

***“Imperfect” Local Search Strategies on Technology Landscapes: Satisficing,
Deliberate Experimentation and Memory Dependence***

Karén Hovhannisian

DISA, Università degli Studi di Trento

Email: hovhannisian@unisi.it

Abstract:

This paper contributes to the recent stream of literature on NK Model’s applications to the field of technological evolution. It is argued that while the model has a great explanatory potential in economics proper, its behavioral foundations are still maladapted for treatment of purportive decision-making strategies for technological innovation. Concentrating on the decision rule for accepting novelties, we first analyze the consequences of intentional and unintentional imprecision in following hill-climbing strategy, highlighting the interplay between rigidity and deliberate experimentation. Building on Simon’s insights on satisficing behavior and designing without final goals we build a simulative model that provides a possibility to compare strategies differing in the desired level of imprecision. Secondly, we shift our attention to the question of organizational memory, analyzing in a simulation setting a fully memory dependent and a fully memory independent innovation-related strategies. The results confirm that from the one hand up to a certain level “imperfection” of rule-following behavior is a virtue rather than a threat, while from the other, that past successes can preclude adaptability of the firm, while disregarding such successes can be very risky.

February, 2004

Introduction¹:

Modeling evolution of boundedly rational agents in economics is a tricky issue in many respects. Grounding such attempts in theories originally designed for analyzing evolutionary processes in a completely different domain makes the issue even more complex. However, coping with the difficulties encountered can well be justified by the rewarding end results of such an endeavor.

The recent and fast growing research on applying Stuart Kauffman's *NK Model* (Kauffman & Levin 1987, Kauffman 1993) to a vast terrain of behavioral, organizational and strategic issues in economics (see Levinthal 1997, Frenken *et al.* 1999, Kauffman *et al.* 2000, Gavetti & Levinthal 2000, Rivkin 2000, Fleming & Sorensen 2001 among many others) seem to provide a valid way to formalize in analytical, or more commonly, simulative models a number of phenomena that formerly have either been disregarded, or remained the prerogative of appreciative theorizing.

Few of the most significant advantages the model provides us with when applied to studies of evolving technological and organizational forms in microeconomics can be seen as meta-theoretical and not confined to any field of science in particular. First and foremost this refers to its graph-theoretical structure and bases in non-integral space, presupposing the connections between the elements in the system under scrutiny to have as much explanatory power in the analysis of its dynamics as the characteristics of those elements themselves. When the system that we analyze is the technology employed by a firm, this would translate into a shift from the commonly used (and as commonly criticized) *production function* approach to that of *production recipes*, acquired ability to include in the analysis apart from the *existing*, the *nascent* technologies, and the evolution taking place on *performance* or *technology landscape* where the *distance* between two technologies determines the ease with which a shift from one to another can be made.²

¹ The author would like to thank Marco Valente, Koen Frenken, Esben S. Andersen, Ugo Pagano, Patrick Llerena, Brian Loasby, Toke Reichstein and the participants of the DRUID 2003 Winter Conference, University of Siena Seminar Series, ETIC-LAB Workshop and WEHIA 2003 Conference for their comments and suggestions on the previous drafts of the paper. Marco Valente's help was especially crucial for coding the model. All the mistakes and omissions however are mine only.

² All those notions and concepts will be defined and explained in detail below

Nevertheless, the behavioral foundations of the model are clearly field-specific, and while fitting well in explaining the evolutionary process of change on microbiological level, they need to be seriously adjusted when we venture into an analysis of agents who are humans or that are man-made artifacts. Unlike genes, humans can think, define and follow different rules, act strategically, improvise, experiment, learn from their mistakes and successes, remember and forget, be biased, use insight, etc, etc... The list would have still been incomplete even if I went on for the rest of my life.

Even more ambitious than “simply” listing all the determinants of behavior in which human or human-shaped evolution differs from evolution on genetic level, would be an attempt to actually incorporate all those in a single model. This is clearly not my intention here.

Instead, I would mainly concentrate on the rule of accepting or rejecting novelty, first introduced to the field of technological change by Kauffman & Macready (1995), and used quite consistently ever since. My aim would be in trying to show the inconsistency of *hill climbing* scenario from a behavioral point of view in its being both *myopic* and *perfect* at the same time. The assumption of absolute myopia has been lifted in our parallel research (Hovhannisian 2003b) to give a room for analyzing breadth and depth of search, while this paper in sequence deals with a scenario of *myopic local search with noise*, or as we call it here, *imperfect local search*.

Doing so opens up space to discuss some further behavioral issues like deliberate experimentation, satisficing, designing without final goals, organizational memory and organizational forgetting. The analysis is formalized in a simulation setting taking as a benchmark the original formulation of the model and showing the differences in dynamics the named modifications bring about.

The paper proceeds as follows. In the next section, the framework of research is defined, the definitions provided, and the original model discussed. After that we concentrate on the search rule used by the agents, and show both desirability and plausibility of introducing imperfections in the myopic local search. Special attention is given to the requirement of consistency in the treatment of bounds on rationality and satisficing behavior. Further on, satisficing behavior is discussed more in detail in

connection to the issue of designing without final goals (Simon 1969). At this point the simulative model, based on the insights thus gained, is built and run in two major settings, differing in the treatment of organizational memory. After the analysis of results, conclusions are drawn, and further research agenda outlined.

Defining the Framework of Analysis:

1.1. Basics of Kauffman's NK Model

As the name suggests, two main components of the model are N and K , where the former represents the number of elements a system is comprised of, and thus defines how large it is, while the latter measures the number of other elements a change in a given element's state affects, and thus defines the level of interdependences within it.

An additional component of the system is A , which measures the number of different *states* each element can occupy. For simplicity, A is normally set to 2, so that we can represent the system as a binary string. For example the string {1110} would represent a system of 4 elements, 1 of which is in the state '0', while the 3 others in the state '1'. If we think of some technology as a system, states of the elements can be seen either as an *on/off* setting of a knob on a machine, or else as a usage of one of the alternative technological processes or features. This is similar to the idea of *morphological analysis of technological trajectories* (Foray & Grübler 1991). In their example "four characteristic parameters [*of molding technology*] are identified and subdivided:

P₁: The nature of the pattern (P₁¹: permanent, P₁²: lost);

P₂: The nature of the mold cavity (P₂¹: hollow, P₂²: full);

P₃: The stabilization force (P₃¹: chemical, P₃²: physical);

P₄: The bonding method (P₄¹: simple, P₄²: complex)."³

Translating this into our framework, the subscript after each P would denote the position of the given element in the string, and the superscript would denote its state ("0" or "1"). In this particular example the elements can be combined in $2^4=16$ different ways

³ Foray, Dominique & Arnulf Grübler (1991) *Morphological Analysis, Diffusion, and Patterns of Technological Evolution: Ferrous Casting in France and the FRG*, in Nakićenović, N. & Grübler, A. (eds.) *Diffusion of Technologies and Social Behavior*. Springer Verlag, page 410

that we would call *technological configurations*.⁴ The efficiency of the system is measured as an average over the efficiencies of the elements it is comprised of.⁵

In a simple system where the elements are independent, the efficiency of each element depends only on its own state, so that the overall efficiency of the system can be optimized in at most N steps through making a pair wise comparison between the efficiencies of the states of the elements one at a time.

In more general and realistic case the elements comprising a system are interdependent, so that the efficiency of each element depends on its own state and on the states of K other elements. The case of $K=0$ would then represent a system of independent elements, while the case of $K=N-1$ would represent a fully interdependent system, with all the other cases falling in-between the two extremes.

1.2. Production Recipes and Technology Landscape

From a formal point of view a system defined in this way is a graph, Γ , composed of two types of sets – *vertices* (the elements of the system) and *edges* (the connections between them), so that we have $\Gamma = (V, E)$. If we want to use that structure in microeconomic theory, we need to make a shift from considering technology (or a firm) as a *production function* to viewing it as a *production recipe*.

The word recipe itself leads to a very intuitive example explaining the importance of that change. It is indeed not enough to know just the list of the ingredients (elements) such as flour, water, salt, yeast and the temperature of the oven, together with their relative quantities to bake a loaf of bread. What we also need to know is how those ingredients are combined, in what sequence they are to be used and so on.

This is just as true when we talk about producing a car rather than baking a loaf of bread, or, indeed, when we talk about designing the whole organizational and technological structure of a firm.

⁴ in general the number of configurations is A^N .

