. Given this distance metric, we can define the set of "neighbors" for any production recipe,

$$\mathcal{N}_d(\omega_i) = \{\omega_j \in \{\Omega - \omega_i\} : d(\omega_i, \omega_j) = d\}, \tag{12}$$

where $\mathcal{N}_d(\omega_i)$ denotes the set of d-*neighbors* of recipe ω_i and $d \in \{0, \dots, N\}$.

With this definition of distance between recipes, it is straightforward to construct the *technological graph*, $\Gamma(V, E)$. The set of nodes or vertices of the graph, V, are the production recipes $\omega_i \in \Omega$. The set of *edges* of the technological graph, E, connect any given

[3] Our view of technological innovation is similar to that of Romer (1990), who remarks that over the past few 100 years, "the raw materials that we use have not changed, but as a result of trial and error, experimentation, refinement and scientific investigation, the instructions that we follow for combining raw materials have become vastly more sophisticated". Our "production recipes" are directly analogous to Romer's "instructions".

recipe to its $d=1$ *neighbors,* i.e., to the elements of $\mathcal{N}_1(\omega_i)$. For any production recipe, the number of one-operation variant neighbors is given by

$$|\mathcal{N}_1(\omega_i)| = (S-1)N \quad \text{for all } \omega_i \in \Omega. \tag{13}$$

Thus each node of Γ is connected to $(S-1)N$ other nodes.

The technology graph Γ and the efficiency map $\theta : \Omega \rightarrow \mathfrak{R}$ (efficiencies can be associated with each node in Γ) constitute a technology landscape. (For a comprehensive discussion of landscape models, see Stadler, 1995.) Assume, for the moment, that the labor efficiencies $\theta(\omega_i)$ are known with certainty for each $\omega_i \in \Omega$. Adopting some method for tie-breaking, we can orient the edges of the graph Γ from vertices with higher labor efficiencies toward vertices associated with lower labor efficiencies. The firm's search problem can then be recast as that of "moving" in the technology landscape (varying the production recipe by changing the state of at least one operation) in order to maximize θ. The "steps" constituting such a walk represent the adoption, by the firm, of the sampled variants for its production recipe.

In the more general (and interesting) case where the efficiency $\theta(\omega)$ associated with each production recipe is not known with certainty, a *random field*, $\mathcal{F}$, can be defined over the production recipes $\omega \in \Omega$ by the joint probability distribution

$$F(\theta_1, \dots, \theta_{S^N}) = \text{Prob}\{\theta(\omega_i) \le \theta_i \quad \text{for } i = 1, \dots, S^N\}, \tag{14}$$

where $\theta(\omega_i) \in \Gamma$ is the labor requirement at vertex i, each θ_i is a positive scalar and S^N is the total number of vertices (i.e., of production recipes).[4] The joint probability distribution in Eq. (14) induces a probability measure μ on $(\theta_1, \dots, \theta_{S^N})$. The mapping implicit in Eq. (2), along with the measure μ, forms a probability space which is a random field on Γ, the technological graph (see Macken and Stadler, 1995). In general then, a technology landscape is a realization of $\mathcal{F}$ (Stadler and Happel, 1995).

3.2. Correlation structure of the technology landscape

Perhaps one of the most important properties of a technology landscape is its *correlation structure*. The correlation of a landscape measures the degree to which nearby locations on the landscape have similar labor efficiencies. A straightforward way to measure the correlation of a landscape is by means of the correlation function

$$\rho(d) = \frac{E(\theta(\omega_i)\theta(\omega_j)|d) - E(\theta(\omega_i))E(\theta(\omega_j))}{\sigma(\theta(\omega_i))\sigma(\theta(\omega_j))}, \tag{15}$$

where $\sigma(\theta)$ is the standard deviation of efficiencies and $\rho(d)$ is the landscape's correlation coefficient for efficiencies corresponding to production recipes ω_i and ω_j which are a distance d apart (Eigen et al., 1989). The expectation $E(\theta(\omega_i)\theta(\omega_j)|d)$ is with respect to the probability distribution $P(\theta(\omega_i), \theta(\omega_j)|d)$ which will be defined later in Section 5.

In the case of an *Ne* technology landscape the level of intranalities characterizing a production recipe induces the *correlation structure* on the landscape. To see this, consider

[4] For a general introduction to random fields, see Griffeath (1976) or Vanmarcke (1983).

the limiting case of a production method characterized by $e = 1$. In this case the contribution made by each operation to overall production cost is independent of the states of the other operations since the contribution to total efficiency made by each operation depends only on the state of that operation. Whether or not each operation in the production recipe makes its highest possible contribution to total efficiency depends in turn on whether or not the operation in question occupies its optimal state. Therefore, there exists a single globally optimal configuration for the firm's production method under which each operation occupies its optimal state. Any other configuration, which must necessarily have lower efficiency, can be sequentially changed to the globally optimal configuration by successively changing the state of each operation. Furthermore, any such suboptimal recipe lies on a connected pathway via more efficient one-operation variants to the single global optimum in the landscape. Given the additive specification for production efficiency, a transition to a one-operation variant neighbor of ω_i (i.e., changing the state of one operation) typically alters the efficiency of the production method by an amount $\mathcal{O}(1/N)$. When $e = 1$, production methods a distance $d = 1$ away in the efficiency landscape therefore have nearly the same efficiency. Consequently, production methods in an $e = 1$ landscape are tightly correlated in their production efficiencies.

In contrast, in the $e = N$ limit the contribution made by each operation to the efficiency of the production method depends on the state of all other operations. The contribution made by each operation is changed when even a single operation is altered. Consider any initial production method among the S^N possible recipes. Alteration of one of the production method's operations alters the combination of the $e = N$ operations that bear on the efficiency of each operation. In turn, this alteration changes the efficiency of each operation to a randomly chosen value from the appropriate distribution. The total production efficiency of the new production recipe is therefore a sum of N new random variables, from which it follows that the new efficiency is entirely uncorrelated with the old efficiency. The efficiency of any given production method is therefore uncorrelated with the efficiencies of its nearest neighbors.

Following Weinberger (1990) and Fontana et al. (1993), we can formally define a *correlation coefficient* for an *Ne* landscape. Suppose the firm moves from production method ω to production method ω', a distance d apart. Let $P(d)$ be the probability for any given operation to be among the d operations that are changed by moving from ω to ω'. The autocorrelation coefficient, $\rho(d)$, for two production methods a distance d apart is then given by

$$\rho(d) = 1 - P(d). \tag{16}$$

The efficiency of an operation is unchanged if it is not one of the d operations that have been changed as the firm moved from ω to ω', and if it is not one of the e neighbors of any of the changed operations. These two events are statistically independent, and thus

$$P(d) = 1 - \left[1 - \frac{d}{N}\right]\left[1 - \frac{e-1}{N-1}\right]^{d}, \tag{17}$$

from which it follows that

$$\rho(d) = \left[1 - \frac{d}{N}\right]\left[1 - \frac{e-1}{N-1}\right]^{d}. \tag{18}$$

When $d = 1$ and there are no production externalities ($e = 1$), $\rho(1) \approx 1$ for large N and when every operation affects every other operation $(e = N)$, $\rho(N) = 0$.

As e increases, the landscape goes from being "smooth" and single peaked to being "rugged" and fully random. For low values of e the correlation spans the entire configuration space and the space is thus nonisotropic. As e increases, the configuration space breaks up into statistically equivalent regions, so the space as a whole becomes isotropic (Kauffman, 1993).

