

Neutrality in Technological Landscapes

José Lobo¹ John H. Miller² Walter Fontana³

January 26, 2004

¹Santa Fe Institute, *email:* jose@santafe.edu.

²Carnegie Mellon University and Santa Fe Institute, *email:* miller@santafe.edu,
and corresponding author.

³Santa Fe Institute, *email:* walter@santafe.edu.

Abstract

Searching for technological and organizational innovations is a fundamental activity in social systems. Recent work in this area has focused on the notion of firms locally searching for high points on a “fitness landscape.” Here we show how the presence of neutral networks forming pathways of temporarily inconsequential changes can alter the potential of innovation in such worlds. In particular, such neutral muddling can ultimately improve the potential of the system for innovative leaps. Neutrality provides organizational and technological systems a degree of robustness against tinkering while simultaneously being a crucial enabler of change and innovation.

1 Introduction

Developing a better understanding of technological and organizational change is a key goal of the social sciences. Recent work in this area has relied on biological metaphors to create new models—in particular, these models have focused on the notion of search across a “fitness landscape.”¹ Here, we discuss how new developments in the understanding of biological evolution, namely the notion of neutral networks (Schuster *et al.*, 1994), may fundamentally alter our view of technological and organizational change.

The concept of a fitness landscape was introduced into biology by Wright (1932) and has become one of the most prevalent metaphors for thinking about biological evolution. A fitness landscape is a representation of a search space in which peaks and valleys are associated with genotypes that give rise to successful and unsuccessful phenotypes respectively (Provine, 1986). In this metaphor, evolution is viewed as an adaptive walk of populations across the fitness landscape. A key characteristic of such landscapes is its level of ruggedness. A smooth fitness landscape is a search space with only a few high-fitness solutions. In contrast, a rugged landscape corresponds to a search space with an irregular topography consisting of numerous peaks (local fitness optima).

Technological and organizational search is widely presumed to take place on rugged landscapes. As discussed below, firms and organizations are thought to rely on local search by sampling variants located near a currently implemented solution. Local search on a rugged landscape tends to result in the relatively quick discovery of a local optimum. Firms are therefore presumed to frequently “get stuck” on sub-optimal solutions, a consequence often emphasized in the economic landscapes literature. Yet, as recently pointed out by Baumol (2002), updating Schumpeter’s (1942) argument, firms in market economies do routinely innovate and improve products and processes. This observation is at odds with a central prediction of current landscape models.

The neutralist perspective of evolution (Maynard-Smith, 1970; Kimura, 1968, 1983) argues that, at the molecular level, most genetic change is non-

¹See, for example, Kauffman, 1995; Westhoff, Yarbrough, and Yarbrough, 1996; Dalle, 1997; Levinthal, 1997; Beinhocker, 1999; Levinthal and Warglien, 1999; McKelvey, 1999; Kauffman, Lobo, and Macready, 2000; McCarthy and Tan, 2000; Ebeling, Karmeshu, and Scharnhorst, 2001; Constant, 2002; Levitan, Lobo, Schuler, and Kauffman, 2002; Rivkin and Siggelkow, 2002; Fleming and Sorenson, 2001, 2003; and Haslett and Osborne, 2003.

adaptive, that is, neutral or quasi-neutral. Under this perspective, neutrality is largely viewed as a manifestation of system robustness against genetic mutations. Yet, neutral configurations do not occur as scattered, isolated points in genetic space, but rather form connected networks where small steps lead from one neutral configuration to another. In this way, a random walk can trace out many neutral pathways reaching across genetic space. As pointed out by Huynen *et al.* (1996) and Fontana and Schuster (1998), connected neutrality enables evolutionary change by providing the specific genetic contexts necessary for subsequent mutations to become consequential. That is, neutral genetic changes can accumulate in the genotype and thereby set the context for a subsequent mutation to cause an advantageous change that would have been unattainable from the earlier genetic configurations. From this point of view, the same mechanism that buffers a biological organism against genetic perturbations not only conveys robustness to a lineage, but also evolvability, that is, the capacity to evolve (Wagner and Altenberg, 1996). Neutral networks transform the typical picture of a rugged landscape with localized fitness peaks into a landscape characterized by vast interconnected systems of ridges corresponding to suboptimal solutions. Critically, this implies that equivalent solutions are no longer localized in genetic space, and thus our usual measures of ruggedness and correlation do not adequately characterize the landscape.

We argue here that neutral networks are an important feature of technological and organizational landscapes. Firms often experiment with modifying extant solutions of organizational or technological problems which frequently do not result in noticeable change or improvement (a point made by Arrow (1974)). The accumulation of temporarily inconsequential (hence neutral) changes may seem useless, yet such changes produce precisely the context needed for innovation. Given sufficient neutrality, landscapes can be searched efficiently by constantly repositioning in configuration space suboptimal solutions. Notwithstanding the apparent futility of this process, such muddling ultimately improves the potential of the system for innovative leaps.

2 Innovation Through Local Search

Successful firms are capable of steady, and occasionally dramatic, improvements in performance across a wide variety of dimensions. Learning and innovation are a quest into the unknown, involving a search of technological,

organizational, and market opportunities.² This search process takes place within a space of possibilities whose elements are all of the potential variations of technologies, production processes, operational routines, engineering designs, organizational forms, inventory methods, scheduling systems, supply chains, managerial strategies, and so on, utilized by the firm. The various means by which a firm moves in this space of possibilities greatly influence the direction, rate, and overall success of learning and innovation. At any time a firm can perform an experiment in hopes of finding new improvements.

The empirical literature on technology management and firm-level technological change emphasizes that, although firms employ a wide range of search strategies, they tend to engage in *local* search—that is, search that enables firms to build upon their established technology and expertise (see, for example, Sahal, 1981; Freeman, 1982; Lee and Allen, 1982; Hannan and Freeman, 1984; Tushman and Anderson, 1986; Boeker, 1989; Henderson and Clark, 1990; Shan, 1990; Barney, 1991; Helfat, 1994; and Herriott, Levinthal and March, 1995). The prevalence of local search stems from the significant effort required for firms to achieve a given 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 knowledge base (see the discussion in Abernathy and Clark, 1985; Cohen and Levinthal, 1989; Stuart and Podolny, 1996; Potts, 2001; Fleming, 2001; and Loasby, 2002).