⁵ simple average is used in the further analysis.

As defined by Auerswald *et al.* (2000): “A production recipe is a complete list of engineering instructions for producing given outputs from given inputs.”⁶

The specific assignment of states to each operation a technology is comprised of is termed a *technological configuration*. The whole set of possible technological configurations then is a multi-dimensional *technology landscape*.⁷ The number of dimensions here depends on the number of elements each configuration is comprised of, while the *ruggedness* of it is a function of how high is the level of interdependence between those elements. The term landscape itself is used due to the way such space looks like on a 3D plot. For $K=0$ such landscape would look like that of and around Mount Fuji, with a single peak representing the global (and unique) optimum. With an increase in K , however, the correlation between the neighboring points on the landscape decreases, and the landscape starts to resemble more that of Alps, with a large number of local peaks of different height, and valleys of different depth between them.

Lobo and Macready (1999) provide the following definition: “A *technology landscape* consists of (1) a profit function assigning a real-valued number to each technology in the space of possible technological configurations; and (2) a metric structure over the space of technological possibilities which reflects whether any two given technologies are “close” to one another or “distant” from each other.”⁸

All the technological configurations that have been employed or at least sampled represent the sub-set of *existing technologies*, while the complement sub-set represents the (yet) undiscovered or *nascent technologies*.

1.3. Strategies of Innovative Change

Generally, firms can innovate either by changing the state of one or several elements (operations) that constitute a part of their current technology, or alternatively,

⁶ 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, page 394

⁷ Technology landscape is $N+1$ dimensional, with one dimension for each of the elements plus one for the efficiency mapping

⁸ Lobo, José & William G. Macready (1999) *Landscapes: A Natural Extension of Search Theory*, Santa Fe Institute Working Paper 99-05-037 E, page 1

by adding, replacing or removing any of those elements. While the latter possibility would be extremely interesting to explore, in line with the previous work on the subject, only the former option is dealt with in the current paper⁹.

Another important assumption concern the distance at which firms search for new technologies. We can define the distance between any two technological configurations on the landscape as the number of elements whose state has to be changed in order to convert from one to another. As Kauffman *et al.* (2000) write: “More precisely, the distance $d(\omega_i, \omega_j)$ between the production recipes ω_i and ω_j is the *minimum* number of operations which must be changed in order to convert ω_i to ω_j .”¹⁰

Implicit in that statement is that distance is a symmetric measure of differences, so, it is as easy to convert *Technology A* into *Technology B* as the other way round. This can be seen as an alarming limitation of the model setting, but, in order to keep the model simple, and again in line with the previous work, the assumption is kept intact¹¹.

More importantly for the purpose of the current paper, the notion of distance between the configurations enables us to distinguish between *local search* strategies and what have been termed the strategies of *long jumps*. The search is local if the state of only one element in the system is changed at a time. At any point then, a firm employing a particular technological configuration can move along the edges of the graph to any of the *adjacent* vertices, or else stay where it was. Making a jump would mean that two or more elements' states are changed at a time, so that a firm acquires the ability to move to a vertice on the graph, not directly connected to the one where it was before the jump.

There are several, both theoretical and empirical justifications of considering solely local search strategies. First of all, as discussed in Barney (1991), Hannan & Freeman (1984), Henderson & Clark (1990) and elsewhere, firms tend to innovate incrementally, building on their current competences. Relatedly, Levinthal & March (1981), March

⁹ There is in fact quite a lot of empirical evidence that this is quite reasonable an assumption. For a straightforward example think of an industry like biotechnology. See, however, Altenberg (1994), (1997) for a model designed to allow for such possibilities.

¹⁰ **Kauffman, Stuart; José Lobo & William G. Macready (2000) *Optimal Search on a Technology Landscape*, Journal of Economic Behavior & Organization, vol. 43, page 146**

¹¹ For an extremely interesting discussion on asymmetric distances see Fontana (2003) and the references therein.

(1991) and Teece (1986) among others show that changing the technological configuration drastically is associated with very high risk and uncertainty, and the attempts to do so have a very large probability to fail.

Moreover, the simulation results of the model by Auerswald *et al.* (2000) confirm the intuition that: “Taking bigger steps on a given landscape is somewhat like walking with smaller steps on a more rugged landscape. Hence, increasing δ [the number of operations altered per trial] should be analogous to increasing e [number of intranalities per operation].”¹²

Setting Up a Model

2.1. Search Rule: Internal (In)consistency

Up until now we have provided a broad picture of the models of search on technology landscapes, discussing in sequence their foundations in more general *graph theory*, their common structure and the set of assumptions underlying that structure. Not much has been said however about the behavioral foundations and heuristics of search.

While in any other respect the model below is identical to a number of other models of search on technology landscape, this is exactly the *search rule* employed where the modifications are made.

In economic interpretation of the *NK Model*'s original setting a firm performs the search on the landscape through a random walk. Kauffman *et al.* (2000) define a search rule in the following simple way: “Let θ_i be the efficiency of the production recipe currently used by the firm, and let θ_j be the efficiency of a newly sampled production recipe; if $\theta_i < \theta_j$, the firm adopts $w_j \in \Omega$ in the next time period; if $\theta_i > \theta_j$, the firm keeps using w_i .”¹³

Connecting this to a broader literature on search theoretic models (see Roberts and Weitzman 1981, Weitzman 1979 or Vishwanath 1992), search strategy put this way can

¹² 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, page 429

¹³ Kauffman, Stuart; José Lobo & William G. Macready (2000) *Optimal Search on a Technology Landscape*, Journal of Economic Behavior & Organization, vol. 43, page 149

be characterized as *local*, *random* and *sequential*, the agents employing it as possessing *myopic perfect foresight*, while the acceptance rule as *overall performance* based.

Some of the *non-local* search strategies have been outlined above. The most profoundly studied *non-local* search strategy is the one of *parallel search*, where more than one local move is performed synchronously (see Macready *et al.* 1996). It has been found however that parallel search strategies are superior to local search only when the current efficiency of the system is very low, becoming increasingly disruptive for efficiency levels above average, especially so when the *degree of parallelism* is high.

The issue of on what level should the decision on whether to accept or reject a possible change in configuration be made has been studied extensively (see e.g. Kauffman *et al.* 1994, Frenken *et al.* 1999, Dosi *et al.* 2001, Rivkin & Siggelkow 2003 and Siggelkow & Levinthal 2003) based on Simon's insight on decomposability and near-decomposability. Among its other merits, this setting allowed for the issue of organizational coordination to enter the picture, and show how different the outcome of search is under different levels of centralization.

Nevertheless, probably the most crucial shortcoming of the original setting when its behavioral plausibility is put under test is in the internal inconsistency of treating the bounds on the rationality of the agents under analysis. There are two different angles to look at the "mental skills" of the modeled agents:

1. From the one hand, the agents are too limited in what they can only observe their current state and the consequences of their action (one at a time) for just one period ahead, while everything that happened before, and everything that might happen after, remains in complete impenetrable darkness. So the agents are bound to make a decision on accepting/rejecting novelty having their aspiration level defined exclusively by the efficiency of the technology currently in use, and the efficiency of the sampled adjacent technology defined exclusively by an extremely short-run assessment. So then, a novelty is accepted if and only if the direct immediate gain is positive. This story depicts an extremely boundedly rational decision making process.

2. From the other hand, paradoxically, the agents are just too bright in what not only they are able to make a precise estimation of the efficiency of the currently used technology at each step of the process, but they are able, on top of it, to estimate with equal precision the efficiency of a technology a shift to which is under consideration at each such step. So they are able to know and not err a tad in knowing how good is something that they have never yet used. And they do not use the new technology before actually shifting to it in the original setting, because even if not stated explicitly, it is implied in the model that the *evaluation* and *action* stages of the process of change are disjoint and come in sequence rather than in parallel. Unlike the one in the previous paragraph, this story depicts an extremely over-rational decision making process.

This might be the right way to represent the evolution of genes in biology¹⁴, but the inherent internal inconsistency of behavioral foundations on which the model is built, from the perspective of man-designed technological evolution, or for that matter, of any evolutionary process in social domain, is extremely dangerous and implausible. We are simultaneously asking *too much* and *too little* from the agents in the model.

There are two alternative ways the author has pursued in an attempt to smooth out the acute edges of the problem mentioned.

