

Modeling Directed Local Search Strategies on Technology Landscapes

Draft

Karén Hovhannisian

**ROCK, Department of Management and Computer Science,
University of Trento**

Email : khovhann@economia.unitn.it

June, 2004

Introduction

Simon (1962) defines a complex system as one that consists of many elements that interact in a non-trivial way. A strategy of a firm then is certainly complex, since it encompasses a number of decisions and since those decisions are interdependent.

Complexity poses a problem of conflicting constraints in designing the optimal strategy, and has been in the focus of economic analysis at least for the last half a century. One of the most extensively studied methods of dealing with complexity, and the one that dates back to Simon's work, is the decomposition of the strategy into quasi-independent modules.

On the organizational level such decomposition would result in division of labour between different units and departments, while on the technological level it would mean independent design of the constituent parts of an artefact.

Modularization of the problem structure has been shown to be an impressive tool for solving complex problems (classical text is Simon 1969, recent contributions include Frenken *et al.* 1999, Baldwin & Clark 2000, Fleming & Sorenson 2001a, Ethiraj & Levinthal 2004). However, there are several serious problems or *dangers of modularity*.

The most obvious one is that complex structures might well just not be decomposable or even near-decomposable, so that the agents would be unable to divide the problem into *patches* or else find valid *interfaces* to start with. Any decomposition in this case would lead to neglecting some fraction of interconnections within the system, and depending on how important those links are, the efficiency achieved might be quite substantially lower than desired.

Moreover, while it is generally true that the average solution of a more complex problem would be worse than that of a simpler one, it is just as true that the relation would be turned upside down when we look at the best possible solution of the respective problems (Fleming & Sorenson 2001ab, 2003 provide the intuition behind this result, Evans & Steinsaltz 2002, Durrett & Limic 2003 provide its formal proof; see also *Figure 7* for simulation analysis results).

More complex settings provide more opportunities in which elements can influence each other, and we do need to remember that despite the fact that more attention has been given to the possible *negative* feedback loops between the elements, modularization of a landscape discards of just as large a fraction of possible *positive* feedbacks. Thus, *ceteris paribus*, if we want to achieve a higher efficiency level in absolute terms (rather than as a fraction of a globally optimal result), more complex landscape is simply a better place to look for it. As Fleming and Sorenson (2001b) point

out: “By placing a premium predictability in their product development efforts, companies create a technology landscape that’s easier to navigate – but one that may produce fewer breakthroughs”¹

Another interesting insight into dangers of a modular design has been provided by Rivkin (2000), who claims that: “As the elements of a firm’s decision problem grow numerous and interdependent, imitation of a successful strategy becomes very difficult. Indeed it can become “intractable” in a technical sense of the word”.² This would suggest that strategies (or technologies) that are more complex are *naturally codified*, and hence their use *hedges* the firms against the risk of being imitated by the competitors.

Finally, as found by Ethiraj and Levinthal (2004): “[...] in the long run, erring on the side of greater integration poses lower performance penalties than erring on the side of greater modularity”³ This suggests that when the underlying structure of the landscape is not perfectly known to a decision maker, and thus the correct decomposition is not attainable, choosing a more complex strategy may well be a better idea than choosing a less complex one.

Summing up all the abovementioned a conclusion can be made that while to a certain degree modularization is a valid strategy of dealing with complexity, it should be used with great caution. Moreover, and what this paper will deal with, the “correct” decomposition will still leave the resulting “blocks” rather large and the elements within them interrelated, so that alternative ways of dealing with complexity have to be sought.

Modularization can be seen as an *objective* way to simplify the problem to be solved. The problem itself changes as a result, and the new problem that we obtain is objectively simpler both for the decision-maker herself, and for the others. Alternatively, in order to still be able to reap the fruits of a more complex structure and prevent the others to imitate it, a *subjective* way to make the problem simpler have to be designed, a way that doesn’t change the problem itself, but our knowledge and abilities to solve it.

As a framework of the analysis we take Stuart Kauffman’s model of fitness landscapes (Kauffman & Levin 1987, Kauffman 1993). The original model deals with the evolution of coupled natural (biological) systems. It is therefore plausible to assume absence of any *strategic intentions* or *foresight* possessed by the agents of the system. However, this is no longer a valid assumption when a transition from natural to social evolution analysis is being made.

¹ Fleming, Lee & Olav Sorenson (2001b) *The Dangers of Modularity*. Harvard Business Review, September, pages 20-21

² Rivkin, Jan W. (2000) *Imitation of Complex Strategies*. Management Science, vol. 46, No. 6 (June), page 825

³ Ethiraj, Sendil K. & Daniel Levinthal (2004) *Modularity and Innovation in Complex Systems*. Management Science, vol. 50, No. 2 (February), page 169

Take a game of chess⁴ as an example of a complex coupled system. Each figure on the board can be placed in a number of different positions, with the “value” of a given figure being in a given position depending on where the rest of the figures are placed. What the model in its original setting, as well as its current applications to the economics proper (see e.g. Levinthal 1997, Frenken *et al.* 1999, Kauffman *et. al.* 2000) assume, is that a player randomly chooses one (and only one) figure on the board at a time and decides whether to move it to a new position or not possessing perfect information of the consequences of that move for the next period, while at the same time being completely ignorant on how that move influences the options a player will have when she is called upon to make subsequent moves.

This might be the “strategy” of someone who has never heard of the game before, but I’d be very skeptical to take it as a reasonable approximation for a firm’s perception of the technological evolution it undergoes.

Even quite a bad chess player would consider several possible moves at a time, and, a relatively better one would also be able to think several steps ahead when making the decision on the move at the current period. Indeed, this is exactly what differentiates a good from a bad player. The game of chess remains the same, the problem itself does not become less complex, but, subjectively, it is simpler for a more experienced player than for the novice precisely for the reason that the former is more able to think in *breadth* and in *depth*. Experience of a player allows her to draw a better “cognitive map” of the problem.⁵

Apart from the mere reason of implausibility (after all a model is not supposed to replicate the real world perfectly), there are much more grounded justifications of enriching the original model in the proposed manner. Introducing depth and breadth of search opens up a way to consider and compare the whole spectrum of local search strategies from extremely basic and myopic to overwhelmingly sophisticated, and measure in effect the level of sophistication needed in order to solve problems of different complexity. Moreover, when strategic intent enters the picture, some of the most important conclusions drawn from the analysis of the economic applications of the model in its original setting fail to hold.

The basic insight gained from the analysis of the *NK*-based models of local search⁶ is that for at least partially interconnected systems that are characterized by the existence of multiple

⁴ Think of it as a variation of a chess game played against nature, rather than an opponent

⁵ See Gavetti & Levinthal 2000 and below for more detailed review of their model.