There is also ample evidence that while firms are sporadic inventors, they are routine innovators (Schmookler, 1966; Rosenberg, 1982; and Baumol, 2002). Although most large companies (particularly those in oligopolistic high-technology markets) frequently invest in research, few make a habit of seeking inventions, preferring instead to tinker, recombine, and perfect.³ The design skills, technical know-how, organizational knowledge, and managerial styles resident in a given company result from the cumulative choices

²Much of the modern macroeconomics and management science literature on technological and organizational innovation is couched in the framework of search theory. See, for example, Kohn and Shavell, 1974; Evenson and Kislev, 1976; Weitzman, 1979; Levinthal and March, 1981; Hey, 1982; Jovanovic, 1982; Nelson and Winter, 1982; Tesler, 1982; Dosi, 1984; Muth, 1986; Sargent, 1987; Cohen and Levinthal, 1989; Jovanovic and Rob, 1990; Marengo, 1992; Ericson and Pakes, 1995; Rosenberg, 1995; Bikhchandani and Sharma, 1996; Klepper, 1996; Hoppe, 2000; and Mahdi, 2003.

³Of relevance here is the observation that understanding often followed practice in the application of science in the United States industry, with analysis and quantification only coming after tinkering and experimentation (see Rosenberg and Birdzell, Jr., 1986, and Mowery and Rosenberg, 1998).

made by the firm’s engineers, scientists, and managers. These choices tend to reinforce successful practices and steer the firm away from “disruptive” changes. Furthermore, as a market matures, the best performing companies tend to be those that are most aggressive in their local search strategies, what Christensen (1997) refers to as “component level innovation.”

Note that local and undirected search is not necessarily antithetical to innovation. In their influential examination of successful business organizations, Collins and Porras (1994, p. 141) observe: “In examining the history of the visionary companies, we were struck by how often they made some of their best moves not by detailed strategic planning, but rather by experimentation, trial and error, opportunism and—quite literally—accident. What looks in hindsight like a brilliant strategy was often the residual result of opportunistic experimentation and ‘purposeful accidents.’ ” Empirical evidence, engineering practice, and historical records all strongly suggest that a firm’s current technological, managerial, and organizational practices greatly constrain its technological search to remain close to what the firm already does and knows (Basalla, 1988; Anderson and Tushman, 1990; Freeman, 1994; Ashmos, Duchon, and McDaniel, 1998; and Caselli, 1999).

3 A General Search-Landscape Framework

We treat technologies as combinatorial systems, whose global behavior results from the interactions among constituent components (similar approaches are used by Romer, 1990, 1996; Weitzman, 1996, 1998; Brusoni and Prencipe, 2001; and Fleming, 2001). We consider the phenomenon of neutrality to be of particular relevance for understanding the evolution of “complex” technologies. Indeed, Iansiti and Khanna (1995) call a technology complex if its functional characteristics can be obtained through multiple configurational, material, and engineering approaches.

A high degree of interconnectivity among constituent components appears to be another characteristic of complex technologies (Potts, 2001; Loasby, 2001). Connectivity causes the performance of a component to affect, or be affected by, other components. In biology, the consequences of a mutation of gene i depend on the particular variant of gene j present in that same genome. This phenomenon is known as epistasis.

The interactions among components influence both the ruggedness of the landscape and the extent and structure of neutral networks. Ruggedness arises due to “frustration,” a term used in physics (Anderson, 1972) to de-

note multiple conflicting interactions on a given component, such that no set of behaviors of the components can simultaneously satisfy (or optimize) all of the constraints. Frustration challenges the ability of firms to find good solutions.⁴ The structure and extent of neutral networks also is tied to component interactions because such interconnections will often mask the consequences of changes in individual components.

The biological metaphor of a “fitness” landscape provides a sensible representation of a technology search space. We can think of each possible variant for a given technology that a firm may want to implement—a configuration—as a location on a landscape. Neighboring locations are “close” in the sense that it is easy for a firm to move from one neighboring configuration to another. Each configuration is associated with a fitness value (payoff or performance), and the height of the corresponding point on the landscape reflects this value. Thus, the “peaks” in the landscape represent configurations yielding good solutions, while the “valleys” correspond to undesirable ones.⁵

An important property of a landscape is its *correlation* (Stadler, 1992; Fontana *et al.*, 1993)—a measure of the extent to which nearby technological variants have similar levels of performance. Landscape correlation is low if slight changes to a solution drastically alter performance; It is high if performance is relatively insensitive to such changes. Such correlations are computed by averaging (squared) performance differences over a large sample of configuration pairs at some distance d .

In traditional models of search on fitness landscapes, the correlation of the landscape is a key factor in determining the effectiveness of search: landscapes with low correlations are more “rugged,” and thus are likely to trap agents on local peaks.

The presence of neutral networks, however, challenges this traditional view of correlation coinciding with the difficulty of search in high-dimensional spaces. Neutral networks are paths of equivalent neighbors percolating across the space. Any such path only requires the proper alignment of points along a few dimensions, and thus it is not directly tied to the average differences

⁴A similar theme is found in the work of Sagan (1993), Kennedy (1994), Dörner (1996), and Perrow (1999).

⁵Landscapes arise in many settings and have been used to model evolutionary genetics, molecular biology, combinatorial optimization, chemistry, economics, and statistical mechanics. For a comprehensive discussion of landscape models see Macken, Hagan, and Perelson (1991); Macken and Stadler (1995); and Stadler (1995).

that are the basis of correlation measures. It is possible, therefore, to have a very “rugged” landscape (in terms of correlation) that can be easily traversed via a neutral network (Huynen *et al.*, 1996). If so, innovation on complex search landscapes may be far easier than previously thought.

3.1 Search as a Walk on a Landscape

We now describe the general features of the technological problem facing the firm in our modeling framework. The firm’s technology is comprised of a number of distinct operations which at each moment can be in one of a finite number of possible, discrete, states. Consequently, improvements in the technology entails changes in the state of the components constituting the technology.⁶ Technological improvements result from the firm finding better configurations for its technology. Thus, the firm’s technological problem is a combinatorial optimization problem. A technological landscape is a means of representing the problem faced by the firm in its search for the optimal technology.

The search by firms for new technological solutions can be conceptualized as a simple *adaptive walk* on a technological landscape. In such a walk, a firm searches for improvements by sampling for new solutions at distance d from the currently available solution. (In the case of local search, d is small relative to the diameter of the configuration space.) Sampling implies an experiment in which the firm changes an existing technology and evaluates its consequences. If the performance of the experimental configuration is equal to or greater than the performance of the current configuration, the firm adopts the experimental configuration and “moves” to the new location on the technological landscape. This sampling step is then iterated starting with the newly adopted location on the landscape. Firms continue making such steps until they reach a solution that cannot be further improved through local search.