An idea of *off-line parallel search with insight* has been recently explored (see Hovhannisian 2003b), showing that when parallelism is not binding, the *evaluation* and *action* stages of the process of change clearly distinguished and consistent with each other, and the assumption of absolute myopia lifted, the agents are able to reach the globally optimal configuration even being way short of possessing perfect foresight. In this way we attempt to solve the problem of inconsistency by loosening up the stringency of the story described in point (1) above.

In the present treatment, consequently, the assumption underpinning the story described in point (2) is relaxed, so that their *myopic foresight* is no longer *perfect*.

¹⁴ although see Levitan & Kauffman (1995) for a different point of view.

2.2. Intentional and Unintentional Imprecision

The “perfectness” part of the assumption of *myopic perfect foresight* has been challenged initially in the literature on *simulated annealing* (see Kirkpatrick *et al.* 1983) and *noisy adaptive walks* (see Levitan & Kauffman 1995). However, due to the fact that the former was related to the field of *evolutionary programming*, while the latter was a contribution to a debate in *evolutionary biology*, in neither of the two cases the motivation and intuition behind the change made was in any way related to the behavioral issues or the treatment of rationality.

Our claim here is that modification of the assumption is absolutely necessary in the social domain for the reason that considering unintentional and intentional imprecision in evaluation of novelty is crucial both for the internal consistency of the model setting, and for plausibility of the results obtained.

Unintentional imprecision is the easy part to explain. Indeed, we are hardly ever able to measure precisely the efficiency of the techniques and technologies in use, and far less so if the efficiency of novel and untried ones is attempted to be estimated. Once we put extreme bounds on agents’ foresight, as discussed above, both for the plausibility and for consistency of our argument we need to refine and narrow the bounds of their analytical and computational skills. It is not to say that the efficiency of a new technological configuration cannot be estimated at all, which would have been analogous to claim that each and every decision on making a technological change is taken purely at random. Instead we claim that while generally capable of making some approximate evaluation on whether change would be beneficial or not, some of the changes that would have been beneficial are being foregone, while some others that would later turn out to be detrimental are being made.

Indeed, the fact that by far not all the decisions made by firms are precise and frictionless, especially when innovation-related decisions are of our concern, comes as no surprise to anybody, is well documented in numerous case study analyses, and makes it into business news headlines ever so often.

Unintentional imprecision is not always a bad thing, however. As a matter of fact, products of such “mistakes” include cheese, Teflon, Coca Cola, potato chips, Guinness beer, aspirin, penicillin, glass, electricity and even America.

Much more interesting, however, is to make a case for intentionality of imprecision. The question we would try to answer below is: even assuming possible a perfectly precise estimation of the extremely short-run efficiency of the technological configuration a shift to which is contemplated on, would it be wise to follow the search rule as defined in the original model?

And my guess would be: no, it would not. First of all, to see why it will not be wise to do so, we need to return to the issue of the rule’s internal (in)consistency.

The overall value of any technological configuration, or indeed of any action, depends in general on two aspects: its current efficiency, and the possibilities for future actions a shift to it creates. This is close to the distinction between the of current and the option value. Say, standing in a long queue in front of a theater can hardly be called an extremely enjoyable way of spending time; however, we do stand in queues, because so doing provides us with an opportunity of watching a superb performance later on -- something that we might associate with a very high value. So, quite often, we are willingly decreasing the efficiency associated with our current state, in order to obtain a higher level of efficiency in the future.

In the myopic setting of the model, the agents do not possess any information about what the future might bring, except for in a very short-run. This does not mean, however that they do not realize the limits of their own long-term assessment skills. Recognition of such limits is an important behavioral factor. As Loasby put it: “[T]he recognition of ignorance changes the logic of choice.”¹⁵

The decision-makers realize that because the simple comparison between the current efficiencies of two alternative technological configurations does not contain all the necessary information to make a choice, rigidly following the hill-climbing policy,

¹⁵ **Loasby, Brian** (1976) *Choice, Complexity and Ignorance*. Cambridge University Press, page 74

they might end up being *precisely wrong* rather than *precisely right*. And as the popular among economists proverb goes: “it is better to be roughly right than precisely wrong”.

Rigidity is the key issue here. There is a strong evidence based on theoretical grounds, on the accounts of the real internal policies within firms, as well as on common sense, that flexibility gained through experimentation, if used within certain limits, is an extremely valuable asset.

Simon writes: “Exposure to new experiences is almost certain to change the criteria of choice, and most human beings deliberately seek out such experience”.¹⁶ This is a central topic for Loasby (1976), March (1978), Weick (1979, 1998), Stacey (1992), Peters (1997) and many others. Brown and Eisenhardt (1998) wrote extensively on organizational balancing between the rigidity trap (too much structure) and chaos trap (too little structure). They provide a beautiful example about how the Naskapi people of the North-Eastern Labrador fight for their survival in that unfriendly environment they live in, through caribou hunting. Many generations long experience provides them with good knowledge of the hunting tactics. Nevertheless, they experiment:

Most days, the Naskapi relied on the experience of the senior hunters in the band. But in times of high uncertainty, when game had been particularly scarce, the Naskapi set aside their experience and turned to magic. [...] So the hunter-dreamer cradled a shoulder blade from a long-dead caribou, attached it to a stick, and put it over a campfire. The band patiently waited for cracks to appear and then hunted in the direction of the cracks¹⁷.

That seems like a completely irrational way of decision-making. But, in reality, it did help them to survive, because exactly through those random trials, the Naskapi people could learn about the new hunting grounds, the ones that would have remained untried if they had persisted in following their experience all the time.

Finally, and especially taking into consideration the above argument, decision makers would rationally avoid too high levels of precision in their estimates also for a

¹⁶ *ibid.*, page 162

¹⁷ *ibid.*, page 96

simple reason that both monetary and time costs associated with further increasing it after some point would become unjustifiably high.

2.3. Satisficing Threshold and Designing without Final Goals

Although in quite a different vein, the idea that imperfect solutions to the problem can be optimally preferred to perfect ones has initially found its way into *NK Model* related literature in the work by Frenken, Marengo and Valente (Frenken *et al.* 1999). Their idea that: “if problem-solvers are ready to accept algorithms which lead to less than optimal (“satisficing”) solutions they can decrease the size (and thus the execution time) of the algorithm required to find it”¹⁸, is relevant especially to the argument on “costs of precision” raised in the previous paragraph. Even more important for the current analysis is the definition of the set of “satisficing” solutions they give: “The set of **ϵ -satisficing solution** is the set of strings whose value is at most ϵ lower than the global optimum.”¹⁹ Following Frenken (2001) we would call ϵ a “satisficing threshold.”²⁰ This idea in sequence bears upon the observation by Herbert Simon that: “In the face of real-world complexity, the business firm turns to procedures that find good enough answers to questions whose best answers are unknowable.”²¹

At least in its unintentional part, and to some extent on the intentional part as well, our vision of the issue comes very much close to the idea of satisficing threshold. However, there are significant differences.

In their treatment, the threshold applies exclusively to the final level of efficiency obtained, while the evolutionary process leading to it leaves the search rule of the original model intact. It can be argued that this is exactly the way Simon was suggesting to treat the issue. But there are several reasons why we hold a different point of view.

¹⁸ 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, page 147

¹⁹ *ibid.*, page 157

²⁰ Frenken, Koen (2001) *Understanding Product Innovation using Complex Systems Theory*. Unpublished Academic Thesis. University of Amsterdam, page 76

²¹ Simon, Herbert (1969 [1996]) *The Sciences of the Artificial*. The MIT Press, Cambridge, MA, page 28

First of all Simon to a large extent dealt with situations when many alternatives exist, some fraction of them can be sampled, but only one has to be chosen, and when it is, the search stops. So the question was mainly about how large a fraction to sample, before making the final choice. On the contrary, in the *NK* setting, we model agents who have to make a choice each period, and regardless of whether the new alternative has been accepted or not, the process continues on, and in the next period another choice situation is faced. So, unlike the former case when there is a single goal, to which an assumption of satisficing threshold can be applied, we have here the case of evolution to (generally unreachable) final goal through accepting or rejecting subgoals in the process.