A related measure of landscape correlation, and one which can be used to compare landscapes, is the *correlation length*. The correlation length, l, of a technology landscape is defined by

$$l^{-1} = \sum_{d \geq 0} \rho(d). \tag{19}$$

For a correlation coefficient which decays exponentially with distance, the correlation length is the distance over which the correlation falls to $1/e$ of its initial value. For the *Ne* technology landscape

$$l = -\frac{1}{\ln \rho}. \tag{20}$$

4. Search on the technology landscape

4.1. Search cost

The firm's walk on a technology landscape is similar to a random search within a fixed population of possibilities (Stone, 1975). [5] In the model presented here the firm seeks technological improvements by sampling d-variants of its currently utilized production recipe. It does this by selecting independent drawings from some distribution F, at a sampling cost of c per drawing (where $c > 0$). The firm's search rule is fairly simple. Consider a firm that is currently utilizing production recipe ω_i and whose labor efficiency is therefore $\theta(\omega_i)$. The firm can take either of two actions: (1) keep using production recipe ω_i, or (2) bear an additional search cost c and sample a new production recipe $\omega_j \in \mathcal{N}_d$ from the technology landscape. The decision rule followed by the firm is to change production recipes when an efficiency improvement is found, but otherwise keep the same recipe. Let θ_i be the efficiency of the production recipe currently used by the firm, and let θ_j be the efficiency of a newly sampled production recipe; if $\theta_j > \theta_i$, the firm adopts $\omega_j \in \Omega$ in the next time period; if $\theta_j \leq \theta_i$, the firm keeps using ω_i. This search rule is in effect an "uphill walk" on the landscape, with each step taken by the firm taking it to a d-operation variant of the firm's current production recipe.

The actual procedures used by the firm when searching for technological improvements can range from the non-intentional (e.g, "learning by doing"), to the strategic (investments

[5] Technological change has often been modeled by economists as a random search within a fixed population of possibilities (see, e.g., Evenson and Kislev, 1976; Weitzman, 1979; Levinthal and March, 1981; Hey, 1982; Nelson and Winter, 1982; Tesler, 1982; Muth, 1986; Cohen and Levinthal, 1989; Jovanovic and Rob, 1990; Marengo, 1992; Adams and Sveikauskas, 1993).

in R&D); technological improvements can result from small scale innovations occurring in the shop-floor or from discoveries originating in a laboratory. The level of sophistication of the firm's search for new technologies is mapped into how many of the operations comprising the currently used production recipe have their states changed as the firm moves on its technology landscape. Production recipes sampled at large distances represent very different production processes while production processes separated by small distances represent similar processes. Improved variants found at large distances from the current recipe represent wholesale changes whereas nearby improved variants constitute refinements rather than large scale alterations.

The many issues of industrial organizational, quality control, managerial intervention and allocation of scarce research resources involved in firm-level technological change are here collapsed into the cost, c, which the firm must pay in order to sample from the space of possible configurations for its production recipe. We assume the unit cost of sampling to be a nondecreasing function of how far away from its current production recipe the firm searches for an improved configuration — recalling that in the metric used here the distance between two configurations in the technology landscape is the number of operations which must be changed in order to turn one production recipe into the other. For present purposes it suffices to have the relationship between search cost and search distance to be a simple linear function of distances

$$c = \alpha d, \tag{21}$$

where $\alpha \in \Re^{+}$ and $1 \leq d \leq N$ is the distance between the currently utilized production recipe, ω_i, and the newly sampled production recipe, ω_j.

4.2. Search distance

At what distance away from its current production recipe should the firm search for technological improvements? In the most "naive" form of search on a technology landscape the firm restricts itself to myopically sampling among nearby variants in order to climb to a local optimum. Might it be better for the firm to search further away? The answer is "yes", but the optimal search distance typically decreases as the labor efficiency of the firm's current production recipe increases.

Consider an *Ne* technology landscape with a moderately long correlation length and suppose that a firm starts production with a production recipe of average efficiency 0.5 (recall that for the *Ne* efficiencies lie between 0 and 1). Then half of the 1-operation variant neighbors of the initial production recipe are expected to have a lower labor efficiency, and half are expected to have higher efficiency. More generally, half of the production recipe variants at any distance $d = 1, \ldots, N$ away from the initial configuration should be more efficient and half should be less efficient. Since the technology landscape is correlated, however, nearby variants of the initial production recipe, those a distance 1 or 2 away, are constrained by the correlation structure of the landscape to be only slightly more or less labor efficient than the starting configuration. In contrast, variants sampled at a distance well beyond the correlation length, l, of the landscape can have efficiencies very much higher or lower than that of the initial production recipe.

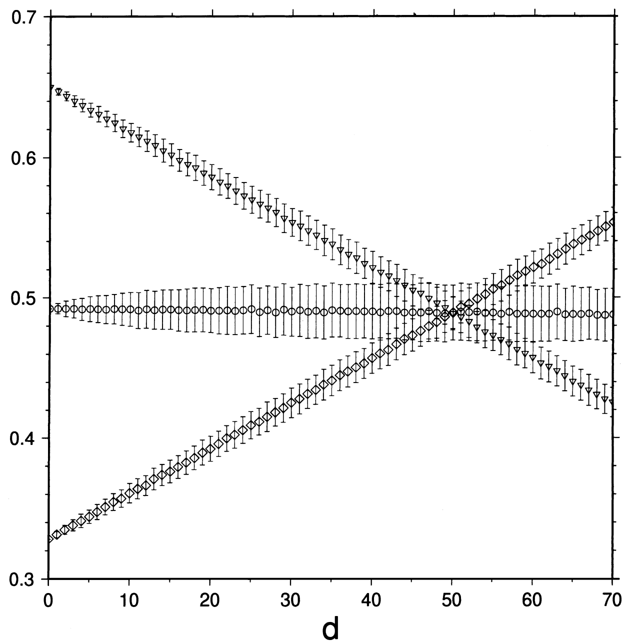

Fig. 1. Mean labor efficiencies ± one standard deviation versus search distance for $N = 100$, $e = 1$, and three different initial efficiencies.

It thus seems plausible to suppose that, early in the firm's search process from a poor or even average initial configuration, the more efficient variants will be found most readily by searching far away on the technology landscape. But as the labor efficiency increases, distant variants are likely to be nearly average in the space of possible efficiencies — hence less efficient — while nearby variants are likely to have efficiencies similar to that of the current, highly efficient, configuration. Thus, distant search will almost certainly fail to find more efficient variants, and search is better confined to the local region of the space.

Figs. 1–3 show the results of simulations exploring this intuition for a technology landscape with $N = 100$, varying e values, $S = 2$ and three different starting labor requirements (near 0.35, 0.50, and 0.70). The e intranalities are assigned at random from any of the other $N - 1$ operations. The number of operations e^j affected by the jth operation is binomially distributed. The labor efficiency ϕ^j of the jth operation is assigned randomly from the uniform distribution $U(0, 1)$. The total labor requirement of a production recipe thus varies from 0 to 1, and for N large enough has a Gaussian distribution with mean 0.5. From each initial position, 5000 variants were sampled at each search distance $d = 1, \dots, 100$. Since $N = 100$, a distance of, e.g., $d = 70$, corresponds to changing the state of 70 of the 100 operations in the binary string representing the firm's current position on the technology landscape. Each set of 5000 samples at each distance yielded a roughly Gaussian distribution of labor requirements encountered at that search distance. Figs. 1–3 show, at each distance, a bar terminating at one standard deviation above and one standard deviation

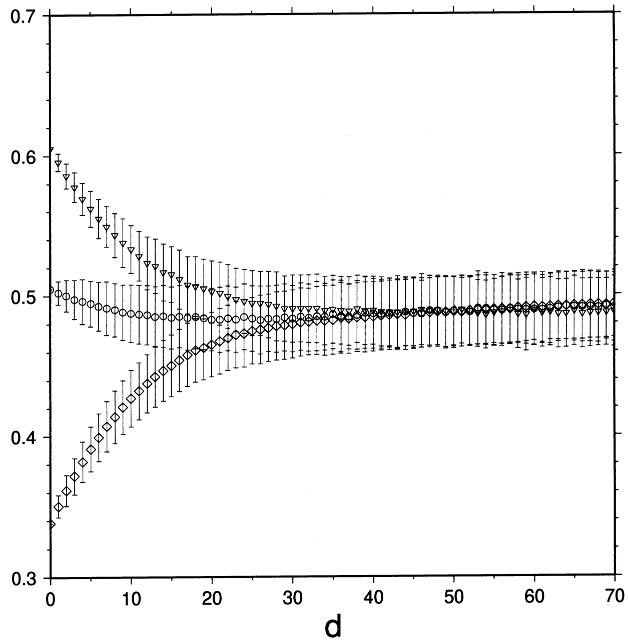

Fig. 2. Mean labor efficiencies ± one standard deviation versus search distance for $N = 100$, $e = 5$, and three different initial efficiencies.

below the mean labor requirement found at that distance. Roughly one-sixth of a Gaussian distribution lies above one standard deviation. Thus, if six samples had been taken at each distance, and the "best" of the six chosen, then the expected increase in labor efficiency at each distance is represented by the envelope following the "plus" one standard deviation marks at each distance.