Note that neutral networks cannot be exploited if the decision to adopt the experimental configuration requires the latter to be *better* than the current configuration. Furthermore, the notion of local optimum is deceptive in

⁶Our view of technological innovation is similar to that of Romer (1990), who remarks that over the past few hundred years the raw materials used in production have not changed, but that as a result of trial and error, experimentation, refinement, and scientific investigation, the “instructions” followed for combining raw materials have become vastly more sophisticated. Our “technologies” are analogous to Romer’s “instructions.”

the presence of neutral networks. While the performance of a configuration may not be amenable to improvement locally or even regionally, the situation is unlike that of a conventional localized peak. A neutral network is a non-local structure, like a ridge-system, which permits level motion into far-away regions of the landscape where improvement may become possible again. Whether adaptive walks starting at different locations on the same technological landscape end up at different local optima, depends therefore on the extent of neutrality.

Figure 1, panel (a), provides a simple illustration of these ideas. Agents must navigate this two-dimensional search space whose configurations have three possible payoffs (low, moderate, or high). Moreover, agents are constrained to move only within a neighborhood of their current location in configuration space. Consider an agent starting in the upper-left corner. Such an agent receives a moderate payoff, and local moves either leave the agent indifferent or worse off. If the search strategy requires improvement, the agent is destined to remain in this corner. Yet, if the strategy only requires avoiding a worsening of the payoff, the agent has an escape path through a narrow neutral ridge, providing moderate payoff, between the upper-left and lower-right corners. On this path, a myopic agent would diffuse back and forth through neutral walks. Eventually, however, the diffusive motion would move the agent into the high payoff domain in the lower-right corner. In this particular example, given sufficient time, such a random walk will always succeed.⁷

Figure 1, panel (b), illustrates the impact of adjacency among neutral networks. Three neutral networks, each capturing one of three different “functionalities,” are shown on the left-hand side of the panel. The likelihood of a transition from, say, the square-functionality to the circle-functionality is measured by the size of the shared boundary between these two networks relative to the total boundary size of the square network. These boundary relationships between the networks determine the “accessibilities” of the corresponding functionalities. It is conceptually convenient (Fontana and Schuster, 1998; Stadler *et al.*, 2001) to define a threshold of accessibility above which a functionality is declared to be a neighbor of another functionality. Note that this relationship of accessibility may not be symmetric.⁸

⁷For a true one-dimensional random walk the asymptotic probability that the system has not moved to the higher payoff by time t scales by l/\sqrt{t} , where l is the length of the path.

⁸In which case the space of functionalities is not a metric space.

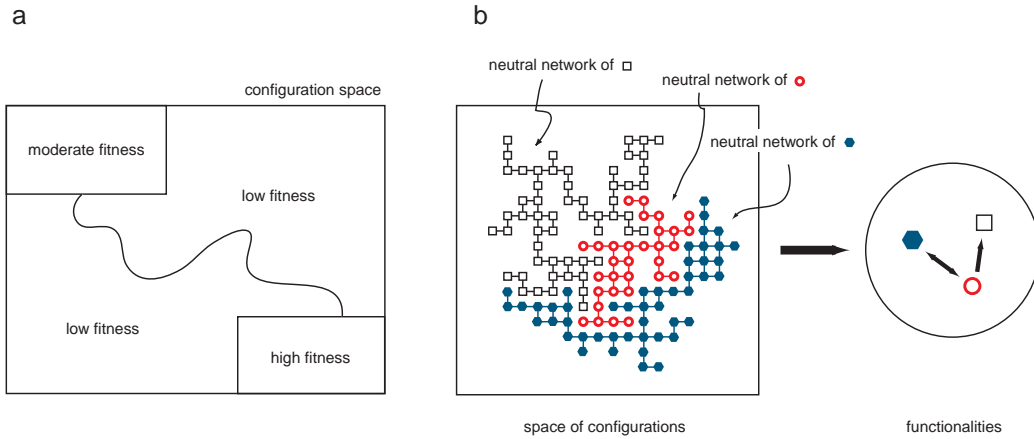


Figure 1: Some neutrality schematics.

Intuitively, it is possible for, say, the country of Monaco to be near France, because a random step out of Monaco is likely to end up in France; Yet, by that same criterion France is not near Monaco.

Organizational and technological search has come to be seen as requiring a judicious choice between exploitation and exploration. In March’s (1991) discussion of organizational search, a distinction is made between the exploration of new possibilities and the exploitation of existing capabilities. Above, we show how an agent can explore (via drifting along a neutral network), while still maintaining the ability to exploit. Thus neutral networks begin to blur the distinction between exploitation and exploration.

4 Sources of Technological Neutrality

The existence of neutrality in a technological landscape (indeed, on any landscape) hinges on distinct combinatorial configurations having, or being assigned, similar performance measures. It is possible that different “material” technologies, instantiating different engineering and physical-chemical processes, can result in the same functionality. Individual modifications to an existing system design often do not have measurable effects on performance, although in the aggregate such modification can indeed result in improved performance. As a familiar example, take the case of automobile design and how standard features have evolved from bare necessities to include gas gauges, heaters, defrosters, and more recently, ergonomically adjustable

seats, GPS devices, rear-view cameras, and micro-processors to monitor the performance of engine components. None of these features greatly affected performance, but taken together, they can substantially alter the effectiveness of the underlying car design.

Neutrality becomes manifest by the co-existence of different technologies that perform similarly. A well-known example of an old technology keeping up for a long time with a novel one is the prolonged existence of sailing ships next to steamships (Rosenberg, 1972). Similarly, consider that in the late 1950s, three different types of airliner—the propeller driven Douglas DC-7, Lockheed’s turboprop Electra, and the jet-propelled Sud Caravelle, could all be found in commercial service.⁹

Another source of neutrality rests on the fact that assigning a performance measure to a new technology, or evaluating the results of a technological or organizational experiment is, ultimately, an exercise in evaluation, and such evaluations often have a limited resolution. Moreover, evaluations are often subject to error. Time restrictions, managerial constraints, opportunity costs, the effort devoted to quality control, and scientific or technological considerations, may all affect the evaluation of the performance of experiments and introduce neutrality by observation. In the same vein, it is often easier to judge whether a variant performs worse than its predecessor, than it is to decide whether the experiment is an improvement.

5 Modeling Neutrality in a Technological Landscape

5.1 The *NK*-Model

For the purpose of discussing the implications of neutrality on economic search, we use a simple model of technological search as movement on a *NK*-landscape (Kauffman, 1993). The model presented here is a variant of the one developed in Auerswald, Kauffman, Lobo, and Shell (2000) and Kauffman, Lobo, and Macready (2000).