Simon himself was quite skeptical about considering any goal final. He writes: “The idea of final goals is inconsistent with our limited ability to foretell or determine the future. The real result of our actions is to establish initial conditions for the next succeeding stage of action”²², and also: “A paradoxical, but perhaps realistic, view of design goals is that their function is to motivate activity which in turn will generate new goals.”²³

Organizational and technological design is indeed an open-end process, and if so, there seems to be no reason, or even a possibility, to view some steps in this process as leading to *more final* goals than others. Therefore it seems to be quite natural in the world of “designing without final goals”²⁴, to augment the idea of accepting *good-enough* end results with a mechanism of taking *good-enough* decisions in each step of organizational and technological design process.

There is actually, one final reason, why such a change calls for being made. It is quite strange indeed to consider a rule that says: continue search until a solution that is at most ε worse than the globally optimal one is found, for a simple reason that this would suggest that we actually know the exact value of such global optimum. And, quite naturally, we never really do.

²² *ibid.*, page 163

²³ *ibid.*, page 162

²⁴ originally a title of a paragraph in Simon’s book. *ibid.*, page 162

2.4. Memory Dependence: When We Was Fab vs. No Regrets

*“Back then long time ago when grass was green
Woke up in a daze, Arrived like strangers in the night
Fab - long time ago when we was fab...”*

George Harrison

*“No regrets, they don’t work
No regrets they only hurt...”*

Robbie Williams

In the previous two sections we dealt predominantly with the question of how the uncertain future returns would influence our current choices. Future consequences of today’s choice, however, are not its sole determinants. Let us now look at the other side of the coin, and see how past choices, successes and failures can hinder or influence what we are willing, and indeed capable to do today.

Despite the fact that imperfectness of local search dispels the nearly path-determined nature of the search process, through endowing the agents with a chance to deviate (be it unintentionally or intentionally) from the originally chosen path, past still plays a significant role in evaluating today’s choices.

In fact, it has in this setting a much higher explanatory power than before. In the original model’s strictly uphill walk scenario, every new technology a shift to which is being made is the most efficient one that has ever been employed by the firm throughout its history. This is pretty obvious, since if it wasn’t, a shift to it simply wouldn’t have been made. In this simplified representation of reality the efficiencies of the technologies used in the past don’t really influence the current aspiration levels by default.

In the modified version of the model presented here, downhill moves are possible, so that the efficiency of the technology currently in use does not necessarily have to be the highest of what has been encountered before. Due to the fact that a firm is not anymore assumed to be able to estimate the efficiency of an untried technology with precision, it might overestimate it, and realize the mistake only when the shift to an inferior technological configuration has already been made.

What if this does happen? Should the firm set as its aspiration level the efficiency of the currently used inferior technology, and try to get away from it to a better one as fast as they can? Should it alternatively, keep the past, more efficient technology as the benchmark with which the possible novel technologies are to be compared with, cashing

in on the knowledge of the existence of more favorable point in the landscape they possess? Should it finally base their decisions on the combination of the two?

The third option is probably the most interesting to consider, and is the optimal choice of action from a decision-maker's perspective, since balancing between the two extremes would have provided a firm with an ability to capitalize on the past, without at the same time being too rigid in treating novelty. However, the case of exploring that option brings the whole complex dynamics of aspiration level adjustments into the picture, and deserves a separate treatment (see Hovhannisian 2003c).

For the purposes of the current analysis, we would explore instead the extreme cases of "memory dependence" that we term after two songs, cited in the beginning of the section.

Agents in *When We Was Fab* setting are fully memory dependent. Once they have made a mistake that brought them to a substantially inferior technology, they will try to review their choice and get back to a better technological neighborhood. However, this would also mean that no matter how low on the technology landscape they have found themselves due to a misjudged shift to an inferior technology, they will stick to it until they find a technological configuration that according to their estimations is more efficient than the maximally efficient previously encountered one. Hence they would not agree to shift to any novel technology even a fraction worse than the one they have experienced using in their times *when they was fab*. And the efficiency of this latter technology would serve as their aspiration level.

On the contrary, the agents in *No Regrets* setting are careless about their past successes and failures, deeming the only pair of efficiencies important in making a decision on accepting/rejecting novelty the one currently in use and the one currently under scrutiny. So, even if because of a miscalculation of new technology's efficiency they actually made a downhill move, they will not *regret* the good times, but rather would do everything to find an improvement over the currently used, inefficient technology. So then, their aspiration level is always kept equal to the their current operational efficiency.

Both alternatives, even if simplifying it quite substantially, reflect the reality of decision-making, and have been analyzed in various settings in March & Olsen (1976), Weick (1976), Harrison & March (1984), Nystrom & Starbuck (1984) and Miller (1994).

The first alternative, even if more rigid, provides more certainty. While letting the agents explore and experiment in the vicinity of the currently chosen path, this setting precludes them from wandering too far away, thus in a way keeping the balance between exploration and exploitation. Second alternative is more dangerous in what a series of even slight miscalculations can lead to a very significant fall in performance of the firm on aggregate. However, the more reckless and novelty seeking behavior can allow the agents to spot far better peaks on the landscape.

Simulation Model

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. Modeling “Imperfect” Local Search Strategies

We suggest here that the search rule of the original *NK model* can be easily reformulated for the purpose of modeling “imperfect” local search strategies. As it was noted before, in the original model, the new configuration w_j is adopted instead of the current one w_i if and only if $\theta_j > \theta_i$, where θ_j is the efficiency of a newly sampled

production recipe, and θ_i is the efficiency of the production recipe currently used by the firm. It has also been noted that because of the particular way in which the original model was set up, at each period the currently used technology w_i is always the most efficient one ever tried, so that

$$\theta_{i,t} = \theta_{max,t} \text{ for } \forall t. \quad (1)$$

Let us instead consider that the firm can observe perfectly only the efficiency of the currently employed technological configuration, θ_i while observing some hypothetical level of efficiency $\bar{\theta}_j$ instead of the real value θ_j , with:

$$\bar{\theta}_j = \theta_j + \chi \varepsilon \quad (2)$$

where ε is randomly distributed in $[-0.5;0.5]$, and χ is a tunable parameter measuring the degree of imprecision (either intentional or unintentional) of the agent's estimation of the new configuration's potential efficiency level. The extreme case of $\chi=0$ reflects a *perfectly myopically rational* strategy of the agents as in the original model, while, on the opposite, a case of $\chi=1$ reflects a situation when the observed efficiency of a new configuration is maximally random.

Due to the possibility of shifting to an inferior technological configuration present in the current model modification, *equation (1)* does not hold with certainty anymore, so that we can differentiate between the rules of accepting/rejecting novel technologies between the agents in *When We Was Fab* and *No Regrets* settings.

The agents using the former strategy would adopt a new configuration if and only if $\theta_{max} < \bar{\theta}_j$, while the ones using the latter, would do so if and only if $\theta_i < \bar{\theta}_j$.

The mean value of ε is 0, so that on average firms in both settings observe the real value of each possible configuration, and are not biased as for the “direction” of imprecision. However, given that $\theta \in [0,1]$, and given the randomness of ε , the modification of the model would lead to cases when for the values of $\chi \neq 0$ and increasing towards $\chi=1$, the agent is more and more likely to either reject a configuration that is more efficient than the current one, and, more importantly, to accept configurations

moving it “downhill”. It is assumed that once the new configuration is accepted, its true efficiency level becomes perfectly observable for the firm.

Due to a restriction that the efficiency cannot be negative or have a value of more than the maximum of I , the algorithm of the model on which the simulations are run is written in such a way to assign a value of 0 for all the values of $\bar{\theta}_j < 0$, and a value of I for all $\bar{\theta}_j > I$.

2.4. Simulation Results

For all the simulation runs the value of N was kept constant equal to 20, while the value of $K \in \{0, \dots, N-1\}$. The value of χ varies for each K in the range $\{0, \dots, 0.5\}$, with an interval of 0.025 to account thus for the probability of making a mistake in evaluation of the efficiency of the new configuration in the range between 0% and 50% in each direction.²⁵ For each value of K the simulation was run 10 times, with different seeds. The efficiency of each strategy was computed as the average over the efficiency obtained by each of the 10 agents employing it. Hence for each combination of K and χ , 100 observations were obtained. Due to physical limitations, the *never-ending* process of technological and organizational change was “stopped” at 5000th step.

The first thing that we want to see is whether our intuition about the imprecision in myopic search being a virtue and not only a threat is backed up by the results of the simulation runs. *Figures 1* and *2* provide the averaged results for all the possible combinations of parameters K and χ in the *When We Was Fab* and *No Regrets* settings respectively.

<Insert Figures 1 and 2 about here>

The figures present on a 3- D plot the average efficiency levels obtained by the agents for each pair of K and χ at the last, 5000th step, as well as its 2- D top projection.