Fig. 1 shows that when $e = 1$ and the initial labor efficiency is near 0.5, the optimal search distance with six samples occurs when around 50 of the 100 operations are altered. When the initial labor efficiency is high, however, the optimal search distance dwindles to the immediate vicinity of the starting configuration. In contrast, when the initial labor efficiency is much lower than the mean, it is optimal for the firm to "jump" (i.e., search far away) instead of "walk" (i.e., search nearby) across the technology landscape. For Fig. 2, where $e = 5$, the correlation length is shorter and as a result the optimal search distance for initial efficiencies near 0.5 is smaller (in this case around $d = 5$). It is still the case that for highly efficient initial recipes, search should be confined to the immediate neighborhood. Very poor initial efficiencies still benefit most from distant search. In Fig. 3, where $e = 11$, the correlation length of the technology landscape is shorter still and optimal search distances shrink further.

The numerical results suggest that on a technology landscape it is optimal to search far away when labor efficiency is low in order to sample beyond the correlation length of the configuration space. As labor efficiency increases, however, optimal search is confined closer to home. These numerical results are intuitively appealing and even common sensical.

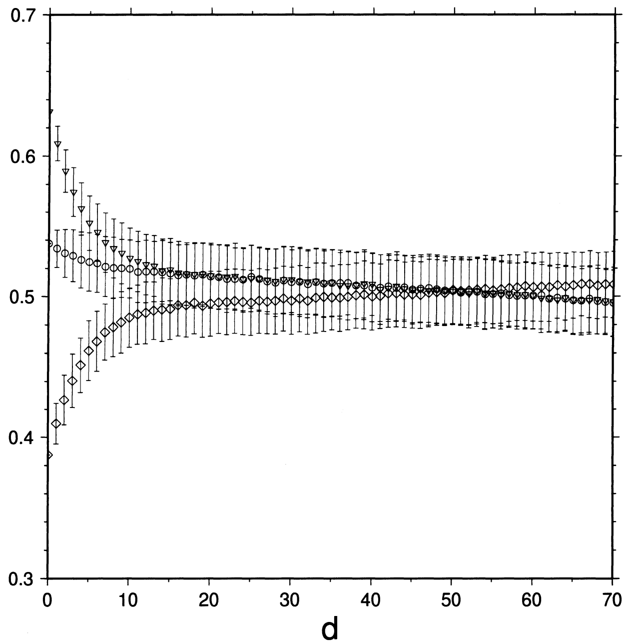

Fig. 3. Mean labor efficiencies ± one standard deviation versus search distance for $N = 100$, $e = 11$, and three different initial efficiencies.

It is appropriate to ask, however, how closely tied are the results to the adoption of a specific technology landscape model. In the next two sections we provide a framework with which to analytically consider the question of optimal search distance not just for an *Ne* landscape but any landscape model. Section 5 outlines a formal framework with which to treat landscapes while Section 6 places search cost within a standard dynamic programming context.

5. Analytic approximation for the distribution of efficiencies

Technology landscapes are very complex entities, characterized by a neighborhood graph Γ and an exponential number of labor efficiencies S^N. In any formal description of technology landscapes we have little hope of treating all of these details. Consequently we adopt a probabilistic approach focusing on the statistical regularities of the landscape and which is applicable to any landscape model.

To treat the technology landscape statistically we follow Macready (1999) and assume that the landscape can be represented using an *annealed approximation*. The annealed approximation (Derrida and Pomeau, 1986) is often used to study systems with disorder (i.e., randomly assigned properties) as is the case with our *Ne* model. (Recall that the labor efficiencies ϕ_i^j are assigned by random sampling uniformly from the interval [0, 1].) In evaluating the statistical properties of the *Ne* landscape one must first sample an entire technology landscape and then measure some property on *that* landscape. Repeated sampling and measuring on many landscapes then yields the desired aggregate statistics. Analytically mimicking this process is difficult, however, because averaging over the landscapes is the final step in the calculation and usually results in an intractable integration. In our annealed approximation the averaging over landscapes is done *before* measuring the desired statistic, resulting in vastly simpler calculations. The annealed approximation will be sufficiently accurate for our purposes and we shall comment on the range of its validity.

As an example of our annealed approximation, let us assume we want to measure the average of a product of four efficiencies along a connected walk in Γ. We label the efficiencies θ_1, θ_2, θ_3, and θ_4. If $P(\theta_1, \dots, \theta_{S^N})$ is the probability distribution for an entire technology landscape this average is calculated as

$$\int \theta_1\theta_2\theta_3\theta_4\, P(\theta_1, \dots, \theta_{S^N})\, \mathrm{d}\theta_1 \cdots d\theta_{S^N} = \int P(\theta_1, \theta_2, \theta_3, \theta_4)\, \theta_1\theta_2\theta_3\theta_4\, \mathrm{d}\theta_1\, \mathrm{d}\theta_2\, \mathrm{d}\theta_3\, \mathrm{d}\theta_4. \tag{22}$$

This integral may be difficult to evaluate depending on the form of $P(\theta_1, \theta_2, \theta_3, \theta_4)$. Under the annealed approximation this integral is instead evaluated as

$$\int P(\theta_1)\theta_1 P(\theta_2|\theta_1)\theta_2 P(\theta_3|\theta_2)\theta_3 P(\theta_4|\theta_3)\theta_4\, \mathrm{d}\theta_1\, \mathrm{d}\theta_2\, \mathrm{d}\theta_3\, \mathrm{d}\theta_4, \tag{23}$$

where $P(\theta|\theta')$ is the probability that a configuration has labor efficiency θ conditioned on the fact that a neighboring configuration has efficiency θ'.

As we have seen, under our annealed approximation the entire landscape is replaced by the joint probability distribution $P(\theta(\omega_i), \theta(\omega_j))$, where production recipes ω_i and ω_j are a distance 1 apart in Γ. For any particular technology landscape the probability that the efficiencies of a randomly chosen pair of configurations a distance d apart have efficiencies θ and θ' is

$$P(\theta, \theta'|d) = \frac{\sum_{\langle\omega_i,\omega_j\rangle_d} \delta(\theta - \theta(\omega_i))\delta(\theta' - \theta(\omega_j))}{\sum_{\langle\omega_i,\omega_j\rangle_d} 1}, \tag{24}$$

where the notation $\langle\omega_i, \omega_j\rangle_d$ requires that production recipes ω_i and ω_j are a distance d apart and δ is the Dirac delta function.[6] Rather than work with the full $P(\theta, \theta'|d)$ we simplify and consider only

$$P(\theta(\omega_i), \theta(\omega_j)) \equiv P(\theta(\omega_i), \theta(\omega_j)|d = 1). \tag{25}$$

For some technology landscape properties we might need the full $P(\theta(\omega_i), \theta(\omega_j)|d)$ distribution but we will approximate it by building up from $P(\theta(\omega_i), \theta(\omega_j))$. More accurate extensions of this annealed approximation may be obtained if $P(\theta(\omega_i), \theta(\omega_j)|d)$ is known.