The firm’s *technology* encompasses all of the deliberate organizational, managerial, and technical practices that, when performed together, result in the production of a specific good or service. We assume, however, that technologies are not fully known even to the firms that use them, much less to outsiders. Technologies are assumed to consist of N distinct *processes*

⁹It was not until Boeing’s 707 in 1958 that the technological superiority of jet propulsion for civilian airliners was clearly established.

that exist in S possible types. These types reflect either qualitative (for example, whether to use a conveyor belt or a forklift for internal transport) or quantitative aspects (such as the setting of a dial on a machine), or a mixture of both. A *technological configuration* (technology for short) consists in a specific choice for each of the N processes in a production recipe. Thus, there are S^N possible technologies. A new technology is represented by a change in at least one process within an existing technological configuration. In this framework, technological improvement takes the form of finding technologies with superior performance.

Typically the processes constituting a technology interact, meaning that the consequences of a particular process choice depends on the other processes in the configuration. Interaction among the component processes often leads to a tradeoff between conflicting criteria. For example, the management decision to buy in bulk can lead to decreasing per unit production costs, but also higher warehousing costs. Or consider the use of gas turbines (which are relatively easy to turn on and off) which can make a power grid more flexible, but at the same time more expensive to run. Within a company, giving greater autonomy to design teams can accelerate the rate at which new ideas are generated, but this can also make product design integration more difficult. Designing an aircraft requires positioning the fuel tanks, creating strong but flexible wings, finding engines that are powerful yet quiet, light, and fuel efficient, separating redundant systems, and insuring sufficient passenger capacity. The best solution with regard to one aspect of the design problem often conflicts with the best solution to other aspects.

In the NK -model each process makes a contribution to the overall performance of the technology that depends on its own type and the types of K other processes ($K \leq N - 1$).¹⁰ The parameter K represents the interactions among the processes constituting a technology and therefore determines the level of conflicting constraints a firm has to face when making choices. When $K = 0$ (no interaction), each processes performs independently of all other processes. In contrast, when $K = N - 1$, the performance of each process depends on itself and all other processes. In this way, K determines the correlation structures of the landscape. When K equals zero, the landscape has a single, smooth-sided, peak. As K increases, the landscape becomes

¹⁰The S^{K+1} possible contributions to total performance made by the j th process are treated as i.i.d. random variables drawn from some distribution F . In the numerical results presented below, F is $U(0, 1)$, although our results are insensitive to this choice.

more rugged. A change in a process affects the contributions of several other processes, causing neighboring technologies to have different performances values. When $K = N - 1$, the landscape is maximally rugged. The correlation coefficient is given by (Weinberger, 1990; Fontana *et al.*, 1993)

$$\rho(d) = \left[1 - \frac{d}{N}\right] \left[1 - \frac{K}{N-1}\right]^d, \quad (1)$$

and thus for $d = 1$ and large N , $K = 0$ implies $\rho(1) \approx 1$ and $K = N - 1$ has $\rho(N) = 0$.

5.2 Neutrality in an NK -World

The raw NK -model has no neutrality. We introduce neutrality by discretizing the continuous performance values produced by the original model. In the standard NK -model, any given configuration on the landscape is assigned a performance (fitness) value between 0.0 and 1.0. Here we collect these continuous values into M bins, with the fitness values between $[0.0, 1/M)$ being mapped to bin 1, $[1/M, 2/M)$ being assigned to bin 2, and so on.¹¹ This simple procedure implements a notion of neutrality in which configurations with closely related fitness values are hard to distinguish from one another.

Intuition about the behavior of the model with neutrality can be gained through the following thought experiment (see Figure 2). Consider taking a thin vertical slice out of a mountain range composed of evenly-distributed, horizontal layers of rock. As we increase the value of M , we increase the number of layers that make up the rock (by decreasing the height of each layer). As M goes to infinity the layers become incredibly thin and we are back in the standard NK -model.

Agents in this world wander the surface of the mountain range in search of increasing heights. A hill-climbing searcher is willing to randomly wander on the surface as long as it is in its current layer of rock (implying neutrality). If during these travels it hits another layer of rock, it will always enter higher layers and always avoid lower ones. If M is very low, then almost all configurations (that is, locations on the surface) become indistinguishable, and the agent ends up wandering randomly across the entire surface. The key question we wish to answer is how do increasing levels of neutrality (decreasing M) affect search on landscapes of differing ruggedness.

¹¹Newman and Englehardt (1998) use a related approach to impose neutrality by discretizing the weights that enter into the fitness calculation.

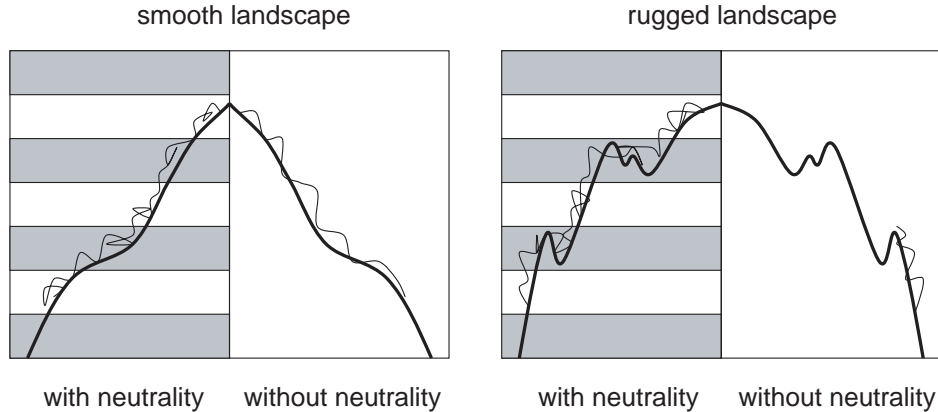


Figure 2: An illustration of the interaction between neutrality (left-hand side of each panel) and ruggedness.

First, consider a smooth, single peak landscape (left panel of Figure 2). If the mountain has a large number of layers (high M , that is, low neutrality), a hill-climbing searcher will quickly ascend to, and remain at, the top of the mountain. As we increase the size of the layers (that is, decrease M or increase neutrality), agents will have a harder time ascending the heights, as they will often be wandering up and down in their current layer. They will, however, eventually ascend to the uppermost layer, as their random wandering within any given layer will at some point put them in contact with the next higher layer. Once they achieve the highest layer, they will randomly wander within it. Thus, we would predict that under smooth landscapes, increasing neutrality will slow adaptation (due to excessive drift) and the final outcome will be close to the highest value possible on the landscape within an error dictated by the height of the final layer.