⁶ The analysis is not limited to the scenario of a search being strictly local, however it has been shown that for the tightly coupled systems the so-called long jumps (changing the state of several elements at a time) are beneficial only for low initial efficiency levels. In the current paper it is therefore assumed that the firms employ only local search strategies.

equilibria, the agents are getting trapped on local efficiency maxima, hardly ever reaching the global optimum. Alternatively it is shown that exhaustive search through the whole space of possible configurations (something only a perfectly rational agent can undertake) guarantees reaching the global optimum, but, due to an astronomic number of such configurations on a multidimensional landscape, is extremely costly.

Bringing the measures of *depth* and *breadth* of search into the analytical framework, the current paper investigates the whole space of the possible strategies of local search between the myopic search scenario, as in the original *NK model* ($breadth=depth=1$), and the perfect foresight scenario ($breadth=depth=N$). The main point of the current research is to show that while quite obviously perfect foresight is a sufficient condition for the attainment of the globally optimal solution, it is by no means a necessary condition for that.

2. Breadth and Depth of Search. Preliminaries

Much of the existing applications of the NK Model of Fitness Landscapes in the fields of strategy and organizational design center around the idea of decomposability and creation of modular quasi-independent structures. Despite being an impressive analytical tool, however, modularization has been shown to have certain limits. Such limits can usually be reached quite quickly, leaving the resulting block-structures still too complex for a usual myopic and purely experiential trial and error strategy. This would lead the agents to reach solutions that while being possibly better than those obtainable with leaving the problem structure purely integral, are still inferior in most of the cases to the globally best solution existing on the landscape.

We can see the whole problem of a decision maker as consisting of two parts: (a) finding the optimal level of decomposition, thus obtaining the block-structures of optimal size and complexity, and, (b) finding a strategy of search that given the resulting characteristics of the block-structures obtained, would lead us to an optimal solution of the problem.

While much attention has been given lately to the first part of the problem, the second part of it is clearly underinvestigated. The task of the current paper is precisely in attempting to close that gap. I assume here that an optimally sized block-structures have already been obtained in the ways proposed e.g. by Frenken *et al* (1999) or Ethiraj & Levinthal (2004), and having that assumption as a starting point of the analysis, extend the model to find the simplest strategy that leads the agents to the global optimum in shortest time and with least possible requirements on the agents' rationality level.

2.1. Breadth of Search on Technology Landscape

The idea of parallelism of search has been analyzed quite extensively in prior research both within mainstream economics (Vishwanath 1988, 1992) and in its evolutionary branch (Kauffman & Macready 1995). In fact, the idea has been analyzed deeper still in the field of genetic algorithm based evolutionary programming (Azencott 1992, Macready *et al.* 1996).

Despite such a wide field of application, there is one particular nuance that at least to my knowledge has been left unchallenged in formal modeling exercises. It is always the case that parallelism refers to the actual moves made, and hence, there is an acquired necessity to treat parallel search as non-local in nature. However, implicit in the models is the fact that at each step of the process of evolutionary search agents undergo two distinct stages: firstly they evaluate the value added of a possible change, and only secondly, if the analysis proves the change to be beneficial, they actually do move. A *move* here is understood to be the action of changing the state of an element. So by a parallel search it is usually meant that the state of more than one element is attempted to be changed at a single step.

Intuitively this should lead to a faster rate of adaptation, but there is a significant problem embedded in such strategy. While at the stage of *evaluation* changes are considered one-by-one (in parallel but independently), at the stage of the actual *move* the changes are made together (simultaneously). In consequence, a situation might arise that due to internal connectivity of the system, while a change in any of those elements' states leads to a higher overall efficiency given the other elements' states are left unchanged, this is no longer true when those other elements' states do change. Thus, the information that is used for making a decision on whether or not to *flip* the state of some particular element becomes outdated in the presence of parallel moves.

An alternative way is proposed here. As in the parallel search scenario several possible changes are considered one-by-one on the *evaluation* stage of the process. However, at the stage when the move is being made, the state of only one single element is being changed.

So, the parallelism implicit in the evaluation stage is meant to determine in which direction it is more rewarding for the agent to move. The agent simply confronts the efficiency of the currently employed technological configuration with a set of $b \in [1;N]$ alternative configurations, all from its direct neighborhood, and chooses to move to the one that has the highest efficiency (or, if all the probed neighbors are less efficient than the current configuration, the agent just stays put).

Going back to the example of chess, breadth of search then is a measure of how many different possibilities are analyzed before the player chooses which figure to move.

2.2. Depth of Search on Technology Landscape

While breadth of search enriches the structure of the model horizontally, it still remains completely “flat”. The efficiency of the alternatives is estimated on the basis of an extremely short-run assessment, so that while gaining in breadth, the strategy remains myopic.

Existence of bounds on rationality and limited organizational abilities to assess far-stretching consequences of today’s choice is a known fact; however, going from an extreme of endowing the agents with perfect foresight, as commonly done in neoclassical microeconomic models, directly to another extreme of providing them with no insight at all, like it is done as commonly in evolutionary economic modeling, we actually “fly over” the most interesting cases.

It serves then for the purposes of the current analysis to add a vertical measure of search, its *depth*. Depth of search here determines how *insightful* the agents conducting the search are. Just like in the case of breadth, we have a set $D \in [1; N]$.

The model thus recognizes the fact that the value of each alternative is comprised of (1) its correspondent efficiency (direct present reward of being in a given position), and (2) the options for further improvements that alternative creates (possible future rewards of it). The original *NK Model* was limited to evaluating alternatives solely on the basis of the former criteria. While it is possible to combine the two, the current paper, alternatively, is focused on strategies that are based on the latter criteria of assessment.

It is important to highlight at this point that even after introduction of a vertical measure of search, it still remains local in nature. An agent characterized by a higher value of depth is capable to base its current decision on shifting to a specific technological configuration on the possibilities for future evolution towards higher fit such shift provides her with. This is not to say, however, that having spotted an attractive high-efficiency configuration, the agent is able to make a shift to it directly, as would have been the case for the scenario of long-jump search with foresight.⁷ The agent would still have to change the state of one element at a time, possibly having to suffer from being positioned in low-efficiency points before the goal is reached. This is one of the ways in which more insightful search is costly.

2.3. Drawing Cognitive Maps

The idea is rather similar to the scenario proposed by Gavetti & Levinthal (2000). Their simulation analysis reveals the interplay between drawing a cognitive map of reality, and the

⁷ This is the scenario proposed in Gavetti & Levinthal 2000. See below for more detailed review of their model.

experience bases possessed by the adapting agents. Thus, the agents are seen as both forward- and backward-looking.