From $P(\theta(\omega_i), \theta(\omega_j))$ we can calculate both $P(\theta(\omega_i))$, the probability of a randomly chosen production recipe ω_i having efficiency $\theta(\omega_i)$, and $P(\theta(\omega_i)|\theta(\omega_j))$, the probability

[6] The Dirac delta function is the continuous analog of the Kronecker delta function: $\delta(x)$ is zero unless $x = 0$ and is defined so that $\int_I \delta(x)\, \mathrm{d}x = 1$ if the region of integration, I, includes zero.

of a production recipe ω_i having labor efficiency $\theta(\omega_i)$ given that a neighboring production recipe ω_j has labor efficiency $\theta(\omega_j)$. Formally these probabilities are defined as

$$P(\theta(\omega_i)) = \int_{-\infty}^{\infty} P(\theta(\omega_i), \theta(\omega_j))\, \mathrm{d}\theta(\omega_j), \tag{26}$$

and

$$P(\theta(\omega_i)|\theta(\omega_j)) = \frac{P(\theta(\omega_i), \theta(\omega_j))}{P(\theta(\omega_j))}. \tag{27}$$

Note that we have assumed, for mathematical convenience, that labor efficiencies range over the entire real line. While efficiencies are no longer bounded from below, the ordering relationship amongst efficiencies is preserved and extreme labor efficiencies are very unlikely.

For *Ne* landscapes the following probability densities may be calculated exactly as:

$$P(\theta(\omega_i)) = \frac{1}{\sqrt{2\pi}} \exp\left[-\frac{\theta^2(\omega_i)}{2}\right], \tag{28}$$

$$P(\theta(\omega_i), \theta(\omega_j)) = \frac{1}{2\pi\sqrt{1-\rho^2}} \exp\left[-\frac{\theta^2(\omega_i) + \theta^2(\omega_j) - 2\rho\theta(\omega_i)\theta(\omega_j)}{2(1-\rho^2)}\right], \tag{29}$$

$$P(\theta(\omega_i)|\theta(\omega_j)) = \frac{1}{\sqrt{2\pi(1-\rho^2)}} \exp\left[\frac{(\theta(\omega_i) - \rho\theta(\omega_i))^2}{2(1-\rho^2)}\right], \tag{30}$$

where $\rho = \rho(1) \approx 1 - e/(N-1)$ (see Eq. (18)) and where we have assumed without loss of generality that the mean $\mu(\theta_i)$ and variance $\sigma^2(\theta_i)$ of the technology landscape are 0 and 1, respectively. This annealed approach approximates the *Ne* technology landscape well when $e/N \sim 1$, i.e., when $\rho \sim 0$, but can deviate in some respects when $e/N \sim 0$, i.e., when $\rho \sim 1$ (see Macready, 1999). Eqs. (28)–(30) define a more general family of landscapes (since ρ can be negative) characterized by arbitrary ρ.

Since we are interested in the effects of search at arbitrary distances d from a production recipe ω_i, we must infer $P(\theta(\omega_j)|\theta(\omega_i), d)$ from $P(\theta(\omega_i), \theta(\omega_j))$. We shall not supply this calculation here but only sketch an outline of how to proceed (for full details, see Macready, 1999). To begin, note that $P(\theta(\omega_j)|\theta(\omega_i), d)$ is easily obtainable from $P(\theta(\omega_i), \theta(\omega_j)|d)$ as

$$P(\theta(\omega_j)|\theta(\omega_i), d) = \frac{P(\theta(\omega_i), \theta(\omega_j)|d)}{P(\theta(\omega_i))}. \tag{31}$$

$P(\theta(\omega_i), \theta(\omega_j)|d)$ is not known but it is related to $P(\theta(\omega_i), \theta(\omega_j)|s)$, the probability that an s-step random walk in the technology graph Γ beginning at ω_i and ending at ω_j has labor efficiencies $\theta(\omega_i)$ and $\theta(\omega_j)$ at the endpoints of the walk. (Each step either increases or decreases the distance from the starting point by 1.) $P(\theta(\omega_i), \theta(\omega_j)|s)$ is straightforward to calculate from Eq. (29). $P(\theta(\omega_i), \theta(\omega_j)|d)$ is then obtained from $P(\theta(\omega_i), \theta(\omega_j)|s)$ by including the probability that an s-step random walk on Γ results in a net displacement of d-steps. The result of this calculation is that $P(\theta(\omega_j)|\theta(\omega_i), d)$ is Gaussianly distributed with a mean and variance given by

$$\mu(\omega_i, d) = \theta(\omega_i)\rho^d, \tag{32}$$

$$\sigma^2(\omega_i, d) = 1 - \rho^{2d}. \tag{33}$$

Eqs. (32) and (33) play an important role in the next section.

6. Optimal search distance

6.1. The firm's search problem

In order to determine the relationship between search cost and optimal search distance on a technology landscape, we recast the firm's search problem into the familiar framework of dynamic programming (Bellman, 1957; Bertsekas, 1976; Sargent, 1987). Recall that each production recipe $\omega_i \in \Omega$ is associated with a labor efficiency θ_i. Production recipes at different locations in the technology landscape — and therefore at different distances from each other — have different Gaussian distributions corresponding to different $\mu(\omega_i, d)$ and $\sigma(\omega_i, d)$. The firm incurs a search cost, $c(d)$, every time it samples a production recipe a distance d away from the current production recipe. The search cost $c(d)$ is a monotonically increasing function of d since more distant production recipes require greater changes to the current recipe. For simplicity we take $c(d) = \alpha d$ (see Eq. (21)) but arbitrary functional forms for $c(d)$ are no more difficult to incorporate within our framework. The firm's problem is to determine the optimal search distance at which to sample the technology landscape for improved production recipes. [7]

To determine the optimal distance at which to search for new production recipes we begin by denoting the firm's current labor efficiency by z and supposing that the firm is considering sampling at a distance d. If $F_d(\theta)$ is the cumulative probability distribution of efficiencies at distance d, the firm's expected labor efficiency, $E(\theta|d)$, searching at distance d is given by

$$E(\theta|d) = -c(d) + \beta\left(z\int_{-\infty}^{z} \mathrm{d}F_d(\theta) + \int_{z}^{\infty}\theta\,\mathrm{d}F_d(\theta)\right), \tag{34}$$

where β is the discount factor. It may be the case that this discount factor is d-dependent since larger changes in the production recipe would likely require more time but we shall assume for simplicity that β is independent of d. The difference in labor efficiencies between searching at distance d and remaining with the current production recipe, $D_d(z)$, is given by

$$D_d(z) = E(\theta|d) - z = -c(d) - (1-\beta)z + \beta\int_{z}^{\infty}(\theta - z)\,\mathrm{d}F_d(\theta). \tag{35}$$

$D_d(z)$ is a monotonically decreasing function of z which crosses zero at $z_c(d)$, determined by $D_d(z_c(d)) = 0$. For $z < z_c(d)$ it is best to sample a new production recipe ω_j since $D_d(z)$ is positive. If $z > z_c(d)$ it is best to remain with the current recipe ω_i because $D_d(z)$ will be negative and the cost will outweigh the potential gain. The zero-crossing value $z_c(d)$ thus

[7] Note that since $E[\theta^2] \ll \infty$, by assumption, an optimal stopping rule exists for the firm's search (DeGroot, 1970, Chapter 13).

plays the role of the firm's *reservation price* (Kohn and Shavell, 1974; Bikhchandani and Sharma, 1996). The reservation price at distance d is determined from the integral equation

$$c(d) + (1-\beta) z_c(d) = \beta \int_{z_c(d)}^{\infty} (\theta - z_c(d))\, \mathrm{d}F_d(\theta). \tag{36}$$

From Eq. (36) it can be seen that, as expected, reservation price decreases with greater search cost.