Next, consider a rugged, multi-peaked landscape (right panel of Figure 2). If the landscape is composed of a large number of layers (that is, has low neutrality), then a hill-climbing agent will quickly ascend, and be trapped at, a local peak. Note that even though higher peaks may exist, the local nature of the agent's adaptation prevents it from jumping across intervening valleys to get onto more fruitful slopes. As we increase neutrality, it is possible for a given layer of the mountain to encompass such valleys. Thus, an agent wandering around in such a layer will be able to bridge the previously impenetrable valley and get onto the better slopes. Thus, we would predict

Neutrality	Landscape Type	
	Smooth	Rugged
Low	Best Outcome, Quick Adaptation	Poor Outcome, Quick Adaptation
High	Good Outcome, Slow Adaptation	Good Outcome, Slow Adaptation

Table 1: Predicted relationship of adaptive behavior given neutrality and landscape ruggedness.


that with rugged landscapes, increasing neutrality will also slow adaptation, yet (unlike before) allow agents to achieve better outcomes than they otherwise would, even when we take into account the inherent noise imposed by the agent randomly wandering in the final layer.

Table 1 summarizes the above predictions. In general, as neutrality increases, we expect adaptation to slow as the agents’ probability of wandering within a fitness layer increases. If landscapes are smooth, we would expect neutrality to hamper the discovery and maintenance of good solutions. On the other hand, when landscapes are rugged, neutrality should allow agents to traverse previously impassable valleys and, in so doing, find better solutions.


We test the above predictions through a series of numerical experiments. In these experiments, we simulate an individual agent’s search as an adaptive walk on NK -landscapes of varying ruggedness and neutrality. Landscapes had $N = 100$ and $S = 3$, and K was chosen to produce four different correlations ($\rho = 0.2, 0.3, 0.5, 0.7$). We considered four different levels of the neutrality-tuning parameter M . A thousand searches, with randomly assigned initial locations, were simulated for each combination of parameters, and the reported results are the averages over the thousand runs.

Figure 3 presents the results. On smooth or somewhat rugged landscapes, the presence of neutrality does not improve on the performance of a simple adaptive walk. Indeed, agents adapt slower and tend to end up at inferior levels of fitness. However, on very rugged landscapes ($\rho = 0.2$), neutrality results in superior outcomes. On such landscapes, an intermediate level of neutrality tends to have the best tradeoff between speed of adaptation and ultimate fitness.


Thus, the introduction of neutrality fundamentally alters our view of search in NK -models. Prior work that applied these models to organiza-




(a) $\rho = 0.7$



(b) $\rho = 0.5$



(c) $\rho = 0.3$



(d) $\rho = 0.2$

Figure 3: Average fitness per iteration as a function of landscape correlation (ρ) and level of neutrality (M) versus a simple adaptive walk (SAW).

tional and technological innovation has emphasize how the tunable correlation structure of these models impacts the traditional notion of ruggedness, and hence search. Neutral networks, however, can completely change the character of a landscape without affecting its correlation structure. These networks allow extensive “ridge” systems to form that can be traversed by suboptimal, local wanderings.

Neutrality depends upon both the interdependence among the component parts of a configuration (tuned by the K parameter in the NK -model) and the connectivity of neutral configurations into extensive percolating paths. Decreasing the number of evaluation bins obviously increases the number of neutral configurations. This, however, does not imply that those neutral configurations form into connected path that percolates across the landscape. In fact, in the present model high “ruggedness” (in correlation terms) of the landscape may actually prevent the formation of such extended and connected neutral networks. Thus, the impacts of neutrality we see above are likely to be attenuated over alternative models (Newman and Engelhardt, 1998) that more explicitly create percolating, neutral networks.

6 Technological Neutrality in Software and Hardware

Examples of the potential of technological neutrality can be drawn from a variety of industries. Here, we focus on a few from the information technology sector, an area of industrial activity that has become increasingly important over the last few decades.

As software developers attempt to create ever larger programs, it is increasingly being recognized that there is a brewing “complexity crisis.” The sheer size of current software, linked with the necessary interactions among various code elements, make programs inherently difficult to create, maintain, and modify.

Software development offers an interesting example of neutral technological networks. If we take the source code as the “genotype” and the way the program interacts with users, input, and output as the “phenotype,” then it is easy to see how a variety of genotypes can lead to the identical phenotype.

Software “refactoring” (see, for example, Fowler *et al.*, 1999) is a process in which the underlying source code is rewritten to improve its design without altering the program’s behavior. Refactoring is intended to be neutral since the genotypic changes should not alter the phenotype. However, neutral refactoring does improve the evolvability of the software. As code

is refactored, key elements of the program's structure are better formulated, code is simplified, and becomes easier to understand. In this way, code acquires a new form that may lead to the discovery of opportunities for change or new functionality. Once fully refactored, the code can more readily adapt to various "environmental" changes, like new operating systems, user needs, etc.

Oftentimes code evolvability becomes a goal intentionally pursued by software developers. In those circumstances, evolvability has become an explicit trait of the code's "phenotype." As a consequence, changes that were neutral prior to the declaration of that goal are no longer neutral. Once a feature has become the target of intentional action, it comes under selection pressure. Nonetheless, at early stages of software development or in very large code bases, overall evolvability is desirable but not specifically designable; It is likely to suddenly, and unintentionally, arise by the accumulation of changes that are valued as neutral by the developer. Furthermore, even when evolvability has become a declared feature of design, there are still many *ex ante* equivalent ways of achieving that goal and neutrality is still a factor in shaping that landscape.

The attempt of making software more evolvable (which means intelligible across programmers, extensible, and robust) has led to design elements like encapsulation that hide implementation details behind public interfaces. While enabling intentional extensions of functionality like plug-ins, such design elements also increase the overall neutrality of the program, and facilitate tinkering that is not goal-oriented but of a more experimental character.

Consider, for example, the following common development scenario from the world of Java. A formal request for changes to the Java language is first made through a Java Specification Request (JSR), which is a description, in very general terms, of a set of functionalities for accomplishing a certain task. The JSR is then followed by an Application Programmer Interface (API), a more detailed specification of the set of functionalities to accomplish the desired task. Developers then design implementations that instantiate the API, and there can be, and usually are, a variety of different implementations. Initially, several different implementations, with similar performance characteristics, of the same API are created and promulgated among users. The choice of one implementation over another is often a matter of aesthetic, and not strict performance, criteria.