In their treatment of the issue, the agents are first building a partial representation of the landscape. Such representation provides them with an opportunity to spot the globally optimal configuration of some particular area of the actual landscape.

So far the strategy is very similar to the one proposed in this paper. However, in their model, the agents move directly to that point in the landscape, and from there start the usual local trial-and-error search. So after the initial *insightful* long jump such a strategy provides for, nothing really changes. If we consider the actual evolution towards better fit after *step 1*, the treatment is identical to the original *NK Model* with the initial point is exogenously set by the modeler.

In the extensions of the model the authors deal with the issue of representations changing with time, but the logic remains: representation building and experiential search are treated rather separately than simultaneously.

Spotting the global optimum on a lower-dimensional sub-landscape in the current model, instead of giving the agents the opportunity to start the local search from a “better neighborhood”, provides them with an insight on what direction of search might prove to be more fruitful. As noted above, no long jumps are allowed for in the present treatment, and thus, having identified the current *goal*, the agents can make just one single local step towards it in each period.

The other important way in which the two models differ, is that *goals* are reconsidered again and again at each step of the process. The agents here are constrained to observe the efficiencies of some given neighborhood around the currently used configuration only. However, at each step of the evolution with a positive probability the agent accepts a different technological configuration, and while the maximum dimension of the “observable” neighborhood stays the same, such shift in a position occupied by an agent would result in a different set of points on the landscape that fall into it. So then, for any dimensionality of the observable neighborhood, there is a positive probability that at any step the agent would evaluate an yet unencountered configuration with an efficiency level higher than that of the current goal. Whenever that happens, the goal, and in consequence the direction in which the search is conducted would change over and over again.

The dimensionality of the observable neighborhood is a function of the depth and the breadth of search, and thus, ultimately, depending on the values of those two parameters, the agents would either reach the global maximum or else find themselves *stuck* on a sub optimal peak.

3. Simulation Model. Technicalities

3.1. Simulation Toolkit

All the simulations below were run using Laboratory for Simulation Development (Lsd) language, developed by Marco Valente. Lsd is a freeware that can be downloaded from <http://www.business.auc.dk/~mv/Lsd/lsd.html>. This simulation language is built on C++ platform, and thus is characterized by the speed and flexibility of a low-level language. However, the layer of interfaces embedded in its structure make it much more user-friendly than the former, and possible to use by non-programmers.

NK Model in the original setting is included in the Lsd package as one of the example models. The code for the modified version of it, used in the subsequent analysis in the present paper is available on demand from the author.

3.2. Model. Formal Setting

As in the original model by Kauffman (1987) we define the landscape with the help of the two main parameters: N and K .

N measures the size of the system, or, more precisely, the number of elements it is comprised with. The “system” in our case is the smallest independent *block structure* obtained through the process of decomposition. K reflects the complexity of such system through measuring the level of interconnectedness between its elements. It is the number of other elements a change in the state of the given element affects (or else the average number of other elements that affect the given element changing their state). Formally the system is represented by a binary string, so a change in the state of an element means a flip from “0” to “1” or vice versa.

At the *zeroth* step of the simulation run all the agents are randomly placed on some point of the landscape from where the search is to be conducted. That starting point is defined by a random assignment of the binary strings.

Each point on the landscape is characterized by some *efficiency value*, $\theta(\omega)$ that is measured as an average over the *efficiency contributions* of the elements’ states in the system.

Just as in the original model of random local search, at each step an agent can change the state of one element only. However, the judgment on whether to change or not the state of a chosen element is made on different grounds in this model, formalized below.

The following table presents a summary of the key parameters used:

Table 1: Summary of parameters

Parameter	Description	Range
N	number of elements (operations) per configuration	<i>positive integers</i>
K	number of intranalities (interdependencies) per element	$\{0, \dots, N-1\}$
b	number of directions sampled per step per element	$\{1, \dots, N\}$
d	current depth of search	$\{1, \dots, d_{max}\}$
d_{max}	current maximum depth of search	$\{1, \dots, D\}$
D	maximum depth of search of the agent defined by nature	$\{1, \dots, N\}$
τ_{max}	frequency of updating of d	<i>positive integers</i>
τ	time steps elapsed since the last update of d	$\{1, \dots, \tau_{max}\}$
H_{i,ω_j}	set of recipes i levels away from ω_j	<i>NA</i>
$\Pi\omega_{cur,\omega_{max}}$	set of elements on the shortest path between ω_{cur} & ω_{max}	<i>NA</i>
θ_{cur}	efficiency of the currently used recipe	$\{0, \dots, 1\}$
θ_{max}	efficiency of the best recipe currently observable (ω_{max})	$\{0, \dots, 1\}$
ω_{cur}	currently used recipe	2^N
ω_{max}	most efficient recipe currently observable	$\bigcup_{d \in \{1, \dots, d_{max}\}} \Omega_d$
Ω_d	set of recipes sampled at the current step at depth d	<i>see below</i>

The search strategy algorithm employed by the agents of the model can be described technically in the following way:

0. Observe the initial values for N , K , D , b , ω_{cur} and τ_{max} , given by nature (supplied by the modeler), and go to step (1)
1. Set initially $d=1$; $d_{max}=1$; $\tau=0$; $\theta_{max}=0$, and go to step (2)
2. If $\tau=\tau_{max}$, increase d_{max} by 1, and go to step (3)
3. Observe the current value of d_{max}
 - a. if $d_{max}>D$ set $d_{max}=1$ and go to step (6)
 - b. else if $d_{max}\leq D$ go to step (4)
4. Choose randomly b recipes from a set $H_{1,\omega_{cur}}$. Denote the resulting set Ω_1
 - a. if $d_{max}=1$ go to step (5)
 - b. else if $d_{max}>1$, for each level $d \in \{2, \dots, d_{max}\}$ and each element in Ω_{d-1} choose randomly b recipes from the corresponding $H_{1,\omega_{i \in d-1}}$ sets of recipes, denoting the resulting set Ω_d . Go to step (5)

5. Observe and compare the efficiencies of the resulting set $\bigcup_{d \in \{1, \dots, d_{max}\}} \Omega_d$ of $b+b^2 \dots +b^{d_{max}}$ sampled recipes with the efficiency of ω_{cur} and identify ω_{max}
 - a. if $\theta_{cur} = \theta_{max}$ increase τ by 1, and go to step (2)
 - b. else if $\theta_{cur} \neq \theta_{max}$
 - i) if $\omega_{max} \in H_{1, \omega_{cur}}$ shift to ω_{max} , set $\omega_{cur} = \omega_{max}$, $\tau = 0$, $d_{max} = 1$, and go to step (2)
 - ii) else if $\omega_{max} \notin H_{1, \omega_{cur}}$ shift to $\omega_i \in [\Omega_1 \cap \Pi \omega_{cur}, \omega_{max}]$, set $\tau = 0$, $\omega_{cur} = \omega_i$, $d_{max} = 1$, and go to step (2)
6. Evolution stops here

4. Simulation Results. Limit Cases

4.1. Strategies of Greedy Myopic Search

Let us first consider a set of firms that are able to evaluate many options for change at every point in time, but lack foresight. So then, the structure of the problem is completely *flat* at every step. In terminology of the current model this would cover the set of strategies with $b \in [1, \dots, N]$ and $D=1$. In broader terminology these strategies are variations of so-called *strategies of greedy local search*. The algorithm above is valid for the case, although much of its loops and cycles become redundant.