The firm's optimal search strategy on its technology landscape can be characterized by *Pandora's Rule:* if a production recipe at some distance is to be sampled, it should be a production recipe at the distance with the highest reservation price. The firm should terminate search and remain with the current production recipe whenever the current labor efficiency is greater than the reservation price of all distances (a proof of this result is found in Weitzman, 1979).

6.2. The reservation price for Gaussian efficiencies

In the case where labor efficiencies at distance d are Gaussianly distributed, Eq. (36) reads as

$$c(d) + (1-\beta) z_c = \frac{\beta}{\sqrt{2\pi}} \int_{z_c}^{\infty} \frac{\mathrm{d}\theta}{\sigma(\omega_i, d)} (\theta - z_c) \exp\left[-\frac{(\theta - \mu(\omega_i, d))^2}{2\sigma^2(\omega_i, d)}\right], \tag{37}$$

$$\Rightarrow c(d) + (1-\beta) z_c = \frac{\beta}{\sqrt{2\pi}} \int_{0}^{\infty} \frac{\mathrm{d}u}{\sigma(\omega_i, d)} u \exp\left[-\frac{(u + z_c - \mu(\omega_i, d))^2}{2\sigma^2(\omega_i, d)}\right]. \tag{38}$$

(For clarity the d dependence of z_c has been omitted.) Using the definite integral

$$\int_0^{\infty} \frac{\mathrm{d}u}{b} u \exp\left[-\frac{(u-a)^2}{2b^2}\right] = b \exp\left[-\frac{a^2}{2b^2}\right] + a\sqrt{\frac{\pi}{2}}\, \mathrm{erfc}\left[\frac{-a}{\sqrt{2}b}\right], \tag{39}$$

where $\mathrm{erf}[\cdot]$ is the error function and $\mathrm{erfc}[\cdot] = 1 - \mathrm{erf}[\cdot]$ is the complimentary error function,[8] we find that the equation determining the reservation price is

$$\begin{aligned} c(d) + (1-\beta) z_c = \beta \Bigg(& \frac{\mu(\omega_i, d) - z_c}{2}\, \mathrm{erfc}\left[-\frac{\mu(\omega_i, d) - z_c}{\sqrt{2}\sigma(\omega_i, d)}\right] \\ & + \frac{\sigma(\omega_i, d)}{\sqrt{2\pi}} \exp\left[-\frac{(\mu(\omega_i, d) - z_c)^2}{2\sigma^2(\omega_i, d)}\right]\Bigg). \end{aligned} \tag{40}$$

To simplify the appearance of this equation we write it using the dimensionless variable

$$\delta = \frac{z_c - \mu(\omega_i, d)}{\sqrt{2}\sigma(\omega_i, d)}, \tag{41}$$

[8] The error function $\mathrm{erf}(x)$ is defined as $(2/\sqrt{\pi})\int_0^x \exp(-t^2)\,\mathrm{d}t$ and the complimentary error function, $\mathrm{erfc}(x)$ is defined as $(2/\sqrt{\pi})\int_x^{\infty} \exp(-t^2)\,\mathrm{d}t$.

in terms of which $z_c = \sqrt{2}\sigma(\omega_i, d)\delta + \mu(\omega_i, d)$. The dimensionless reservation price δ is then determined by

$$\frac{\sqrt{2}(c(d) + (1-\beta)\mu(\omega_i, d))}{\sigma(\omega_i, d)} = \beta\left(\frac{\exp[-\delta^2]}{\sqrt{\pi}} - \delta\,\mathrm{erfc}[\delta]\right) - 2(1-\beta)\delta, \tag{42}$$

$$\Rightarrow \frac{\sqrt{2}(c(d) + (1-\beta)\mu(\omega_i, d))}{\sigma(\omega_i, d)} = \beta\left(\frac{\exp[-\delta^2]}{\sqrt{\pi}} + \delta\,\mathrm{erfc}[-\delta]\right) - 2\delta. \tag{43}$$

Defining

$$A(\omega_i, d) \equiv \frac{\sqrt{2}(c(d) + (1-\beta)\mu(\omega_i, d))}{\sigma(\omega_i, d)}, \tag{44}$$

the equation which must be solved for δ is therefore

$$A(\omega_i, d) = \beta\left(\frac{\exp[-\delta^2]}{\sqrt{\pi}} + \delta\,\mathrm{erfc}[-\delta]\right) - 2\delta. \tag{45}$$

The explicit ω_i and d dependence of A is obtained by plugging Eqs. (32) and (33) into Eq. (44). Eq. (45) is the central equation determining the reservation price $z_c(d)$. Approximate solutions to this equation are considered in Section 6.3.

The optimal search distance, d^*, is now determined as

$$d^* = \arg\max_d z_c(d), \tag{46}$$

where the d-dependence of $z_c(d)$ is implicitly determined by Eq. (45). As a function of d, z_c is well behaved with a single maximum so that d^* is the integer nearest to the d which solves $\partial_d z_c = 0$. We now proceed to find the equation which d^* satisfies.

To begin, recall the definition of δ given in Eq. (41). Taking the d derivative of δ yields

$$\partial_d z_c = \sqrt{2}(\delta\partial_d\sigma(\omega_i, d) + \sigma(\omega_i, d)\partial_d\delta) + \partial_d\mu(\omega_i, d). \tag{47}$$

The partial derivatives $\partial_d\mu$ and $\partial_d\sigma$ are given by

$$\partial_d\mu(\omega_i, d) = d\theta(\omega_i)\rho^{d-1}, \tag{48}$$

$$\partial_d\sigma(\omega_i, d) = -2d\rho^{2d-1}, \tag{49}$$

respectively, and we wish to express $\partial_d\delta$ in terms of these known quantities. Differentiating Eq. (45) with respect to d yields

$$\partial_d\delta = \frac{\partial_d A(\omega_i, d)}{\beta\,\mathrm{erfc}[-\delta] - 2} \tag{50}$$

(assuming β is not d-dependent). Thus d^* is determined by

$$0 = \sqrt{2}\left(\delta\partial_d\sigma + \frac{\sigma\partial_d A}{\beta\,\mathrm{erfc}[-\delta] - 2}\right) + \partial_d\mu. \tag{51}$$

Using the definition of A in Eq. (44) its derivative is easily found as

$$\partial_d A = \frac{\sqrt{2}}{\sigma}\partial_d c + \frac{\sqrt{2}(1-\beta)}{\sigma}\partial_d \mu - \frac{A}{\sigma}\partial_d \sigma. \tag{52}$$

Plugging this result in Eq. (51) we find

$$0 = \sqrt{2}\left(\delta\partial_d\sigma + \frac{\sqrt{2}\partial_d c + \sqrt{2}(1-\beta)\partial_d\mu - A\partial_d\sigma}{\beta\,\mathrm{erfc}[-\delta] - 2}\right) + \partial_d\mu, \tag{53}$$

which can be rearranged to give

$$0 = 2\partial_d c + \sqrt{2}(\beta\delta\,\mathrm{erfc}[-\delta] - 2\delta - A)\partial_d\sigma + \beta(\mathrm{erfc}[-\delta] - 2)\partial_d\mu. \tag{54}$$

Finally, we use Eq. (45) to simplify this to

$$\frac{2}{\beta}\partial_d c = \sqrt{\frac{2}{\pi}}\exp[-\delta^2]\partial_d\sigma + \mathrm{erfc}[\delta]\partial_d\mu, \tag{55}$$

where $\partial_d\mu$ and $\partial_d\sigma$ are given in Eq. (49).