Christensen's (1997) work provides some other examples of neutral technologies in the computer industry. Consider the average data recording

density achieved by different head and disk-drive technologies. During the late 1980's, particulate-oxide-disk technology and ferrite-head technology had comparable densities, while in the mid 1990's thin-film heads and magneto-resistive heads were comparable. As discussed by Christensen, these technologies went on to experience very different product development trajectories. Similar performance, as measured by storage capacity, were simultaneously achieved by four different rigid-disk-drive technologies (14-, 8-, 5.25-, and 3.5-inch drives) during the fifteen year span between 1975 and 1990 (Christensen 1993). Each of these technologies arose via different developmental trajectories. Laser and ink-jet printer technologies (originally developed by Hewlett Packard) are quite distinct, and each has experienced different rates of development. Yet, it is currently possible to find pairs of laser and ink-jet printers (manufactured by HP, Cannon and IBM) of comparable performance, as measured by printing speed, print quality, and price. In all of these examples, drift among functionally equivalent technologies often leads to accumulated changes by which small changes in one of the technologies secures a major innovation and shift in market leadership.

7 Conclusion

Technological and organizational innovation is a key feature in social systems. The introduction of the biological notion of a fitness landscape provided a useful substrate by which to model the process of complex search in such domains. Work in this area has shown how the spatial correlation of the fitness landscape can be tied to a notion of ruggedness, and that search agents will tend to get trapped at inferior outcomes on more rugged landscapes.

Here, we extend this view of search by incorporating recent work in biology on the idea of neutral networks. Neutral networks form pathways of temporarily inconsequential changes that, while *a priori* may be thought of as at best non-productive and at worst a serious nuisance or distraction, fundamentally alter the potential of search in such worlds. In particular, muddling through the landscape on neutral networks ultimately improves the potential of the system for innovative leaps, and offers an escape route by which to avoid the usual pitfalls of ruggedness.

Neutral networks form an extensive, connected system of ridges along which search agents cannot sense a difference in performance. This allows agents to drift on the network across the search space without worsening (or improving) their currently best solution. In so doing, however, agents can

reach new, far-away configurations that may yield productive, previously inaccessible, innovations. Such networks are not symmetric, in the sense that it may be easier to transit from one organizational form to another, than the other way around. Moreover, neutrality is always evaluated within a given search context. Thus, the same change that keeps software neutral in the sense that from the user’s perspective it performs the same, may have been non-neutral from the perspective of improving the ability of future programmers to modify and maintain the code.

As demonstrated above, we can build in notions of neutrality into standard models of fitness landscapes. Here we did so just to illustrate the underlying ideas, and we suspect that further theoretical investigations of this phenomenon would be fruitful. In particular, having models with more direct ways to form connected neutral pathways may be useful. We also examined some examples from software and manufactured goods where neutrality may have arisen in industry. Obviously, many more such examples can be found, and it is likely that more extensive case studies would be of interest.

The presence of neutrality provides two important, and somewhat paradoxical, features to organizational and technological systems. First, it gives such systems a degree of robustness against tinkering. Indeed, without such robustness it is hard to imagine how such systems could be productively embraced. Second, as shown above, neutrality is also a crucial enabler of change and innovation. Thus, neutrality is able to simultaneously convey both robustness and innovation—two features that are of fundamental importance to complex systems.

References

- [1] Abernathy, W. J. and K. Clark (1985) "Innovation: Mapping the Winds of Creative Destruction," *Research Policy*, 14, 3-22.
- [2] Anderson, P. W. (1972) "More is different", *Science*, 177, 393-396.
- [3] Anderson, P. and M. L. Tushman (1990) "Technological Discontinuities and Dominant Designs: A Cyclical Model of Technological Change," *Administrative Science Quarterly*, 35, 604-633.
- [4] Arrow, K. J. (1974) *The Limits of Organization*. New York: W.W. Norton & Company.
- [5] Ashmos, D. P., D. Duchon, and R. R. McDaniel, Jr. (1998) "Participation in Strategic Decision Making: The Role of Organizational Predisposition and Issue Interpretation," *Decision Sciences*, 2, 25-51.
- [6] Auerswald, P., S. Kauffman, J. Lobo, and K. Shell (2000) "The Production Recipes Approach to Modeling Technological Innovation: An Application to Learning by Doing," *Journal of Economic Dynamics and Control*, 24, 389-450.
- [7] Barney, J. B. (1991) "Firm Resources and Sustained Competitive Advantage," *Journal of Management*, 17, 99-120.
- [8] Basalla, G. (1988) *The Evolution of Technology*. Cambridge: Cambridge University Press.
- [9] Baumol, W. J. (2002) *The Free-Market Innovation Machine*. Princeton: Princeton University Press.
- [10] Beinhocker, E. (1999) "Robust Adaptive Strategies," *Sloan Management Review*, 40, 95-106.
- [11] Bikhchandani, S. and S. Sharma (1996) "Optimal Search with Learning," *Journal of Economic Dynamics and Control*, 20, 333-359.
- [12] Boeker, W. (1989) "Strategic Change: The Effects of Founding and History," *Academy of Management Journal*, 32, 489-515.

- [13] Brusoni, S. and A. Prencipe (2001) “Unpacking the Black Box of Modularity: Technologies, Products and Organizations,” *Industrial and Corporate Change*, 10, 179-205.
- [14] Caselli, F. (1999) “Technological Revolutions,” *American Economic Review*, 89, 78-102.
- [15] Christensen, C. (1993) “The Rigid Disk Drive Industry: A History of Commercial and Technological Turbulence,” *Business History Review*, 67, 531-588.
- [16] Christensen, C. (1997) *The Innovator’s Dilemma: When New Technologies Cause Great Firms to Fail*. Boston: Harvard Business School Press.
- [17] Cohen, W. M. and D. A. Levinthal (1989) “Innovation and Learning: The Two Faces of R&D,” *Economic Journal*, 99, 569-596.
- [18] Collins, J. C. and J. I. Porras (1994) *Built to Last: Successful Habits of Visionary Companies*. New York: Harper Business.
- [19] Constant II, E. W. (2002) “Why Evolution is a Theory About Stability: Constraint, Causation, and Ecology in Technological Change,” *Research Policy*, 31, 1241-1256.
- [20] Dalle, J. M. (1997) “Heterogeneity vs. Externalities in Technological Competition: A Tale of Possible Technological Landscapes,” *Journal of Evolutionary Economics*, 7, 395-413.
- [21] Dosi, G. (1984) *Technical Change and Economic Transformation*. London: Macmillan Books.
- [22] Dörner, D. (1996) *The Logic of Failure: Recognizing and Avoiding Error in Complex Situations*. Cambridge: Perseus Books.
- [23] Ebeling, W, Karmeshu, and A. Scharnhorst (2001) “Dynamics of Economic and Technological Search Processes in Complex Adaptive Landscapes,” *Advances in Complex Systems*, 4, 71-88.
- [24] Ericson, R. and A. Pakes (1995) “Markov-Perfect Industry Dynamics: A Framework for Empirical Work,” *Review of Economic Studies*, 62, 53-82.