²⁵ Larger values of χ were not taken into consideration for the reason of their being clearly inferior for each K .

The first striking result of the modified model is that for no values of K , apart from the uninteresting case of $K=0$, is the perfectly precise rigid behavior (corresponding to the case of $\chi=0$) optimal.

As it can be seen from the figures, only for the cases of low complexity, corresponding to low values of K , the optimal level of χ is just marginally higher than 0%. With increasing K , the optimal levels of χ reach for the both settings a level of 10% already for $K=3$, slowly and steadily rising from 12,5% to 22,5% for more complex systems. The optimal values of χ show quite similar dynamics between the two settings, with the values for the *No Regrets* setting just slightly lower for the corresponding values of K .

In *Figures 3 and 4*, we analyze instead the efficiency obtained by the agents averaged over 5000 time steps for all the combinations of K and χ .

<Insert Figures 3 and 4 about here>

This is done in order to control for the possible change in dynamics when the whole evolution of the efficiency levels is taken into consideration rather than only the position of the agents on the technology landscape at some particular time step. The results confirm that the story behind the *Figures 1 and 2* is just as valid in this case.

However, as a prescriptive tool just telling the optimal level of imprecision is not enough. As has been discussed in length above, the level of imprecision is both a result of intentional strategic action by the agents and unintentional consequence stemming from the bounds on their abilities to evaluate the novelty. Hence, the actual value of χ can only be partially controlled by a firm. In this circumstances, we would be interested to see how large is the range of imprecision levels that lead to attainment of higher efficiency points on the landscape than it would have been in the case of employing a $\chi=0$ rigid strategy.

The results differ substantially for the two settings analyzed. In the *When We Was Fab* scenario, the deviations from the strict uphill walk are controlled to a higher extent than for the case of *non-regretting* agents. Hence even having a level of imprecision level substantially higher than the optimal, the results obtained are still preferred to the rigid case. Indeed the highest level of imprecision still superior to the case of $\chi=0$ case

reaches 25-30% already for the K as low as 3, and stays at about 40% for average-to-high levels of complexity. Taking the terminal efficiency level at step 5000 instead of the average reinforces the observed dynamics.

Alternatively, in the *No Regrets* setting, values of χ too high are too risky, because the lack of control for the recurrence of mistake-making means that the agents can drift too far downhill, and never be able to recover from the loss of efficiency. So, the intuition that this setting is more dangerous is being confirmed just as well.

Nevertheless, the dangers of this setting are being paid off by the fact that for all the levels of complexity, the technological efficiency corresponding to the optimal imprecision level is always higher than for the case of less flexible *memory dependent* case of the *When We Were Fab* way of strategizing.

So, just as was intuitively stated in the above section, there is a trade-off between the relatively higher certainty of memory dependent way of action, and the relatively higher flexibility the non regretting strategy provides for. So, once the firms are sure enough they can tune the imprecision to the optimal level, the more risky strategy can be applied in order to gain higher returns, but once the imprecision is more of an unintentional outcome of bounds on rationality, and cannot be perfectly controlled for, the relatively more rigid *When We Was Fab* strategy is preferred. The results are presented graphically in *Figure 5* for some values of K .

<Insert Figure 5 about here>

The number of local peaks rises with the complexity of the technology landscape, and as can also be observed from that figure, the higher is K the larger is the range of χ for which the *No Regrets* setting is the winning choice. This is quite an obvious result, since, the larger amount of local peaks leads to a higher chance of getting caught-up on a sub-optimal one. But then, the more is that danger, the more valuable becomes the flexibility that results from the higher frequency and boldness of experimentation the setting is characterized with.

Finally, an interesting perspective opens up when we look at *Figure 6*. What the figure shows is the comparison of the performance of the agents employing the search

heuristic of the original model ($\chi=0$), and for the both settings, the “super agents” that are able to tune their parameter χ to each of the levels of complexity, given by the parameter K , so that to attain the optimal fitness in each of those cases (χ optimal).

<Insert Figure 6 about here>

What is interesting here is not only that the figure confirms the previous results; this had to be the case for obvious reasons. The interesting observation can be made comparing the behavior of from the one hand the line representing the $\chi=0$ case, and, from the other, the lines representing the remaining two cases.

As in the original Kauffman’s model, after a very short initial increase in the average fitness, with the growing complexity of the system, its value steadily declines. This result was taken to suggest that the agents have to work on decreasing the complexity of the system through different mechanisms in order to hope for a better overall performance.

Now, this is not at all necessarily the case if we take the value of χ positive. For both of the remaining two cases in the graph, the initial growth is much longer and much steeper. Reaching its maximum at $K=4$, the average efficiency remains pretty high for the values of K up to $K=7$ (and arguably even longer so for the *No Regrets* case), and does not go below the case of $K=0$ for no values of K , however large it is. On the other hand, for the case of $\chi=0$ for all the values of $K>8$, we observe the average fitness lower than that the agent attains in the completely unconnected system, that a $K=0$ case represents.

This is important in two ways. First of all, the mechanism that the agents can design in order to decrease the complexity of the system are costly, and hence *ceteris paribus* are not desirable.²⁶ Now, using a scheme with a positive value of χ , apart from the other pluses, discussed above, thus, lowers that cost just as well. Secondly, the so-called *new economy* calls for an increased emphasis that has to be put on the cases of average and high complexity, exactly where the search heuristics employing positive values of the parameter χ are performing especially good compared to the case of the original setting of the model.

²⁶ for more detailed discussion on dangers of modular design see Hovhannisian (2003b)

Generalizations and Discussion

There are several points of possible concern that can be raised regarding the above model modifications, and in this section I would try to discuss some of them.

First of all, one might ask, what actually does the trick? In the modified version of the new configuration adopting rule, the factor of randomness seems to play a major part. The mechanism seems to be very reminiscent of the *simulated annealing* principle, well known in the literature on genetic programming²⁷.

In fact, some of the literature on business case studies that served as the starting point for the current paper indeed give a lot of attention to randomness as the surprising force helping to run the business better. Such is the example in Brown & Eisenhardt cited above.

However, my belief, and my aim in this paper was not to accentuate the role of randomness in decision making. The main point was to see how true indeed is the proposition of Simon on *deliberate experimentation* as the guiding force in decision making when facing a complex changing world where no goal is ultimately final.

To check whether random variable actually does play a role, a slight modification of the model was developed, in which randomness was absent. It was noticed that while in the model modification presented above the errors in precise estimation of the relative efficiency of the novel technological configurations could have been both in the sense of accepting a configuration that in fact was inferior to the currently employed one, or rejecting the ones in fact superior to it, the whole idea of experimentation as the guiding force suggests putting more emphasis towards the so to say “optimistic” errors.

Rejecting what might have been a better way of running the business just does not seem to be a good an idea intuitively. So then, would it be right to say that this is exactly the acquired option of accepting modifications even if they are slightly inferior to the present state of affairs that makes the difference?

The results of the simulation runs show that this is indeed the case. What was changed is again the mechanism of accepting or rejecting new adjacent configurations.

²⁷ I would like to thank Koen Frenken for bringing my attention to this point.

Instead of the one used above, the following rule was suggested: accept a new configuration w_j instead of the currently employed configuration w_i if and only if $\theta_i < \bar{\theta}_j$, where²⁸:

$$\bar{\theta}_j = \theta_j + \chi/2 \quad (8)$$

The results are extremely similar to the ones discussed above both in terms of the optimal level of χ , and the highest levels of it still superior to the case of $\chi=0$.

This suggests that this is not the randomization of the strategy, and neither the fact of mistake-making *per se* that is responsible for the results discussed above, but indeed it is the case that overly rigid structures of the perfect myopic optimization technique just do not let the decision maker gain from the advantages a more flexible scheme of a dynamic, experimenting decision making provides.

Another concern might be raised in this respect. It could be the case that omission of such factor as the *search costs* can benefit overly explorative activities, while a modification of the model to a one that accounts for those costs would also allow a conservative strategy of accepting only the configurations by some fraction superior to the current one becoming a winning one.

This is a more difficult concern to answer to, because of the difficulties of directly measuring such costs in the present algorithm of the model. Nevertheless, some conjectures can still be made in this respect. First of all, it has to be noted that search costs can be classified in two major groups:

- The shifting costs between the two adjacent technological or organizational configurations
- The costs of actually estimating whether such change is desirable.