What do we want to see at this point is whether increased *breadth of search* can be of any help even when the process lacks any *depth*. For that purpose we would run simulations for each of the combinations of $b \in [1, \dots, N]$ and $K \in [0, \dots, N-1]$. Setting $N=20$ for all the simulations this gives us $20^2=400$ combinations. To avoid having results biased due to some particular random event, the simulations are run for 10 different initial seeds of random numbers (thus for different realizations of the landscape structure), and within each seed we have 10 agents differing in starting point of the evolution. That gives us a total of 4.000 observations, 100 for each node in the graph. The figures below represent the results of the simulation runs.

<<Insert Figure 1 about here>>

Figure 1 shows that the implied ability to conduct off-line search in several distinct directions in parallel improves the overall performance. However, the improvements are not very substantial: for no combination of b and K it is over 5% from the benchmark case, and for the main mass of the possible combinations it is just about 1-2%. These improvements are definitely not even close to be

of the magnitudes needed to level off the deleterious effects on the efficiency that more complex landscapes bring about. This is clearly visible in *Figure 2*, where the average levels of terminal efficiency are plotted as a fraction of the global maximum. In fact, because of the changed scale, the high peaks clearly visible on *Figure 1* are no longer obvious.

<<Insert Figure 2 about here>>

As we can also observe from the *Figure 2*, for any technology with interrelated elements (any technological landscape with $K \neq 0$), breadth alone, no matter small or large does not guarantee the firms employing the strategy to achieve the globally optimal efficiency other than by chance.

Another interesting result is presented in *Figure 3*. Here we can see that while the improvements over the benchmark case are positively correlated with an increase in b as an average over all K , increasing K contributes to improvements in the efficiency up to some point only, making an overall effect of increasing both b and K ambiguous. Indeed, the highest improvements are achieved for large b and average values of K .

<<Insert Figure 3 about here>>

Increased breadth of search has a double effect. From the one hand it has a positive effect in what rather than just sampling a single adjacent technology at a time and making a decision on whether to accept or reject the shift, the firms in this setting are able to view several options for change, and choose the best one.

However, on the other hand, the more *greedy* the search process becomes, the more crucial difference does the assumption of it being completely *myopic* makes, so, the more is the probability that a firm would end up on a local peak after a very short time. Indeed, as the *Figure 4* shows search time is a strictly decreasing function of both b and K .

<<Insert Figure 4 about here>>

This means, however, that strategies with higher breadth are faster on average, and thus, for a market populated by firms differing only in this characteristic, this would be a clear-cut advantage if competition enters the general picture.

4.2. Strategies of Search with Narrow Insight

Let us now go to the other extreme, and consider the limit case when the insight into the future that the firms possess is a tunable parameter, while only one direction for change is

investigated at each time step. So, while the firms are not myopic anymore, their search instead is set to be extremely narrow.

The intuition behind such set of the strategies is that while being *rigid* in defining a strategy, the firms nevertheless are able to think the strategy through for more than one step ahead. So the firm designs a long-term plan at each given step, analyzing whether the direction of change chosen would be fruitful for the future growth, but, is flexible enough to reconsider the exact direction of change in the next period, if the new information that became available calls for such action. If that happens a new long-term plan is designed, and is accepted as a guideline for future change unless at any future period an alternative is found that has a higher maximum payoff in the future.

Technically, similarly to the previously analyzed limit case of greedy myopic search, it would mean that this would cover the set of strategies with $D \in [1, \dots, N]$ and $b=1$ on the set of the landscapes with varying complexity, so that $K \in [0, \dots, N-1]$.

All the rest of the settings are left like before, so once again that gives us a total of 4.000 observations, 100 for each node in the graphs below. While the setting remains almost unchanged, the results we obtain from running the simulations differ quite substantially.

<<Insert Figure 5 about here>>

Figure 5 (compare with *Figure 1*) shows us the magnitude of improvements over the benchmark case that such strategies bring about. The results are much more impressive than in the previous setting, with the improvements of about 15% not being stand outs, and an average (over all combinations of $D > 1$ and $K > 0$) of about 7% observed.

This is being further confirmed by the results summarized in *Figure 6* (compare with *Figure 2*) where the average levels of terminal efficiency are plotted as a fraction of the global maximum. Here, we see that while the improvements are still not enough to cover the gap with the global maximum, the firms are getting much closer to it even for very complex landscapes.

<<Insert Figure 6 about here>>

However, as it had been hinted above in the discussion on dangers of modularity, in absolute terms the average efficiency of the global maximum on the technology landscape is positively correlated with the complexity of the latter.

This is obvious from the *Figure 7* where the average results over 400 (20 for each value of K) landscapes are presented. This means that the results reported as a fraction of the corresponding global maximum are biased towards lower levels of complexity. If we adjust the results accordingly, we observe in *Figure 8* that when applying a strategy that is not myopic, the firms

would quite often prefer using more complex and interconnected technologies rather than modular ones. As the graph shows this is actually true even for the benchmark case. Despite the fact that the global maximum is obtained only for $K=0$ landscapes, in the absolute terms the highest efficiency is reached when K is as high as 5.

<<Insert Figures 7 and 8 about here>>

However, not all is as bright as it seems. While clearly having a significant positive impact on the terminal efficiency of the technological configuration, a strategy with *insight* take longer to actually get us there. Such strategies lead to an increase in checking and double checking whether the technological configuration attained is indeed the best of what's around, and, given the structure of the algorithm, the agents would leave a relatively high local peak even if they spotted another one only marginally better, no matter how far that new target is. All this leads to increased search times, and thus, higher costs and the danger of being outcompeted by less “insightful” but faster to adapt rivals.