6.3. Determination of the reservation price

It is desirable to have an explicit solution for δ (implicitly determined by Eq. (45)). To this end we note some features of the function

$$D_A(\delta) \equiv \beta\left(\frac{\exp[-\delta^2]}{\sqrt{\pi}} + \delta\,\mathrm{erfc}[-\delta]\right) - 2\delta - A(\omega_i, d). \tag{56}$$

Firstly, note that

$$\lim_{\delta\to-\infty} D_A(\delta) = \infty, \qquad \lim_{\delta\to\infty} D_A(\delta) = -A \tag{57}$$

and that $D_A(\delta)$ is monotonic. Thus, there is no solution to $D_A(\delta) = 0$ unless $A > 0$. If $A < 0$ then it is always profitable to try new production recipes. This is the case, e.g., when $c(d)$ is negative and is sufficiently large in magnitude. We assume that the firm is not paid to try new production recipes and confine ourselves to the case $A > 0$.

In the case $A \gg 1$, the solution δ of $D_A(\delta) = 0$ is large and negative. In this case the term multiplying β is almost zero and to a very good approximation the solution of $D_A(\delta) = 0$ is

$$\delta = -\tfrac{1}{2}A, \tag{58}$$

or $z_c(d) = -c(d) + \beta\mu(\omega_i, d)$. The d dependence of the reservation price in this limit is particularly simple

$$z_c(d) = \beta\theta\rho^d - \alpha d. \tag{59}$$

This is maximal for $d = 0$, corresponding to terminating the search. This result makes intuitive sense because if A is large then either search costs are high and additional sampling is too expensive or labor efficiencies are high and it is unlikely to find improved production recipes. We thus find that there are diminishing returns to search depending upon the firm's current location in the technological landscape.

In the opposite limit, $0 < A \ll 1$, the solution is at δ large and positive. In this case we use the asymptotic expansion [9]

$$\mathrm{erfc}[-\delta] = 2 - \frac{\exp[-\delta^2]}{\pi}\sum_{k=0}^{n-1}\frac{(-1)^k \Gamma(k+\frac{1}{2})}{\delta^{2k+1}} + \frac{\exp[-\delta^2]}{\pi} R_n, \tag{60}$$

where $|R_n| < \Gamma(n+\frac{1}{2})/\delta^{2n+1}$. Working to third order in $1/\delta$ and recalling that $\Gamma(\frac{1}{2}) = \sqrt{\pi}$ gives the approximate equation

$$A(\omega_i, d) = 2(\beta - 1)\delta + \frac{\beta}{2\sqrt{\pi}\delta^2}\exp[-\delta^2]. \tag{61}$$

In the special case $\beta = 1$, δ is determined by

$$\delta^2 \exp[\delta^2] = \frac{1}{2\sqrt{\pi}A}, \tag{62}$$

which has the solution

$$\delta = \sqrt{W\left[\frac{1}{2\sqrt{\pi}A}\right]}, \tag{63}$$

where $W[\cdot]$ is Lambert's W function [10] defined implicitly by $W[x]\exp W[x] = x$. For small A we can use the asymptotic expansion $W(x) \sim \ln x$ (see Corless et al., 1996) to write

$$\delta \sim \sqrt{-\ln[2\sqrt{\pi}A]} = \sqrt{\ln\left[\frac{2\sqrt{2\pi}\sigma(\omega, d)}{c(d)}\right]}. \tag{64}$$

6.4. Numerical results

In this section we present results for the optimal search distance as a function of: (i) the initial labor efficiency of the firm, (ii) the cost of search as represented by α in $c(d) = \alpha d$; and (iii) the correlation ρ of the technology landscape. For brevity we will not present the β dependence but note that $\beta < 1$ decreases the optimal search distance.

In appropriate parameter regimes we have used the approximations in Eqs. (58) and (61), elsewhere we have resorted to a numerical solution to Eqs. (45) and (55).

Figs. 4 and 5 present the optimal search distance d^* as a function of the firm's current efficiency and the search cost parameter, α. In regions of parameter space in which the optimal search distance is zero it is best to terminate the search and not search for more

[9] The Γ function is defined by $\Gamma(x) = \int_0^\infty \mathrm{d}t \exp[-t]t^{x-1}$. For integer x, $\Gamma(x) = (x-1)!$.

[10] See Corless et al. (1996) for a good introduction to Lambert's W function.

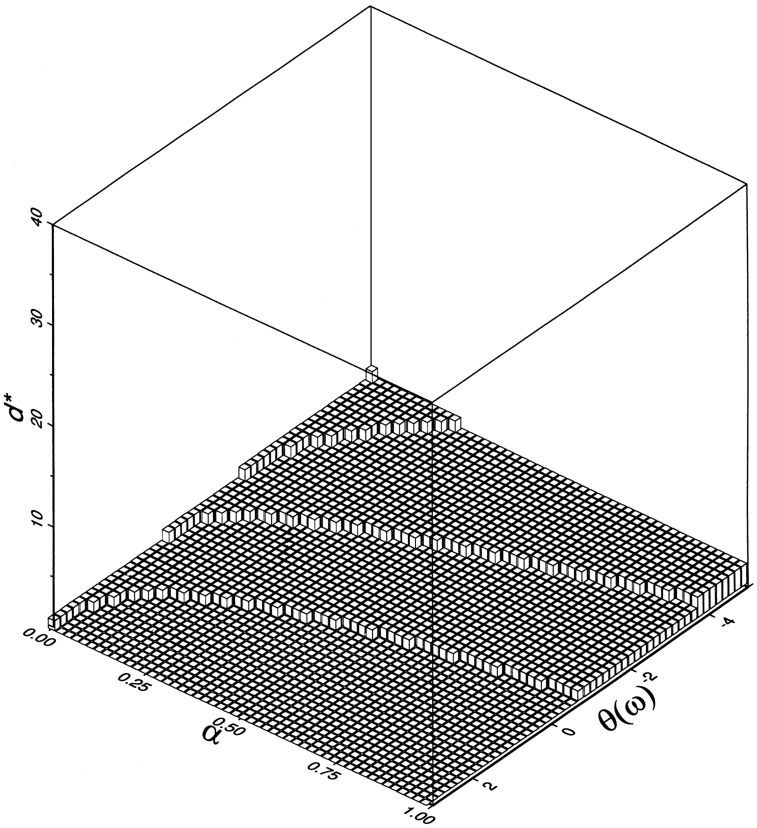

Fig. 4. Optimal search distance d^* as a function of the search cost α and the initial labor efficiency $\theta(\omega)$ for a landscape with correlation coefficient of $\rho = 0.3$.

efficient production recipes. We note a number of features paralleling the simulation results presented in Section 4.2. In general, for low initial efficiencies it is better for the firm to search for improved production recipes farther away. As search costs increase (i.e., as α increases), the additional cost limits optimal search closer to the firm's current production recipe. For production recipes which are initially efficient, the advantages of search are much less pronounced and for high enough initial efficiencies it is best to consider only single-operation variants. Again, a higher cost of search results in even smaller optimal search distances.

The effects of landscape correlation (as measured by ρ) on optimal search distance are dramatic. On highly correlated technology landscapes (e.g., $\rho = 0.9$), correlation extends across large distances and as a result large optimal search distances are obtained (see Fig. 5). For a technology landscape with little correlation ($\rho = 0.3$), optimal search distances shrink (see Fig. 4). In the limiting case of a completely uncorrelated technology landscape ($\rho = 0$), all search distances are equivalent since no landscape correlation exists to exploit during the search.