Now, with no doubts, the search heuristic applied in the present modification of the model calls for an increase in the first group of the costs. Experimenting means more

²⁸ $\chi/2$ is taken instead of χ to account for the same magnitude of the effect when we change from a stochastic to a deterministic case.

often changes, and hence more resources have to be directed towards *shifting costs*. The magnitude of that increase can be measured in the simulation runs by the *successful mutations* parameter. And it does increase substantially when we move from the case of $\chi=0$ to positive values of χ .

Thus, depending on K , for its high enough values, the search heuristic corresponding to the optimal values of χ presupposes about 10-30 times more shifts than in the case of $\chi=0$ for the *No Regrets*, and about 5-10 times more for the *When We Was Fab* setting, and hence, the overall shift costs should be significantly higher.

However, this is not as alarming as it might seem. The thing is that for high enough values of K , the optimal value of χ is between 15-25% in both directions, which accounts for the double of that values range of desired imprecision. But then, as it has been noted above, that imprecision is a way of economizing on the costs of evaluating the relative efficiency of the yet untried adjacent technologies.

Thus, while dragging the shifting costs up, such search heuristic lowers significantly the calculation and estimation costs. Measuring the relative magnitude of these two effects unfortunately is an unsolved problem yet, but the fact that we are considering here the day-by-day small organizational and technological shifts, which are not great in magnitude, leaves us to think that the costs associated with such shifts should not be too high, and should well be balanced by the decrease in the presumed precision in estimating the yet unknown.

Consequently, one might ask, why do the agents necessarily have to stick to some given magnitude of imprecision, rather than trying to tune that parameter in accordance with the level of the complexity of the system?

Indeed, as it has been noted in the discussion of the *Figure 6* above, the “super agents” able to do so apparently perform better than the agents applying any other given search scheme. So, in a way the model restricts the possibilities of the agents. Quite obviously, a simpler system calls for more rigid scheme, because of the tradeoff between the costs of rigidity (low in this case) and the possibility it gives to reach a higher level of

efficiency, and that scheme has to become more and more flexible with an increase in the complexity.

Moreover, the possibility of tuning the parameter χ is an important advantage in other, probably even more important respect. With the evolution of technology, different and quite distinct phases change each other. Apparently then, different values of χ would be optimal depending on whether the agents are in the phase of fast and booming development of the given technology, or the technology is mature, and only slight and slow changes are being made to it.

Finally, a concern might be raised regarding the question of why don't the agents change the state of several elements at the same time. One of the reasons of not including that possibility in the model is that what I was focusing the attention on are the small day-by-day decisions, and especially since the agents are given possibility to deviate from the rigid rule of accepting new technological configurations proposed in earlier papers on the subject, a change in the state of more than one element at a time was not seen important. From the other hand, Auerswald *et al.* (2000) confirm the intuition that an effect "taking bigger steps on a given landscape" has is "like walking with smaller steps on a more rugged landscape."²⁹ So then, because we consider here the landscapes of all the possible levels of ruggedness, introducing a possibility of taking larger steps would just mean doing a double work, without expecting any new results.

Turning from *forward-looking* to the question of memory dependence, it is obvious to claim that the two cases discussed are not fully representative of how the past is weighed in real-life business practice. However, the aim of this paper was not in discussing all the range of possible strategies of organizational memory and organizational forgetting.

Conclusions

²⁹ 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, page 427

The current paper has dealt with broadly two related issues in innovation-related business strategizing.

First we discussed in much detail the question of how the uncertain future returns influence our current choices. From the one hand, the internal inconsistency of endowing the agents with too much and too little sense-making abilities, present in the original model has been criticized. It has been noted that the assumption of firm's inherent ability to precisely estimate the efficiency of a novel technology without bearing the consequences of actually having shifted to it, is quite questionable. In this way the unintentionality of imprecision entered the picture.

It was noted further that imprecision is not only a consequence of limits on human analytical skills, but a remedy from the rigidity of not realizing our own ignorance. Our claim was that once the limits to our foresight precludes the possibility to see such long-run consequences, firms might be well better off deviating from the strict rule-following behavior of accepting only uphill leading modifications of the technology in use. This statement goes in line with Weick's claim that: "Loosely coupled systems may be elegant solutions to the problem that adaptation can preclude adaptability."³⁰

In a newspaper article I read a while ago it was discussed how people turn back from the high-precision all-autonomous housing plans, to the more conventional ones. The reason was in that the hi-tech insulation of the walls and windows led to very substantial increase of allergies the owners of those houses acquired. Now this is exactly the same with overly-rigid search schemes, where the imposed precision of the rule being followed leads to the case where the tiny "viruses" that once were introduced (even if not intentionally) in the scheme, are not given a chance of leaving the system ever after, and ultimately result in its failure or stagnation.

The analysis of virtues of intentional imprecision has been present in Simon's and March's works on *deliberate experimentation* and throughout Weick's, Loasby's and Stacey's books and articles. On a more "applied" level, Nystrom and Starbuck (1984) observed that: "Experimentation offers many benefits as a central frame of reference for

³⁰ **Weick, Karl E.** (1976) *Educational Organizations as Loosely Coupled Systems*, Administrative Science Quarterly, March, vol. 21, page 7

top managers. People who see themselves as experimenting are willing to deviate temporarily from practices they consider optimal in order to test the validity of their assumption”³¹

Again based on Simon-March-Olsen line of work on “designing without final goals”, a necessity to adjust the idea of Frenken and colleagues on “satisficing threshold” has been proposed. It was argued that technological evolution is an open-end process and a shift from applying the satisficing rule to the end results to applying it to every single step of technological evolution can be necessary.

The second key task raised in the current paper was to attempt to answer how much indeed do the firms have to be dependant on their own past experience, especially so, their past successes. Guided by empirical observations such as: “Organizations succumb to crises largely because their top managers, bolstered by recollections of past successes, live in world circumscribed by their cognitive structures”³², as well as by insights gained from the work on “core rigidities” (Leonard-Barton 2000) and “perils of excellence” (Miller 1994) it was proposed to analyze memory dependence in two extreme settings.

One of them, termed *When We Was Fab* strategy, puts too much weight on the past successes, and disregards the possible current *dire straits*. The other, termed a strategy of *No Regrets*, on the opposite, treats the memory “as a pest”, as suggested by Weick (1976), and considers the only relevant data for making a choice the efficiency of the technology currently in use, and the efficiency of the technology a shift to which is possible. While the former brings more certainty into the structure, the latter gains on increased flexibility.

All the above intuitive considerations were put under test in simulation setting, using as a benchmark case Kauffman’s *NK Model* that has attracted recently much interest in evolutionary modeling. In our view while having an extremely high potential for economic applications this model suffers from the side-effects of direct translation of biological concepts into the field of social evolution.

³¹ Nystrom, Paul C. & William H. Starbuck (1984) *To Avoid Organizational Crises, Unlearn*, Organizational Dynamics, Spring 1984, page 62

³² *ibid*, page 57-58

The main point of the current analysis thus has been to provide behavioral background nested to a larger degree in economic theory and business literature. As the analysis in the current paper has shown, when some of the behavioral assumptions are tuned to the case of human-influenced (if not to say human-determined) technological evolution, a number of conclusions valid for the natural systems fail to hold, or else become locally valid, while the others grow in importance.

In the end, remembering once again Simon's idea on design without final goals, we would like to note that the model discussed in the current paper is by far not meant to be considered such a final goal. On the contrary, our belief is that *NK*-inspired research is able to give much more important and general results if we note that probably the greatest advantage it provides us with is in the fact that here, unlike either the mainstream neoclassical theory, where too much stress is put on rationality and intentionality, or the "mainstream" evolutionary economics, where the agents too often are seen as possessing no rationality or insight at all, and the purportive strategizing is substituted to a high degree with the idea of them adapting blindly to external environmental changes, here we are able to bring comparisons between agents employing different (and changing) strategies, and putting them in different environments, get the idea on how much valuable one or the other is in different circumstances.