5. Simulation Results. General Case

5.1. Striving for the Global Optimum

What we have seen until now is that both breadth and depth of search taken separately improve the terminal efficiency of the technology used by the firms, albeit to a different extent and with differing side effects. It has also been shown that while in relative terms, measured as a fraction of the global optimum, the improvements in the efficiency level due to an increase in either depth or breadth of search taken separately do not manage to outshadow the negative effects of growing complexity, in absolute terms, the levels of efficiency achieved, with an increase in complexity first grow and then only start to fall. These results already cast some doubt on whether an extreme modularization of the landscape as proposed in the related literature is indeed justified.

However, to make the picture complete we would try here to answer whether there is any combination of b and D other than that corresponding to the strategy with perfect foresight that leads even the firms adapting on quite complex technology landscapes to the global maximum.

For that purpose we would leave the limit cases, and consider the whole family of strategies with both $b \in [1, \dots, N]$ and $D \in [1, \dots, N]$. For brevity, we would not however consider all the possible combinations of b , D and K , running the simulations for $K=1$, $K=4$ and $K=11$.

The results summarized in the *Figures 9, 10 and 11* speak for themselves: not only perfect foresight is not a necessary condition for reaching the global maximum, but in fact there is a very large set of less “perfect” alternative strategies that do just as good.

<<Insert Figures 9, 10 and 11 about here>>

It comes as no surprise that for more complex landscapes reaching the global peak requires more sophisticated strategy. More interesting is the observation we can make that perfect knowledge of the underlying structure of the connections between different elements or operations comprising the technology set is not a necessary condition even where such connections create a tremendously perplex web.

5.2. Actual vs. Maximum Depth

Do we require too much rationality from the agents? Is it true that having even a modest D would be too much to ask for? The answer is probably negative.

Non-stringency of the strategy sophistication is reinforced first of all by an observation we can make with the help of a more thorough analysis of the results obtained. Recall from the setup of the *search algorithm* above that D is the maximum depth of search of the agent defined by nature rather than the depth of search used at each step of the evolutionary process. Indeed if we observe the actual *depth* used throughout the evolution towards the ultimate *sticking point*, its average value even for the firms endowed by the modeler with high levels of D is much lower on average. For the large part of the life-cycle the firms still use the low-cost myopic reasoning that suffices for dismissing clearly inferior directions of search, while going deep into analyzing only several promising ones. Indeed the value of d (see *Table 1* for definition) is being increased only when its current value is not enough to justify the dismissal of a particular direction; it is only in this case that a firm endowed with $D > 1$ uses the insight, and even then, the increase in d is gradual, rather than abrupt, and as soon as its level is sufficient to spot an alternative direction leading to a more efficient technological configuration starts the new cycle of evaluation process from $d=1$.

To show this we ran simulations for the landscape with $N=20$, $K=10$, $D \in [1, \dots, N]$, and $b \in [1, \dots, 6]$ and take the average over the actual depth d for each class of agents. As the results in *Figure 12* show, only for the case of $b=1$ and high values of D do we observe average actual values of d higher than 2, and even in this case the highest value observed is just over 3 for the case when D was set to 20.

<<Insert Figure 12 about here>>

Another important thing that this figure tells us is that d grows with the growth of D only for the classes of agents whose breadth and depth of search are not sufficient to guarantee reaching the global optimum in all instances. When for a given b , the value of D exceeds that threshold, the further increase of the latter does influence the average value of d . This leads us to suggest that for a given b all the strategies with the value of D higher than the lowest needed to guarantee 100% of the agents to find themselves at the global peak, are almost identical. All of them spot the global maximum at exactly the same time, and reach it after exactly the same number of steps, the only difference being that more “insightful” agents take more time to double check the fact that this indeed is the global maximum, putting it through a test of confronting to a larger number of points on the landscape. Before reaching the global maximum such strategies are identical in all respects.

Indeed, the only reason the lines on the graph are not completely flat after some point lies in stochasticity of the landscape structure.

5.3. Further Arguments on Non-Stringency of Rationality Assumptions

Another important characteristic implicit in the current design of the search algorithm is that the search is *redundant*, in the sense that the firm can encounter exactly the same configuration of the technology over and over again. This is due to two main factors:

- (1) given the definition of the technological distance of any two recipes as being symmetric, if we made a local move from *technology 1* to *technology 2* at a given step, the former becomes a part of the *adjacent technologies set* for the latter, and given the random nature of choosing the b alternative elements with which the firm can *experiment* at each step, with a positive probability, increasing with an increase in the value of b , it well might also become a part of the set Ω_1 as defined above.
- (2) given the definition of the technology landscape as a *graph*, each two technological configurations can be converted into each other in a huge number of ways different in length. The shortest way to reach *technology 2* from the *technology 1* is termed the Hamming distance between the technologies. However, a variety of other, longer ways to do exactly the same exist, which leads to further redundancy of the search process.

Thus, the actual number of new technological configurations sampled at each step is much lower than the equation in the search algorithm would suggest, implying in sequence that just as in the case with high values of D , high values of b do not necessarily depict overly-sophisticated nature of the strategy in use.

Another way in which the strategies are clearly *boundedly rational* is that nothing prevents a firm to *lose the track* of the exact path that leads it to the targeted high-efficiency technological configuration. While its efficiency would become in a way the *aspiration level* that the other sampled technological opportunities would be compared with, the actual direction the target lies in can become blurred by the random choice of the directions for change in the next period.

Finally, we can also observe from the graphs that even when the strategy does not suffice for attainment of the globally optimal technological configuration for *all* the studied realizations of the landscape structure and all the sampled starting points of the evolutionary process, it usually does so in *most* of the cases, leading to average results just 1-2% lower than the global maximum. This would let us suggest that in a world where the ε -*satisficing solutions*⁸ (even if only the ones with ε very small) are acceptable alternatives to the globally optimal result, the requirements necessary to impose on the level of sophistication of the strategy in use are far from being stringent.

6. Enter the Competition

In such an abundance of strategies leading to the globally efficient solution or to a one just marginally inferior to it, we have to specify which combination of b and D we would actually want to choose. We have concluded so far that up to some limit the more insight is used in the search process, the more likely it is to end up at the global maximum. However, as we have also seen, after the minimal level of D sufficient for the attainment of the global maximum for all the agents in the class is reached, the further increase in D no longer has a positive effect.

Costs *per se* do not enter the analysis in this paper explicitly. Nevertheless, the costliness of search here can be implicitly measured by the average efficiency of the technologies used throughout the evolution towards the peak, as well as the length of time it takes the firms to get there.

Simulations ran and discussed above lead to suggest that search time is an increasing function of D , and decreasing function of b . So, if we necessarily want to secure the global efficiency for all the agents, and have no budget constraints, an obvious choice of the strategy would be to maximize b , and use the smallest $D \in D(b)^{max}$, where $D(b)^{max}$ is the set of depth values that lead to the global maximum for a given b .