- [25] Evenson, R. E. and Y. Kislev (1976) "A Stochastic Model of Applied Research," *Journal of Political Economy*, 84, 265-281.
- [26] Fleming, L. (2001) "Recombinant Uncertainty in Technological Search," *Management Science*, 47, 117-132.
- [27] Fleming, L. and O. Sorenson (2001) "Technology as a Complex Adaptive System: Evidence from Patent Data," *Research Policy*, 30, 1019-1039.
- [28] Fleming, L. and O. Sorenson (2003) "Navigating the Technology Landscape of Innovation," *Sloan Management Review*, 44, 15-23.
- [29] Fontana, W. and P. Schuster (1998) "Continuity in Evolution: On the Nature of Transitions," *Science*, 280, 1451-1455.
- [30] Fontana, W., P. F. Stadler, E. G. Bornberg-Bauer, T. Griesmacher, I. L. Hofacker, M. Tacker, P. Tarazona, E. D. Weinberger, and P. Schuster (1993) "RNA Folding and Combinatory Landscapes," *Physical Review E*, 47, 2083-2099.
- [31] Fowler, M., K. Beck, J. Brant, W. Opdyke, and D. Roberts (1999) *Refactoring: Improving the Design of Existing Code* Addison-Wesley.
- [32] Freeman, C. (1982) *The Economics of Industrial Innovation*. Cambridge, MA: The MIT Press.
- [33] Freeman, C. (1994) "Critical Survey: The Economics of Technical Change," *Cambridge Journal of Economics*, 18, 463-514.
- [34] Hannan, M. T. and J. Freeman (1984) "Structural Inertia and Organizational Change," *American Sociological Review*, 49, 149-164.
- [35] Haslett, T. and C. Osborne (2003) "Local Rules: Emergence on Organizational Landscapes," *Nonlinear Dynamics, Psychology, and Life Sciences*, 7, 87-98.
- [36] Helfat, C. E. (1994) "Firm Specificity and Corporate Applied R&D," *Organization Science*, 5, 173-184.
- [37] Henderson, R. M. and K. B. Clark (1990) "Architectural Innovation: The Reconfiguration of Existing Product Technology and the Failure of Established Firms," *Administrative Science Quarterly*, 35, 9-31.

- [38] Herriott, S. R., D. A. Levinthal, and J. G. March (1985) "Learning from Experience in Organizations," *American Economic Review*, 75, 298-302.
- [39] Hey, J. D. (1982) "Search for Rules of Search," *Journal of Economic Behavior and Organization*, 3, 65-81.
- [40] Hoppe, H. C. (2000) "A Strategic Search Model of Technology Adoption and Policy," *Industrial Organization*, 9, 197-214.
- [41] Huynen, M. A., P. F. Stadler, and W. Fontana (1996) "Smoothness Within Ruggedness: The Role of Neutrality in Adaptation," *Proceedings of the National Academy of Sciences USA*, 93, 397-401.
- [42] Jovanovic, B. (1982) "Selection and the Evolution of an Industry," *Econometrica*, 50, 659-670.
- [43] Jovanovic, B. and R. Rob (1990) "Long Waves and Short Waves: Growth Through Intensive and Extensive Search," *Econometrica*, 58, 1391-1409.
- [44] Iansiti, M. and T. Khanna (1995) "Technological Evolution, System Architecture and the Obsolescence of Firm Capabilities," *Industrial and Corporate Change*, 4, 333-361.
- [45] Kauffman, S. (1993) *Origins of Order: Self-Organization and Selection in Evolution*. New York: Oxford University Press.
- [46] Kauffman, S. (1995) "Escaping the Red Queen Effect," *The McKinsey Quarterly*, 1, 119-129.
- [47] Kauffman, S., J. Lobo, and W. Macready (2000) "Optimal Search on a Technology Landscape," *Journal of Economic Behavior and Organization*, 43, 141-166.
- [48] Kennedy, P. M. (1994) "Information Processing and Organizational Design," *Journal of Economic Behavior and Organization*, 25 37-51.
- [49] Kimura, M. (1968) "Evolutionary Rate at the Molecular Level," *Nature*, 217, 624-626.
- [50] Kimura, M. (1983) *The Neutral Theory of Molecular Evolution*. Cambridge, UK: Cambridge University Press.

- [51] Klepper, S. (1996) "Entry, Exit, Growth and Innovation Over the Product Life Cycle," *American Economic Review*, 86, 562-583.
- [52] Kohn, M. and S. Shavell (1974) "The Theory of Search," *Journal of Economic Theory*, 9, 93-123.
- [53] Lee, D. M. and T. J. Allen (1982) "Integrating New Technical Staff: Implications for Acquiring New Technology," *Management Science*, 28, 1405-1420.
- [54] Levinthal, D. A. (1997) "Adaptation on Rugged Landscapes," *Management Science*, 43, 934-950.
- [55] Levinthal, D. A. and J. G. March (1981) "A Model of Adaptive Organizational Search," *Journal of Economic Behavior and Organization*, 2, 307-333.
- [56] Levinthal, D. and M. Warglien (1999) "Landscape Design: Designing for Local Action in Complex Worlds," *Organization Science*, 10, 342-357.
- [57] Levitan, B, J. Lobo, R. Schuler, and S. Kauffman (2002) "Evolution of Organizational Performance and Stability in a Stochastic Environment," *Computational & Mathematical Organization Theory*, 8, 281-313.
- [58] Loasby, B. J. (2001) "Time, Knowledge, and Evolutionary Dynamics: Why Connections Matter," *Journal of Evolutionary Economics*, 11, 393-412.
- [59] Loasby, B. J. (2002) "The Evolution of Knowledge: Beyond the Biological Model," *Research Policy*, 31, 1227-1239.
- [60] Macken, C. A., P. S. Hagan, and A. S. Perelson (1991) "Evolutionary Walks on Rugged Landscapes," *SIAM Journal of Applied Mathematics*, 51, 799-827.
- [61] Macken, C. A. and P. F. Stadler (1995) "Evolution on Fitness Landscapes," in Nadel, L. and D. Stein, eds., *1993 Lectures in Complex Systems*. Reading, MA: Addison-Wesley Publishing Company.
- [62] Mahdi, S. (2003) "Search Strategy in Product Innovation Process: Theory and Evidence from the Evolution of Agrochemical Lead Discovery Process," *Industrial and Corporate Change*, 12, 235-270.