References:

1. **Altenberg, Lee** (1994) *Evolving Better Representations through Selective Genome Growth*, in Proceedings of the IEEE World Congress on Computational Intelligence, page 182-187
2. **Altenberg, Lee** (1997), *Genome Growth and the Evolution of the Genotype-Phenotype Map*, page 2 of 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, pp. 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. **Brown, Shona L. & Kathleen M. Eisenhardt** (1998) *Competing on the Edge: Strategy as Structured Chaos*, Harvard Business School Press, Boston, MA
5. **Cyert, Richard M. & James G. March** (1963) *A Behavioral Theory of the Firm*. Blackwell Business
6. **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
7. **Foray, Dominique & Arnulf Grübler** (1991) *Morphological Analysis, Diffusion, and Patterns of Technological Evolution: Ferrous Casting in France and the FRG*, in **Nakićenović, N. & Grübler, A.** (eds.) *Diffusion of Technologies and Social Behavior*. Springer Verlag, pages 407-431
8. **Fleming, Lee & Olav Sorenson** (2001a) *Technology as a Complex Adaptive System: Evidence from Patent Data*. Research Policy 30, pages 1019-1039
9. **Fleming, Lee & Olav Sorenson** (2001b) *The Dangers of Modularity*. Harvard Business Review, September, pages 20-21
10. **Fleming, Lee & Olav Sorenson** (2003) *Navigating the Technology Landscape of Innovation*. MIT Sloan Management Review, Vol. 44 Issue 2, pages 15-24
11. **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
12. **Gavetti, Giovanni & Daniel Levinthal** (2000) *Looking Forward and Looking Backward: Cognitive and Experiential Search*, Administrative Science Quarterly, 45, pages 113-137
13. **Hannan, Michael T. & John H. Freeman** (1984) *Structural Inertia and Organizational Change*, American Sociological Review, 49, pages 149-164
14. **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

15. **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
16. **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
17. **Hovhannisian, Karén** (2002c) *The Virtue of Being Imprecise: Modeling Innovation-Related Business Strategies through Simulation Analysis*. Paper presented at an annual PhD meeting at Certosa di Pontignano, Siena, June 24-25
18. **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
19. **Hovhannisian, Karén** (2003b) *Breadth and Depth of Search on Technology Landscape*. Paper presented at MERIT Seminar, Maastricht, April 29
20. **Kauffman, Stuart; José Lobo & William G. Macready** (2000) *Optimal Search on a Technology Landscape*, Journal of Economic Behavior & Organization, vol. 43, pages 141-166
21. **Kauffman, Stuart** (1993) *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press. New York
22. **Kauffman, Stuart & Simon Levin** (1987) *Towards a General Theory of Adaptive Walks on Rugged Landscapes*. Journal of Theoretical Biology 128, pages 11-45
23. **Kauffman, Stuart; William G. Macready & Emily Dickinson** (1994) *Divide to Coordinate: Coevolutionary Problem Solving*, Unpublished Manuscript, Santa Fe Institute
24. **Kauffman, Stuart & William G. Macready** (1995) *Technological Evolution and Adaptive Organizations*, Santa Fe Institute Working Paper 95-02-008
25. **Kirkpatrick, S., C. D. Gelatt, Jr. & M.P. Vecchi** (1983) *Optimization by Simulated Annealing*, Science, Number 4598, pages 671-680
26. **Leonard-Barton, Dorothy** (2000) *Core Capabilities and Core Rigidities: A Paradox in Managing New Product Development*, Strategic Management Journal, Vol. 13 Issue 5, pages 111-128
27. **Levinthal, Daniel** (1997) *Adaptation on Rugged Landscapes*, Management Science, 43, pages 934-950
28. **Levinthal, Daniel & James G. March** (1981) *A Model of Adaptive Organizational Search*, Journal of Economic Behavior and Organization, 2, pages 307-333
29. **Levitan, Bennett & Stuart Kauffman** (1995) *Adaptive Walks with Noisy Fitness Measurements*. Molecular Diversity, 1(1), pages 53-68
30. **Loasby, Brian** (1976) *Choice, Complexity and Ignorance*. Cambridge University Press
31. **Lobo, José & William G. Macready** (1999) *Landscapes: A Natural Extension of Search Theory*, Santa Fe Institute Working Paper 99-05-037 E
32. **March, James G.** (1978) *Bounded Rationality, Ambiguity and the Engineering of Choice* Bell Journal of Economics, Vol. 9 Issue 2, pages 587-599
33. **March, James G.** (1991) *Exploration and Exploitation in Organizational Learning*, Organizational Science, vol. 2, No. 1, pages 71-87

34. **March, James G. & Herbert Simon** (1958) *Organizations*, New York: J. Wiley
35. **March, James G. & Olsen, J.P.** (1976) *Ambiguity and Choice in Organizations*. Bergen, Norway: Universitetsforlaget
36. **March, James G. & Richard J. Harrison** (1984) *Decision Making and Post-Decision Surprises*. *Administrative Science Quarterly*, Vol. 29 Issue 1, pages 26-43
37. **Miller, Danny** (1994) *What Happens After Success: The Perils of Excellence*, *Journal of Management Studies* 31:3, pages 325-358
38. **Nelson, Richard R. & Sidney G. Winter** (1982) *An Evolutionary Theory of Economic Change*, Harvard University Press
39. **Nystrom, Paul C. & William H. Starbuck** (1984) *To Avoid Organizational Crises, Unlearn*, *Organizational Dynamics*, Spring, pages 53-65
40. **Peters, Tom** (1997) *The Circle of Innovation*. Random House Trade Paperbacks
41. **Potts, Jason** (2000) *The New Evolutionary Microeconomics: Complexity, Competence and Adaptive Behaviour*. Cheltenham: Edward Elgar
42. **Rivkin, Jan W. & Nicolaj Siggelkow** (2003) *Balancing Search and Stability: Interdependencies among Elements of Organizational Design*, *Management Science* 49(3), pages 290-311
43. **Roberts, Kevin, Martin L. Weitzman** (1981) *Funding Criteria for Research, Development, and Exploration Projects*, *Econometrica*, Vol. 49 Issue 5, pages 1261-1294
44. **Siggelkow, Nicolaj & Daniel Levinthal** (2003) *Temporarily Divide to Conquer: Centralized, Decentralized and Reintegrated Organizational Approaches to Exploration and Adaptation*, unpublished manuscript
45. **Simon, Herbert** (1947) *Administrative Behavior: A Study of Decision-Making Processes in Administrative Organizations*. The Free Press
46. **Simon, Herbert** (1969) *The Sciences of the Artificial*. The MIT Press
47. **Stacey, Ralph** (1992) *Managing Chaos: Dynamic Business Strategies in an Unpredictable World*, Kogan Page, London
48. **Teece, David J.** (1986), *Profiting from Technological Innovation*, *Research Policy*, 15:6, pages 285-305
49. **Valente Marco** (2000) *Evolutionary Economics and Computer Simulations: A Model for the Evolution of Markets*, University of Aalborg, PhD Thesis
50. **Vishwanath Tara** (1992), *Parallel Search for the Best Alternative*, *Economic Theory*, Vol. 2 Issue 4, pages 495-508
51. **Weick, Karl E.** (1976) *Educational Organizations as Loosely Coupled Systems*, *Administrative Science Quarterly*, March, vol. 21, pages 1-19
52. **Weick, Karl E.** (1979) *The Social Psychology of Organizing*. Reading, Mass.: Addison-Wesley
53. **Weick, Karl E.** (1998) *Improvisation as a Mindset for Organizational Analysis*, *Organization Science*, Vol. 9 Issue 5, pages 543-556
54. **Weitzman, Martin L.** (1979) *Optimal Search for the Best Alternative*, *Econometrica*, Vol. 47 Issue 3, pages 641-656

Figure 1: When We Was Fab Setting. Average Efficiency as a Function of K and ksi (period 5000)