⁸ Defined by Frenken *et al.* (1998) in the following way: “The set of ε -*satisficing solutions* is the set of strings whose value is at most ε lower than the global optimum”, page 157

However, from the one hand, the breadth of search is costly, and thus, we should be more interested in modest values of b , and from the other, as was noted above, there exist a whole set of strategies that although do not belong to $D(b)^{max}$ set, allow a large fraction of the agents in the respective classes to reach the global maximum, or else to come very close to it, while, at the same time economizing on search time. Moreover, this set of ε -satisficing strategies that we can term $D(b)^{sat}$ have another important advantage. Due to a lower D the agents suffer less from the low-efficiency intermediate positions. This effect was not pronounced in our simulations, since the terminal efficiency only was measured, and it was not of importance how well do the agents do a the intermediate steps of the simulation run.

In the real world, however, not only the final result matters, but also, the intermediate ones. To be able to study this factor, I introduce competition between the agents. The simulation proceeds as follows. At its start we have two classes, each containing 50 agents and they run on a landscape with $N=20$, $K=10$. For each pair of competitors simulations were repeated for 20 times. With frequency F , a certain number S of worst-performing agents are being replaced by the copies of the agents that perform the best. With probability P_{rand} the agents copy just the strategy of their *parents*, while with probability $1-P_{rand}$ they copy as well the position of their *parents* on the landscape.

We set the breadth of search, $b=3$. The classes differ in the value of D . One of the classes is characterized by a lowest value of $D \in D(3)^{max}$ which is $D=11$ and thus is the quickest strategy of the $D(3)^{max}$ set. The other class, alternatively is populated by agents using ε -satisficing strategies. First we ran simulations for agents with $D=8$ against the ones with $D=11$. The controls were:

- Speed of competitive pressure $F=[10, 25, 50, 100, 250]$ steps
- Strength of competitive pressure $S=[1, 2, 5, 10]$ agents
- Probability of random relocation $P_{rand}=[0, 1/2, 1]$

Changing those three controls didn't alter the results in any significant manner, for which reason the rest of simulation runs were performed for $F=50$, $S=2$ and $P_{rand}=1/2$.

We ran simulations to compare all the one-to-one combinations of $D \in [2; 3; 5; 8; 11]$. The simulation run would stop either if one of the classes achieved 100% of the *market share*, or if all the agents in both classes reached the global maximum, or if neither happened for more than 10000 steps. The table below summarizes the results:

Table 2: Market Share (Fraction of Agents Belonging to the Class) at the End of the Simulation Run

	$D=2$	$D=3$	$D=5$	$D=8$	$D=11$
$D=2$	--XXX--	24,4 \ 75,6	0 \ 100	0 \ 100	3,5 \ 96,5
$D=3$	75,6 \ 24,4	--XXX--	0 \ 100	0 \ 100	15,4 \ 84,6
$D=5$	100 \ 0	100 \ 0	--XXX--	50,4 \ 49,6	64,3 \ 35,7
$D=8$	100 \ 0	100 \ 0	49,6 \ 50,4	--XXX--	61,3 \ 38,7
$D=11$	96,5 \ 3,5	84,6 \ 15,4	35,7 \ 64,3	38,7 \ 61,3	--XXX--

As the results show, we observe that both $D=5$ and $D=8$ agents outcompete more insightful $D=11$ ones, leaving them just a little more than a third of the market share. This means that the *value added* of the increased insight after some point is out shadowed by the slower speed of adaptation and low intermediate efficiency levels.

$D=5$ agents perform marginally better than $D=8$ ones against more insightful agents, while in direct competition they each get about 50% of the share. $D=2$ and $D=3$ agents lose against *higher-depth* agents, but even they find a small market share when competing against $D=11$ agents, while losing it all for the ε -satisficing $D=5$ and $D=8$ ones. The reason for a positive market share values for $D=2$ and $D=3$ is to a large degree explained again by the slower adaptation speed of $D=11$ agents. Running simulation for longer than 10000 periods let the latter get the 100% market share in most of the cases, especially so for the case of competition with $D=2$ agents.

Conclusions and Further Research Agenda

Trying to build a theory or a model in one scientific field using the framework adapted from another is a challenging and a very dangerous venture. Both evolutionary economics, and even if we consider it as such to a lesser extent, the mainstream neoclassical economics are firmly

grounded in evolutionary biology and classical physics respectively. However, what works for biological and physical systems, might well not be appropriate for the social domain, and even if it is, major changes have to be made in order to fine tune the borrowed insights into the new field of application.

Although being itself designed as a contribution to the evolutionary thinking in economics, this paper is critical to both the competing fields. Concentrating on the behavioral part of the discourse, the starting point of the analysis has been made with the claim that evolutionary and neoclassical economic theories find themselves in the opposite extremes as far as the ability to foresee and the intentionality of the actions by economic agents are treated. Both extremes can sometimes be valid simplifications, but, too often indeed, they lead to very local results.

Behavioral assumptions of the neoclassical theory have been challenged from within the evolutionary stream on many occasions and although the analysis in this paper is meant to shed some more light on that issue, the main motivation behind is to analyze whether the assumption of very limited myopic foresight of the agents in the evolutionary economics itself does indeed have to be revised.

The paper addresses primarily the recent and fast growing stream of evolutionary modeling exercises based on Kauffman's *NK Model*, initially designed to study genetic evolution in microbiology. The previous applications of the model to economics proper have addressed a variety of topics in organizational and technological change, and have provided a number of extremely interesting results and insights.

One particular issue that attracted a substantial part of the research efforts in the field is the idea of connecting the model with Herbert Simon's insight on decomposition of complex systems. Indeed, with the rise in complexity of the problem stemming from the increase in the level of interconnectedness of the elements it is composed of, it becomes more and more difficult to find an efficient solution to it through local random search of the alternatives.

A way to tackle the issue proposed was to try to break the big problem into a number of independent or almost independent *patches* that can be solved separately without affecting each other in any significant way. This way of decomposing complex landscapes into more modular ones has been shown to be an effective strategy for success.

Modularity however has its own substantial drawbacks. While making it easier to find the global optimum on that *simplified* landscape, discarding from the negative externalities present on more complex landscape, modularization discards just as well of the possible positive externalities. As Fleming and Sorenson (2001) note: "Although the average peak height declines as

interdependence rises, some of the ‘good’ positions on the high- K landscape dominate the best points on the low- K landscape.”⁹

The problem with decomposition strategies is that they are aimed to find an objective way to make the landscape less complex, or, in other words, instead of trying to learn how to solve a more complex problem, the agents in that setting simply substitute that problem with a less complex one. However, taking technology as an example of a complex system, we might notice that apart from the problem addressed in the previous paragraph, decomposing the system (1) can be simply impossible to do, because of the difficulties of truncating the technology in question, and (2) makes the technology more imitable, thus increasing the dangers posed by the competitors (Rivkin 2000).