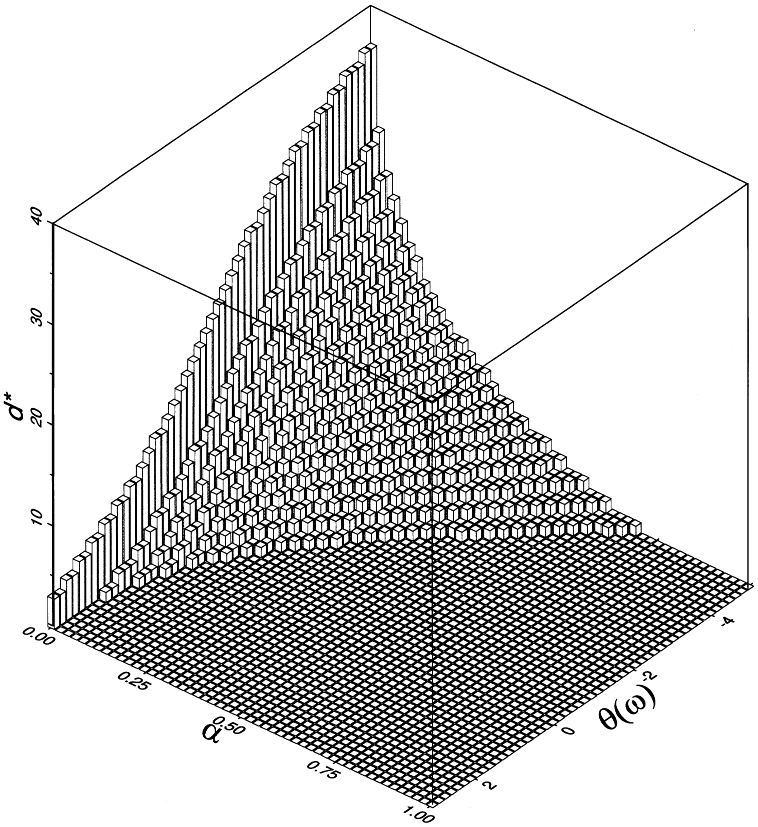

Fig. 5. Optimal search distance d^* as a function of the search cost α and the initial labor efficiency $\theta(\omega)$ for a landscape with correlation coefficient of $\rho = 0.9$.

7. Conclusion

In this discussion we have been concerned with the determination of the optimal distance at which a firm should seek technological improvements in a space of possible technologies. In our model the firm's technology is determined by its organizational capital which in turn is represented by a production recipe whose N constituent operations can occupy S discrete states. Different configurations for a production recipe represent different technologies. Production recipes are also characterized by the level of external economies and diseconomies among the recipe's operations; the parameter e measures the level of "intranalities" of a production recipe. The distance between any two distinct production recipes in the space of technological possibilities is naturally determined by the number of operations whose states need to be changed in order to turn one configuration into another. Our use of production recipes is related to — but goes beyond — the work on activity analysis by Koopmans (1951) and Leontief (1953) and the work of Chenery (1949), Smith (1961) and Marsden et al. (1974) on engineering production functions. As with this previous work, our approach grounds the modeling of productive activity on engineering practice but unlike the early work attempts to provide a sufficient basis for modeling technological evolution.

In order to study how the current location of the firm in the space of technological possibilities affects the firm's search for technological improvements, we model the firm's search as movement on a "technology landscape". The locations in the landscape correspond to different configurations for the firm's production recipe. Local maxima and minima for the labor efficiency associated with each production recipe are represented by "peaks" and "valleys" in the landscape. The "ruggedness" of the landscape is in turn determined by the landscape's correlation coefficient, ρ.

Our initial investigation about the firm's optimal search distance involved computational exploration of the *Ne* technology landscape. The obtained simulation results prompted the development of a formal framework in which a technology landscape was incorporated into a standard dynamics programming model of search. The resulting framework abstracts away from all landscape detail except the important statistical structure which is captured in relatively simple probability distributions. As our main result we find that early in the search for technological improvements, if the initial position is poor or average, it is optimal to search far away on the technology landscape. As the firm succeeds in finding technological improvements, however, it is optimal to confine search to a local region of the technology landscape. Our modeling framework results in an intuitive and satisfying picture of optimal search as a function of the cost of search (which is itself a function of the distance between the firm's currently utilized production recipe and the newly sampled recipe), the firm's current location on the space of technological possibilities and the correlation structure of the technology landscape.

The general features of the story told in this paper — that early search can give rise to dramatic improvements via significant alterations found far away across the space of possibilities but that later search closer to home yields finer and finer twiddling with the details — suggests a possible application of our model to treat the development of "design types". Among the stylized facts accepted by most engineers is the view that, soon after a major design innovation, improvement occurs by the emergence of dramatic alterations in the fundamental design. Later, as improvements continue to accumulate, variations settle down to minor fiddling with design details. We need only to think of the variety of forms of the early bicycles — big-front-wheel–small-back-wheel, small-front-wheel–big-back-wheel, various handle-bars — or of the forms of aircraft populating the skies in the early decades of the century. [11]

We believe that technology landscapes as introduced here can be a useful tool to study firm behavior. However, much future work clearly remains. Perhaps the most direct extension of our model would be to treat landscapes as Markov random fields where the full neighborhood $\mathcal{N}_1$ around any particular configuration is included and results from the study of Markov random fields can be exploited (see, e.g., Kindermann and Snell, 1980). It would be desirable to build a model in which the correlation ρ of the technology landscape arises endogenously rather than treating it as an external parameter as we have done here.

[11] Dyson (1997) estimates that there were literally thousands of aircraft designs flown during the 1920s and 1930s of which only a few hundred survived to form the basis of modern aviation.

In this paper we have studied the optimal search distance for a single firm to sample its technology landscape. But as remarked by Stuart and Podolny (1996), firms do not search in isolation, rather they search as members of a population of simultaneously searching organizations. How is the optimal search distance for an individual firm affected by the presence of other firms exploring the same technology landscape? If the cost of search increases with distance and optimal search distance decreases with increasing efficiency, how often will firms get "trapped" in suboptimal procedures or products? Since, in general, the structure of the technology landscape is only know locally, can a firm search in such a way so as to optimize both improvements on the landscape and learning about the landscape's structure in order to guide further search? These and related questions await further investigation.

Acknowledgements

The authors thank Phil Auerswald, Richard Day, Richard Durrett, Alfred Nucci, Richard Schuler, Robert Solow, Nancy Stokey, Willard Zangwill and an anonymous referee for their helpful comments. The authors also gratefully acknowledge research support provided by the Santa Fe Institute and Bios Group LP.

References

Adams, J.D., Sveikauskas, L.,1993. Academic science, industrial R&D, and the growth of inputs. Center for Economic Studies Discussion Paper 93-1. US Bureau of the Census, Washington, DC.
Audretsch, D., 1991. New firm survival and the technological regime. Review of Economics and Statistics 73, 441–450.
Audretsch, D., 1994. Business survival and the decision to exit. Journal of Business Economics 1, 125–138.
Auerswald, P., Lobo, J., 1996. Learning by doing, technological regimes and industry evolution. Presented at the 71st Annual Meeting of the Western Economic Association, San Francisco, CA.
Auerswald, P., Kauffman, S., Lobo, J., Shell, K., 2000. The production recipes approach to modeling technological innovation: an application to learning by doing. Journal of Economic Dynamics and Control 24, 389–450.
Bailey, M.N., Bartelsman, E.J., Haltiwanger, J., 1994. Downsizing and productivity growth: myth or reality? Center for Economic Studies Discussion paper 94-4. US Bureau of the Census, Washington, DC.
Barney, J.B., 1991. Firm resources and sustained competitive advantage. Journal of Management 17, 99–120.
Bellman, R., 1957. Dynamic Programming. Princeton University Press, Princeton, NJ.
Bertsekas, D.P., 1976. Dynamic Programming and Stochastic Control. Academic Press, New York.
Bikhchandani, S., Sharma, S., 1996. Optimal search with learning. Journal of Economic Dynamics and Control 20, 333–359.
Boeker, W., 1989. Strategic change: the effects of founding and history. Academy of Management Journal 32, 489–515.
Cameron, P.J., 1994. Combinatorics: Topics, Techniques and Algorithms. Cambridge University Press, New York.
Chenery, H.B., 1949. Engineering production functions. Quarterly Journal of Economics 63, 507–531.
Coase, R., 1937. The nature of the firm. Economica 4, 386–405.
Cohen, W.M., Levinthal, D.A., 1989. Innovation and learning: the two faces of R&D. Economic Journal 99, 569–596.
Corless, R.M., Gonnet, G.H., Hare, D.E.G., Jeffrey, D.J., Knuth, D.E., 1996. On the Lambert *W* function. Advances in Computational Mathematics 5, 329–359.
Davis, S.J., Haltiwanger, J., 1992. Gross job creation, gross job destruction, and employment reallocation. Quarterly Journal of Economics 107, 819–863.
DeGroot, M.H., 1970. Optimal Statistical Decisions. McGraw-Hill, New York.

