# NK Modeling Methodology in the Strategy Literature: Bounded Search on a Rugged Landscape

***Martin Ganco***
College of Business
University of Illinois at Urbana-Champaign
350 Wohlers Hall, 1206 S. Sixth Street
Champaign, IL 61820

***Glenn Hoetker***
College of Business
University of Illinois at Urbana-Champaign
350 Wohlers Hall, 1206 S. Sixth Street
Champaign, IL 61820

Version: December 2008

***Abstract***

We discuss recent methodological advances in the NK modeling in the Strategy literature and analyze issues related to its current use including different implementation algorithms, relative versus absolute performance, establishing significance of simulation results and long- versus short-term performance measurements. To facilitate cross-pollination of ideas, we point to advances and extensions of the model developed in other fields that could be effectively utilized to answer Strategy-related questions. These include modeling the strength of interaction, varying the importance of decision elements, utilizing alternative functional forms, incorporating endogeneity in *N* and *K* parameters, and embedding the NK model in a broader dynamic framework.

## Introduction

Since the publication of the seminal NK model (Kauffman, 1993, 1995), the number of studies in the field of Strategy utilizing it has experienced a steady growth (e.g. Ethiraj and Levinthal, 2004; Ethiraj *et al.*, 2008; Frenken, 2001; Ganco and Agarwal, 2008; Gavetti and Levinthal, 2000; Lenox *et al.*, 2006; 2007; Levinthal, 1997; Rivkin, 2000; Rivkin and Siggelkow, 2003; Siggelkow, 2002; Siggelkow and Levinthal, 2003; Siggelkow and Rivkin; 2006). The NK model has been introduced to the field by Dan Levinthal (1997) who showed that the existence of interdependencies among firm choices can explain persistent organizational heterogeneity. This influential paper gave impetus to 30 modeling papers aimed at answering Strategy-related questions published to date in leading management journals.[1] These 30 modeling papers have so far received about 800 citations. The approach using the NK model has been successful in defining a small but growing niche answering questions arising from decision interdependence and complexity in various settings. The strength of the NK modeling has been in providing a relatively simple and replicable methodology addressing problems that are typically poorly answered empirically.

Despite this success, there was virtually no systematic methodological work on the NK model published in Strategy journals. The researchers typically develop modifications and extensions of the original framework developed by Kauffman (1993, 1995) with relatively little lateral conversation related to model methodology. The issues emerging from modifications seem to be often addressed through solutions idiosyncratic to a particular researcher. Such approach may be natural in the early stages of the NK model use within our field but as it slowly matures the benefits arising from the cross-pollination of ideas at the technical and methodological level may be significant. The purpose of our paper is to enable such knowledge spillovers in search for the best practices. Consequently, our paper has several objectives. First, we review and summarize the NK methodology as a modeling

[1] From these, 9 were published in Management Science, 7 in Organization Science, 4 in Administrative Science Quarterly, 4 in Research Policy, 3 in Strategic Management Journal, 2 in Academy of Management Journal and 1 in Academy of Management Review.

framework in an accessible manner. Second, we discuss its use in the Strategy literature from the methodological perspective. Third, we discuss advances and modifications of the NK model developed in other fields, evaluate them and propose their application within our field.

Our paper can be also seen as building on the extant discussion on how to best utilize the simulation modeling frameworks (Davis *et al*., 2007) and is primarily targeted to readers already having some familiarity with the NK modeling literature.

For readers interested in the theoretical contributions of the NK model, several excellent reviews of the modeling literature have been recently published (Chang and Harrington, 2006; Sorenson, 2002). Thorough introductions of the NK model are also readily available (Kauffman, 1993, Ch. 6; Levinthal, 1997; Rivkin, 2000).

## The Background of the NK Model

The NK model consists of two main components – the NK landscape and the agent(s) that searches the landscape. The characteristics of both, the NK landscape and the agent are controlled by model parameters. The parameters *N* and *K* characterize the landscape – the problem space that has to be searched by the agent(s). The agent(s) is controlled by defining its search behavior – by setting the rules on how it performs the search through the landscape.