It was argued above that noticing that complexity is partly a subjective matter, an alternative way to cope with the issue can be proposed, and namely, that of trying to endogenize some part of the connections by learning their effects.

It has been claimed here that the treatment of the agents in the group of models as being able to evaluate only a single alternative at a time step, and then only to see just the immediate direct effects of a possible shift to it, while valid in the original biological domain of the model, is extremely unrealistic when we shift our attention to the issue of how firms develop their technologies.

So then, by introducing the notions of breadth and depth of search on a technology landscape, the current model has dealt with the ways to simplify the complex landscapes subjectively through learning of their underlying structure by sampling in several directions in parallel and giving a weight to longer-term effects a shift to an alternative technology can have.

It has been shown through simulation analysis that while breadth and depth of search taken separately contribute to an increase in the efficiency of the terminal “sticking point” technology, it is only with both factors present when the agents acquire the ability to effectively find the global optimum even on a very rugged landscape.

The costs of increasing the “observable” region of the overall landscape have not at this stage been modeled explicitly. However, it has been shown that through increasing the probability of getting stuck on a sub-optimal peak too early (breadth), and increasing the length of time required to find a sticking point, thus exposing themselves more to the dangers of competitive pressure (depth), such costs entered the picture implicitly.

⁹ Fleming, Lee & Olav Sorenson (2001) *Technology as a Complex Adaptive System: Evidence from Patent Data*. Research Policy 30, page 1022

Introducing the competitive pressure was the most direct way to extend the model, and the results shown that less “insightful” firms quite often outcompete the more “insightful” ones, but the myopic firms of the benchmark model still die out first.

There are several other ways in our research agenda in which the model can be modified. From the one hand, lifting the assumptions of extreme redundancy of search and inability to keep the track of the direction in which the current goal lies, would reinforce the results of the current model, letting firms possessing even lower levels of b and D to effectively reach the globally optimal configuration. From the other hand however, lifting the quite strong assumption of the agents being able to estimate the efficiency of the technologies that are within their “eyesight” precisely, no matter how distant they are, should work in the opposite direction¹⁰.

Moreover, several different ways of evaluating the fruitfulness of some particular direction of change might be interesting to explore, including basing the decision on the average (or weighted average) efficiency over the resulted *path* of change, or a minimax criteria to deal with risk-averse decision makers.

¹⁰ even if such noisy evaluation in the myopic case has proven to be efficiency-enhancing to some extent (see Hovhannisian 2003abc)

References:

1. **Altenberg, Lee** (1994) *Evolving Better Representations through Selective Genome Growth*, in Proceedings of the IEEE World Congress on Computational Intelligence, pages 182-187
2. **Altenberg, Lee** (1997), *Genome Growth and the Evolution of the Genotype-Phenotype Map*, the corrected and revised version of the paper available at: http://dynamics.org/_altenber/. Originally appeared in: **W. Banzhaf & Frank H. Eeckman** (1993), *Evolution and Biocomputation: Computational Models of Evolution*, Springer Verlag Lecture Notes in Computer Sciences vol. 899, 1995, pages 205-259
3. **Auerswald, Phillip; Stuart Kauffman, José Lobo & Karl Shell** (2000) *The Production Recipes Approach to Modeling Technological Innovation: An Application to Learning by Doing*, Journal of Economic Dynamics and Control, 24, pages 389-450
4. **Azencott, Robert** (1992) *Simulated Annealing: Parallelization Techniques*. John Wiley & Son, USA
5. **Baldwin, Carliss & Clark Kim B.** (2000) *Design Rules: The Power of Modularity*. MIT Press
6. **Barney, Jay B.** (1991) *Firm Resources and Sustained Competitive Advantage*, Journal of Management, 17, pages 99-120
7. **Carley, Kathleen M. & D. Svoboda** (1996), *Modeling Organizational Adaptation as a Simulated Annealing Process*, Sociological Methods and Research 25(1), pages 138-168
8. **Cyert, Richard M. & James G. March** (1963) *A Behavioral Theory of the Firm*. Blackwell Business
9. **Dosi, Giovanni; Daniel Levinthal & Luigi Marengo** (2001) *Bridging Contested Terrain: Linking Incentive-Based and Learning Perspectives on Organizational Evolution*. Paper presented at the DRUID Nelson-Winter Conference, Aalborg, June
10. **Durrett, Richard & Vlada Limic** (2003) *Rigorous Results for the NK Model*, The Annals of Probability, vol. 31, No. 4, pages 1713-1753
11. **Ethiraj, Sendil K. & Daniel Levinthal** (2004) *Modularity and Innovation in Complex Systems*. Management Science, vol. 50, No. 2 (February)
12. **Evans, Steven N. & David Steinsaltz** (2002) *Estimating Some Features of NK Fitness Landscapes*, The Annals of Probability, vol. 14, No. 4, pages 1299-1321
13. **Fleming, Lee & Olav Sorenson** (2001a) *Technology as a Complex Adaptive System: Evidence from Patent Data*. Research Policy 30, pages 1019-1039
14. **Fleming, Lee & Olav Sorenson** (2001b) *The Dangers of Modularity*. Harvard Business Review, September, pages 20-21
15. **Fleming, Lee & Olav Sorenson** (2003) *Navigating the Technology Landscape of Innovation*. MIT Sloan Management Review, Vol. 44 Issue 2, pages 15-24
16. **Fontana, Walter** (2003) *The Topology of the Possible*, to appear in *Paradigms of Change*, **A. Wimmer and R. Koessler** (eds.)
17. **Frenken, Koen, Luigi Marengo & Marco Valente** (1999) *Interdependencies, Near-Decomposability and Adaptation*. In: **Brenner, T.** (ed.) *Computational Techniques for Modeling Learning in Economics*. Kluwer Academic Publishers, pages 145-165
18. **Gavetti, Giovanni & Daniel Levinthal** (2000) *Looking Forward and Looking Backward: Cognitive and Experiential Search*, Administrative Science Quarterly, 45, pages 113-137
19. **Hannan, Michael T. & John H. Freeman** (1984) *Structural Inertia and Organizational Change*, American Sociological Review, 49, pages 149-164