Derrida, B., Pomeau, Y., 1986. Random networks of automata: a simple annealed approximation. Europhysics Letters 1, 45–49.
Dunne, T., Roberts, M., Samuelson, L., 1988. Patterns of firm entry and exit in US manufacturing industries. RAND Journal of Economics 19, 495–515.
Dunne, T., Roberts, M., Samuelson, L., 1989. The growth and failure of US manufacturing plants. Quarterly Journal of Economics 104, 671–698.
Dunne, T., Haltiwanger, J., Troske, K.R., 1996. Technology and jobs: secular change and cyclical dynamics. NBER Working Paper 5656. National Bureau of Economic Research, Cambridge, MA.
Dwyer, D.W., 1995. Technology locks, creative, destruction and non-convergence in productivity levels. Center for Economic Studies Discussion Paper 95-6. US Bureau of the Census, Washington, DC.
Dyson, F., 1997. Imagined Worlds: The Jerusalem–Harvard Lectures. Harvard University Press, Cambridge, MA.
Eigen, M., McCaskill, J., Schuster, P., 1989. The molecular quasispecies. Advances in Chemical Physics 75, 149–263.
Ericson, R., Pakes, A., 1995. Markov — perfect industry dynamics: a framework for empirical work. Review of Economic Studies 62, 53–82.
Evenson, R.E., Kislev, Y., 1976. A stochastic model of applied research. Journal of Political Economy 84, 265–281.
Fontana, W., Stadler, P.F., Bornberg-Bauer, E.G., Griesmacher, T., Hofacker, I.L., Tacker, M., Tarazona, P., Weinberger, E.D., Schuster, P., 1993. RNA folding and combinatory landscapes. Physical Review E 47, 2083–2099.
Griffeath, D., 1976. Introduction to random fields. In: Kemeny, J., Snell, J., Knapp, A. (Eds.), Denumerable Markov Chains. Springer, New York.
Hall, R.E., 1993. Labor demand, labor supply, and employment volatility. NBER Macroeconomics Annual 6, 17–47.
Helfat, C.E., 1994. Firm specificity and corporate applied R&D. Organization Science 5, 173–184.
Henderson, R.M., Clark, K.B., 1990. Architectural innovation: the reconfiguration of existing product technology and the failure of established firms. Administrative Science Quarterly 35, 9–31.
Herriott, S.R., Levinthal, D.A., March, J.G., 1985. Learning from experience in organizations. American Economic Review 75, 298–302.
Hey, J.D., 1982. Search for rules of search. Journal of Economic Behavior and Organization 3, 65–81.
Hopenhayn, H., 1992. Exit, entry, and firm dynamics in long run equilibrium. Econometrica 60, 1127–1150.
Jovanovic, B., 1982. Selection and the evolution of an industry. Econometrica 50, 659–670.
Jovanovic, B., Rob, R., 1990. Long waves and short waves: growth through intensive and extensive search. Econometrica 58, 1391–1409.
Kauffman, S., 1993. Origins of Order: Self-organization and Selection in Evolution. Oxford University Press, New York.
Kennedy, P.M., 1994. Information processing and organizational design. Journal of Economic Behavior and Organization 25, 37–51.
Kindermann, R., Snell, J.L., 1980. Markov Random Fields and their Applications. American Mathematical Society, Providence, RI.
Kohn, M., Shavell, S., 1974. The theory of search. Journal of Economic Theory 9, 93–123.
Koopmans, T. (Ed.), 1951. Activity Analysis of Production and Allocation. Wiley, New York.
Lee, D.M., Allen, T.J., 1982. Integrating new technical staff: implications for acquiring new technology. Management Science 28, 1405–1420.
Leontief, W.L., 1953. Studies in the Structure of the American Economy: Theoretical and Empirical Explorations in Input–Output Analysis. Oxford University Press, New York.
Levinthal, D.A., March, J.G., 1981. A model of adaptive organizational search. Journal of Economic Behavior and Organization 2, 307–333.
Lobo, J., Macready, W.G., 1999. Landscapes: a natural extension of search theory. Santa Fe Institute Working Paper 99-05-037E. The Santa Fe Institute, Santa Fe, NM.
Macken, C.A., Stadler, P.F., 1995. Evolution on fitness landscapes. In: Nadel, L., Stein, D. (Eds.), 1993 Lectures in Complex Systems. Addison-Wesley, Reading, MA.
March, J.G., 1991. Exploration and exploitation in organizational learning. Organization Science 2, 71–87.
Marengo, L., 1992. Coordination and organizational learning in the firm. Journal of Evolutionary Economics 2, 313–326.

Marsden, J.R., Pingry, D., Whinston, D., 1974. Engineering foundations of production functions. Journal of Economic Theory 9, 124–140.

Muth, J.F., 1986. Search theory and the manufacturing progress function. Management Science 32, 948–962.

Nelson, R.R., Winter, S.G., 1982. An Evolutionary Theory of Economic Change. Belknap Press, Cambridge, MA.

Papadimitriou, C.H., Steiglitz, K., 1982. Combinatorial Optimization: Algorithms and Complexity. Prentice-Hall, Englewood Cliffs, NJ.

Presscott, E., Visscher, M., 1980. Organization capital. Journal of Political Economy 88, 446–461.

Reiter, S., Sherman, G.R., 1962. Allocating indivisible resources affording external economies of diseconomies. International Economic Review 3, 108–135.

Reiter, S., Sherman, G.R., 1965. Discrete optimizing. SIAM Journal 13, 864–889.

Romer, P.M., 1990. Endogenous technological change. Journal of Political Economy 98, 71–102.

Sahal, D., 1985. Technological guideposts and innovation avenues. Research Policy 14, 61–82.

Sargent, T.J., 1987. Dynamic Macroeconomic Theory. Harvard University Press, Cambridge, MA.

Shan, W., 1990. An empirical analysis of organizatonal strategies by entrepreneurial high-technology firms. Strategic Management Journal 11, 129–139.

Smith, V.L., 1961. Investment and Production: A Study in the Theory of the Capital-using Enterprise. Harvard University Press, Cambridge, MA.

Stadler, P.F., 1995. Towards a theory of landscapes. Social Systems Research Institute Working Paper No. 9506. University of Wisconsin, Madison, WI.

Stadler, P.F., Happel, R., 1995. Random field models for fitness landscapes. Santa Fe Institute Working Paper 95-07-069. The Santa Fe Institute, Santa Fe, NM.

Stone, L.D., 1975. Theory of Optimal Search. Academic Press, New York.

Stuart, T.E., Podolny, J.M., 1996. Local search and the evolution of technological capabilities. Strategic Management Journal 17, 21–38.

Tesler, L.G., 1982. A theory of innovation and its effects. The Bell Journal of Economics 13, 69–92.

Tushman, M.L., Anderson, P., 1986. Technological discontinuities and organizational environments. Administrative Science Quarterly 14, 311–347.

Vanmarcke, E., 1983. Random Fields: Analysis and Synthesis. MIT Press, Cambridge, MA.

Weinberger, E.D., 1990. Correlated and uncorrelated fitness landscapes and how to tell the difference. Biological Cybernetics 63, 325–336.

Weitzman, M.L., 1979. Optimal search for the best alternative. Econometrica 47, 641–654.

Williamson, O., 1985. The Economic Institution of Capitalism: Firms, Markets and Relational Contracting. Free Press, New York.