2-D Projection of Figure 1

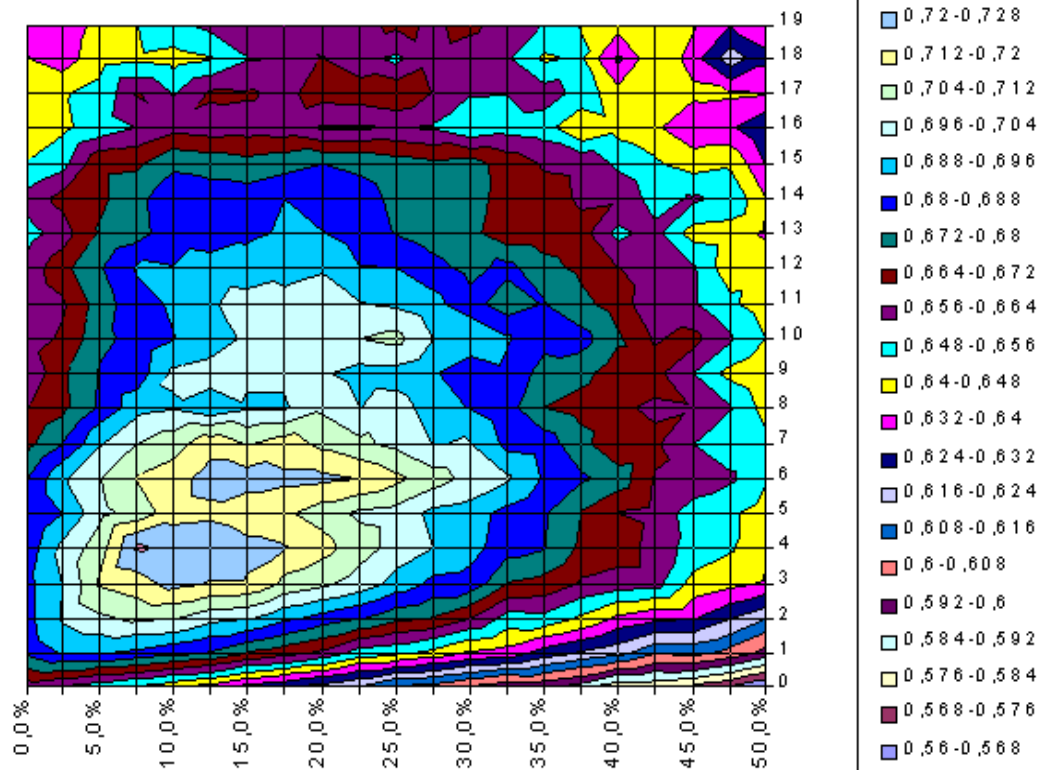



Figure 2: No Regrets Setting. Average Efficiency as a Function of K and ksi
 (period 5000)



2-D Projection of Figure 2

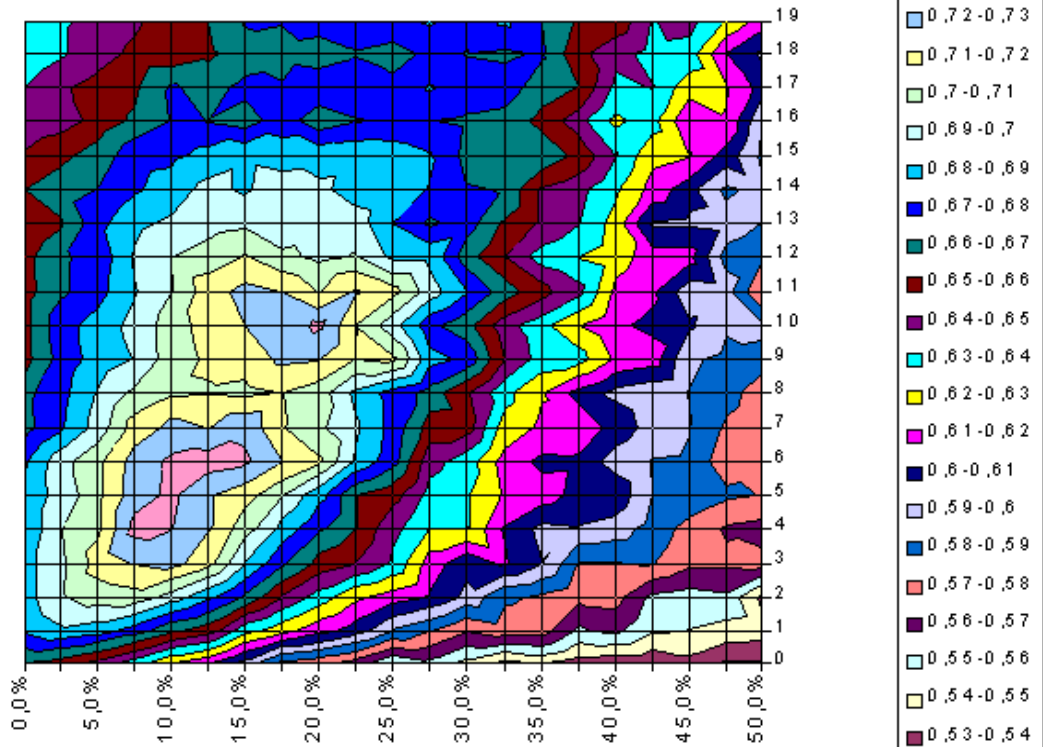




Figure 3: When We Was Fab Setting. Average Efficiency as a Function of K and ksi (average over 5000 periods)




2-D Projection of Figure 3



| |
|--------------|
| 0.71-0.7175 |
| 0.7025-0.71 |
| 0.695-0.7025 |
| 0.6875-0.695 |
| 0.68-0.6875 |
| 0.6725-0.68 |
| 0.665-0.6725 |
| 0.6575-0.665 |
| 0.65-0.6575 |
| 0.6425-0.65 |
| 0.635-0.6425 |
| 0.6275-0.635 |
| 0.62-0.6275 |
| 0.6125-0.62 |
| 0.605-0.6125 |
| 0.5975-0.605 |
| 0.59-0.5975 |
| 0.5825-0.59 |
| 0.575-0.5825 |
| 0.5675-0.575 |
| 0.56-0.5675 |

Figure 4: No Regrets Setting. Average Efficiency as a Function of K and k_{si}
 (average over 5000 periods)



2-D Projection of Figure 4

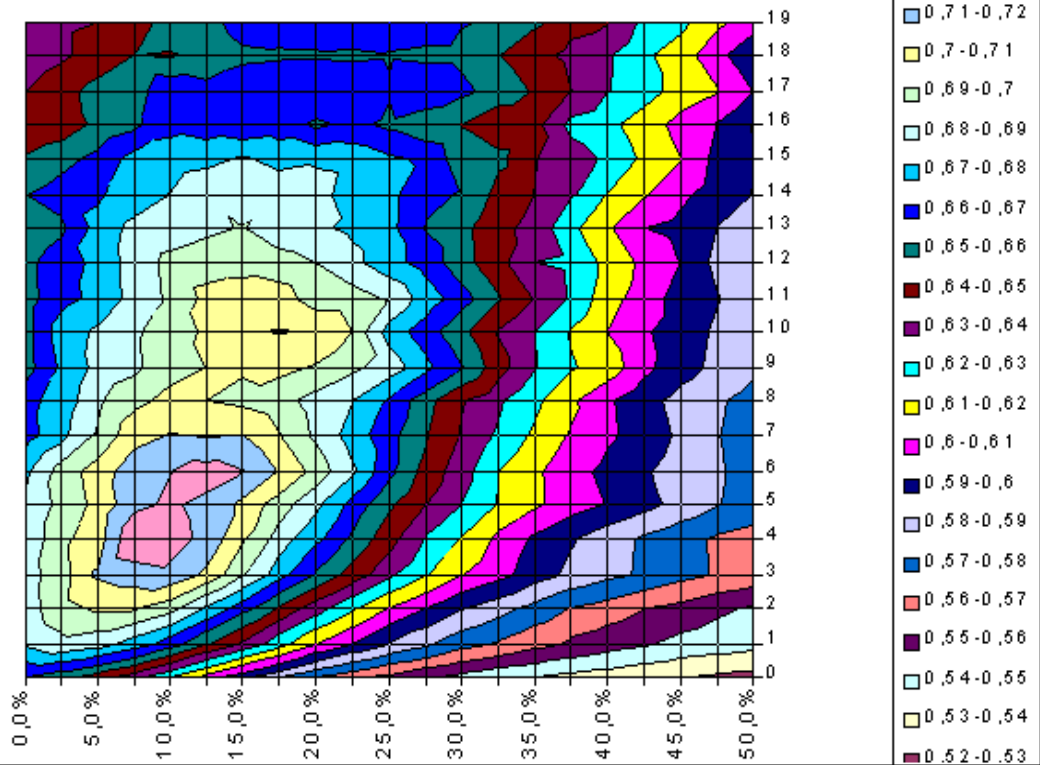


Figure 5: No Regrets vs. When We Was Fab. Average Efficiency over 5000 Periods for some K
 (Series 1: No Regrets Setting; Series 2: When We Was Fab Setting)




Figure 6: Comparison of Highest Average Efficiency Levels Attained by Agents Using Different Strategies (as a Function of K)