- [63] March, J. G. (1991) "Exploration and Exploitation in Organizational Learning," *Organization Science*, 2, 71-87.
- [64] Marengo, L. (1992) "Coordination and Organizational Learning in the Firm," *Journal of Evolutionary Economics*, 2, 313-326.
- [65] Maynard-Smith, J. (1970) "Natural Selection and the Concept of a Protein Space," *Nature*, 255, 563-564.
- [66] McCarthy I. P. and Y. K. Tan (2000) "Manufacturing Competitiveness and Fitness Landscape Theory," *Journal of Materials Processing Technology*, 107, 347-352.
- [67] McKelvey, B. (1999) "Avoiding Complexity Catastrophe in Coevolutionary Pockets: Strategies for Rugged Landscapes," *Organization Science*, 10, 294-321.
- [68] Mowery, D. and N. Rosenberg (1998) *Paths of Innovation: Technological Change in 20th Century America*. New York: Cambridge University Press.
- [69] Muth, J. F. (1986) "Search Theory and the Manufacturing Progress Function," *Management Science*, 32, 948-962.
- [70] Nelson, R. R. and S. G. Winter (1982) *An Evolutionary Theory of Economic Change*. Cambridge, MA: Belknap Press.
- [71] Newman, M. E. J. and R. Engelhardt (1998) "Effects of Neutral Selection on the Evolution of Molecular Species", *Proceedings of the Royal Society of London B*, 256, 1333-1338.
- [72] Perrow, C. (1999) *Normal Accidents: Living with High Risk Technologies*. Princeton: Princeton University Press.
- [73] Potts, J. (2001) *The New Evolutionary Microeconomics: Complexity, Competence and Adaptive Behavior*. Northampton: Edward Elgar.
- [74] Provine, W. B. (1986) *Sewall Wright and Evolutionary Biology*. Chicago: University of Chicago Press.
- [75] Rivkin, J. and Siggelkow, N. (2002) "Organizational Sticking Points on NK Landscapes," *Complexity*, 7, 31-43.

- [76] Romer, P. M. (1990) “Endogenous Technological Change,” *Journal of Political Economy*, 98, 71-102.
- [77] Romer, P. M. (1996) “Why, Indeed, in America? Theory, History, and the Origins of Modern Economic Growth,” *American Economic Review*, 86, 202-206.
- [78] Rosenberg, N. (1972) *Technology and American Economic Growth*. White Plains: Sharp Company.
- [79] Rosenberg, N. (1982) *Inside the Black Box: Technology and Economics*. Cambridge: Cambridge University Press.
- [80] Rosenberg, N. and L.E. Birdzell, Jr. (1986) *How the West Grew Rich: The Economic Transformation of the Industrial World*. New York: Basic Books.
- [81] Sagan, S.D. (1993) *The Limits of Safety: Organizations, Accidents, and Nuclear Weapons*. Princeton: Princeton University Press.
- [82] Sahal, D. (1981) *Patterns of Technological Innovation*. Reading: Addison-Wesley.
- [83] Sargent, T. J. (1987) *Dynamic Macroeconomic Theory*. Cambridge, MA: Harvard University Press.
- [84] Schmookler, J. (1966) *Invention and Economic Growth*. Cambridge: Harvard University Press.
- [85] Schumpeter, J. A. (1942) *Capitalism, Socialism and Democracy*. New York: Harper & Brothers.
- [86] Schuster, P., W. Fontana, P.F. Stadler, and I. Hofacker (1994) “From Sequences to Shapes and Back: A Case Study in RNA Secondary Structures,” *Proceedings of the Royal Society (London) B*, 255, 279-284.
- [87] Shan, W. (1990) “An Empirical Analysis of Organizational Strategies by Entrepreneurial High-Technology Firms,” *Strategic Management Journal*, 11, 129-139.
- [88] Stadler, P. F. (1992) “Correlation in Landscapes of Combinatorial Optimization Problems,” *Europhysics Letters*, 20, 479-482.

- [89] Stadler, P. F. (1995) "Towards a Theory of Landscapes," Social Systems Research Institute Working Paper Number 9506. Madison: University of Wisconsin.
- [90] Stadler, M. R., P. F. Stadler, G. P. Wagner, and W. Fontana (2001) "The Topology of the Possible: Formal Spaces Underlying patterns of Evolutionary Change," *Journal of Theoretical Biology*, 213, 241-274.
- [91] Stuart, T. E. and J. M. Podolny (1996) "Local Search and the Evolution of Technological Capabilities," *Strategic Management Journal*, 17, 21-38.
- [92] Tesler, L. G. (1982) "A Theory of Innovation and its Effects," *The Bell Journal of Economics*, 13, 69-92.
- [93] Tushman, M. L. and P. Anderson (1986) "Technological Discontinuities and Organizational Environments," *Administrative Science Quarterly*, 14, 311-347.
- [94] Wagner, G. and L. Altenberg (1996) "Complex adaptations and the evolution of evolvability," *Evolution*, 50, 967-976.
- [95] Weinberger, E. D. (1990) "Correlated and Uncorrelated Fitness Landscapes and How to Tell the Difference," *Biological Cybernetics*, 63, 325-336.
- [96] Weitzman, M. L. (1979) "Optimal Search for the Best Alternative," *Econometrica*, 47, 641-654.
- [97] Weitzman, M. L. (1996) "Hybridizing Growth Theory," *The American Economic Review*, 86, 207-212.
- [98] Weitzman, M. L. (1998) "Recombinant Growth," *The Quarterly Journal of Economics*, 113, 331-360.
- [99] Westhoff, F. H., B. V. Yarbrough, and R. M. Yarbrough (1996) "Complexity, Organization, and Stuart Kauffman's The Origins of Order," *Journal of Economic Behavior and Organization*, 29, 1-25.
- [100] Wright, S. (1932) "The Roles of Mutation, Inbreeding, Crossbreeding and Selection in Evolution," in D.F. Jones, editor, *Proceedings of the Sixth International Congress on Genetics, Vol. 1*.