20. **Henderson, Rebecca M., Kim B. Clark** (1990) *Architectural Innovation: The Reconfiguration of Existing Product Technology and the Failure of Established Firms*, Administrative Science Quarterly, 35, pages 9-31
21. **Hovhannisian, Karén** (2002a) *Exploring on the Technology Landscapes: Real Options Thinking in the Context of the Complexity Theory*. Paper presented at the DRUID Conference, Aalborg, January 17-19
22. **Hovhannisian, Karén** (2002b) *Exploring on Evolving Landscapes: A Quest for Strategy Heuristic in Complex Worlds*, Paper presented at Richard M. Goodwin Seminar Series, University of Siena, March 7
23. **Hovhannisian, Karén** (2003a) *Satisficing, Deliberate Experimentation and Designing without Final Goals: Modeling Innovation-Related Business Strategy Choice through Simulation Analysis*. Paper presented at the DRUID Conference, Aalborg, January 16-18
24. **Hovhannisian, Karén** (2003b) *"Imperfect" Local Search Strategies on Technology Landscapes: Satisficing, Deliberate Experimentation and Memory Dependence*. Paper to be presented at the WEHIA Workshop, Kiel, May 29-31
25. **Kauffman, Stuart; José Lobo & William G. Macready** (2000) *Optimal Search on a Technology Landscape*, Journal of Economic Behavior & Organization, vol. 43, pages 141-166
26. **Kauffman, Stuart** (1993) *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press. New York
27. **Kauffman, Stuart & Simon Levin** (1987) *Towards a General Theory of Adaptive Walks on Rugged Landscapes*. Journal of Theoretical Biology 128, pages 11-45
28. **Kauffman, Stuart; William G. Macready & Emily Dickinson** (1994) *Divide to Coordinate: Coevolutionary Problem Solving*, Unpublished Manuscript, Santa Fe Institute
29. **Kauffman, Stuart & William G. Macready** (1995) *Technological Evolution and Adaptive Organizations*, Santa Fe Institute Working Paper 95-02-008
30. **Kimura, Motoo** (1983) *The Neutral Theory of Molecular Evolution*. Cambridge University Press, Cambridge, UK
31. **Kirkpatrick, S., C. D. Gelatt, Jr. & M.P. Vecchi** (1983) *Optimization by Simulated Annealing*, Science, Number 4598, pages 671-680
32. **Kirman, Alan** (1987) *Graph Theory*, in **J. Eatwell, M. Milgate & P. Newman** (eds.) *The New Palgrave: A Dictionary of Economics*. Vol. 2. London: Macmillan, pages 558-559
33. **Levinthal, Daniel** (1997) *Adaptation on Rugged Landscapes*, Management Science, 43, pages 934-950
34. **Levinthal, Daniel & James G. March** (1981) *A Model of Adaptive Organizational Search*, Journal of Economic Behavior and Organization, 2, pages 307-333
35. **Levitan, Bennett & Stuart Kauffman** (1995) *Adaptive Walks with Noisy Fitness Measurements*. Molecular Diversity, 1(1), pages 53-68
36. **Loasby, Brian** (1976) *Choice, Complexity and Ignorance*. Cambridge University Press
37. **Lobo, José & William G. Macready** (1999) *Landscapes: A Natural Extension of Search Theory*, Santa Fe Institute Working Paper 99-05-037 E
38. **Macready, William G, Atahanassios G. Siapas & Stuart A. Kauffman**, (1996) *Criticality and Parallelism in Combinatorial Optimization*. Science, 271, 56
39. **March, James G.** (1991) *Exploration and Exploitation in Organizational Learning*, Organizational Science, vol. 2, No. 1, pages 71-87
40. **March, James G. Herbert Simon** (1958) *Organizations*, New York: John Wiley
41. **Nelson, Richard R. & Sidney G. Winter** (1982) *An Evolutionary Theory of Economic Change*, Harvard University Press

42. **Potts, Jason** (2000) *The New Evolutionary Microeconomics: Complexity, Competence and Adaptive Behaviour*. Cheltenham: Edward Elgar
43. **Rivkin, Jan W.** (2000) *Imitation of Complex Strategies*. *Management Science*, vol. 46, No. 6 (June), page 824-844
44. **Rivkin, Jan W. & Nicolaj Siggelkow** (2003) *Balancing Search and Stability: Interdependencies among Elements of Organizational Design*, *Management Science* 49(3), pages 290-311
45. **Roberts, Kevin, Martin L. Weitzman** (1981) *Funding Criteria for Research, Development, and Exploration Projects*, *Econometrica*, Vol. 49 Issue 5, pages 1261-1294
46. **Siggelkow, Nicolaj & Daniel Levinthal** (2003) *Temporarily Divide to Conquer: Centralized, Decentralized and Reintegrated Organizational Approaches to Exploration and Adaptation*, unpublished manuscript
47. **Simon, Herbert** (1947) *Administrative Behavior: A Study of Decision-Making Processes in Administrative Organizations*. The Free Press
48. **Simon, Herbert** (1969) *The Sciences of the Artificial*. The MIT Press
49. **Teece, David J.** (1986), *Profiting from Technological Innovation*, *Research Policy*, 15:6, pages 285-305
50. **Valente Marco** (2000) *Evolutionary Economics and Computer Simulations: A Model for the Evolution of Markets*, University of Aalborg, PhD Thesis
51. **Vishwanath Tara** (1988) *Parallel Search and Information Gathering*, *American Economic Review*, Vol. 78 Issue 2, pages 110-117
52. **Vishwanath Tara** (1992), *Parallel Search for the Best Alternative*, *Economic Theory*, Vol. 2 Issue 4, pages 495-508
53. **Weitzman, Martin L.** (1979) *Optimal Search for the Best Alternative*, *Econometrica*, Vol. 47 Issue 3, pages 641-656

Figure 1: Greedy Myopic Search. Improvements over Benchmark Case




Figure 2: Greedy Myopic Search. Terminal Efficiency as a Fraction of the Global Maximum




Figure 3: Average Improvements over Benchmark Case as a Function of K and Breadth




Figure 4: Search Time as a Function of b and K




Figure 5: Search with Narrow Insight. Improvements over Benchmark Case




Figure 6: Search with Narrow Insight. Terminal Efficiency as a Fraction of the Global Maximum




Figure 7: Global Maximum as a Function of K




Figure 8: Search with Narrow Insight. Terminal Efficiency as a Function of K





Figure 9: Terminal Efficiency as a Fraction of the Global Maximum. $N=20, K=1$



2D Top Projection of Figure 9





Figure 10: Terminal Efficiency as a Fraction of the Global Maximum. $N=20, K=4$



2D Top Projection of Figure 10





Figure 11: Terminal Efficiency Level as a Fraction of the Global Maximum . $N=20, K=10$



2D Top Projection of Figure 11




Figure 12