In the original NK model (Kauffman, 1993) the parameter *N* stands for the number of alleles in the genome that can be either turned on or off. The *K* then controls the density of epistatic connections (called *pleiotropy*) – the linkages between the individual alleles. Within the social science applications the notion of alleles is replaced by decisions and epistasis by interdependence between decision elements. *N* then represents the number of decisions that have to be made and *K* controls how connected they are. The binary bits (0 or 1) of the decision vector (there are *N* bits) may represent broadly defined organizational decisions (Rivkin, 2000) such as those related to firm strategy, organizational form, product design or, more specifically, only to one type of decision (e.g. product design in Ethiraj and Levinthal, 2004 or organizational form in Levinthal, 1997). The value of each binary bit – either zero or one – represents a taken course of action about the particular decision (e.g. choice A vs. choice B is chosen).

The notion of search in the original NK model represents the process of population-level genetic mutation and recombination. Within the social science applications the notion of genetic mutation has been replaced by adaptive purposive search behavior of agents though it may be complemented by population level dynamics driven by selection (Levinthal, 1997). The agents making the decisions typically represent an individual (Gavetti and Levinthal, 2000), set of individuals (Rivkin and Siggelkow, 2003) or an entire organization (Levinthal, 1997).

## NK Model Components and Their Implementation: Brief Description

We follow with a brief discussion of the NK model components – the NK landscape, the agent search as well as the issues related to their implementation. Discussing implementation is necessary before we can tackle either the model use within the field or its possible extensions since many modeling choices are driven by implementation issues.

At the same time, the challenges arising from the model implementation have rarely been discussed in depth and the current paper is an excellent venue to accomplish such goal.

### ***NK Landscape***

The NK landscape defines the space which agents search for better payoffs. There are two main parameters of the space, $N$ and $K$. The parameter $K$ measures the degree of interdependence or coupling between the $N$ elements of the decision vector, i.e. the performance contribution of each element of the decision vector $x_i$, $i=1 \ldots N$, is affected by $K$ other elements $x_j$, where $j$ is not equal to $i$. The performance contributions of each decision element are determined by the payoff function which works as following: for instance, when there exists a coupling between the decision element $x_1$ and element $x_3$ where $x_1$ is the focal element ($x_3$ affects $x_1$) then a change in decision $x_3$ (alternative B instead of A is chosen) will cause the change in the payoff contribution of the element $x_1$ (the value is simply redrawn from the underlying distribution). When the focal decision $x_1$ is coupled to many ($K$) other decisions, its payoff will be redrawn whenever any of the coupled decisions change. The overall payoff of the entire decision vector is the mean of the payoff contributions of its individual decisions elements. The mapping from the space of the combinations (ordered strings) of zero and ones to payoffs define the NK landscape (mapping over $N$-dimensional binary space). The distance on the landscape is typically defined in simple metric using the number of bits in which the two vectors differ (at most $N$). The immediate neighborhood (local) for each vector (point in the space) is defined as all the vectors that differ at most by one bit ($N$ neighbors). High value of $K$ then implies a "rugged" landscape arising from the underlying high interdependence of decision elements (which can be also termed as coupling or complexity). In such case, a change of single bit may cause a dramatic change in the overall payoff since it triggers redrawing of the payoffs of $K$ other elements in addition to changing its own performance contribution. On the other hand, when $K$=0 the space is smooth since the change of any single bit causes only its own contribution to be redrawn.[2]

To describe the model more formally, the NK landscape is characterized by the correspondence mapping of the vector $\boldsymbol{x}$ in the decision space to the outcomes (payoffs). The landscape is a mapping from the set $\boldsymbol{X} = \{0,1\}^N$ to $R_+$ An element $\boldsymbol{x} \in \boldsymbol{X}$ is a vector of binary digits of length $N$. The mapping assigns to each $\boldsymbol{x} \in \boldsymbol{X}$ a payoff $\pi(x) \in R_+$. The mapping $\pi$ depends on the parameter $K$, with $\pi(x, K)$ reflecting the interdependence of the individual components of $\boldsymbol{x}$. The change in the payoff contribution of the $i^{th}$ component is influenced by the change in the $i^{th}$ decision $x_i$, and by the changes in K other components of $\boldsymbol{x}$. If $K = 0$, there are no interdependencies and the $\pi(.)$ function is additive. The mapping (and the landscape) is generated by assigning a payoff $\pi_i(.)$ which is a random number from a given distribution to each decision $x_i$, $i = 1,\ldots,N$ and each instance when either $x_i$ changes or some of the $K$ decisions that are associated with $x_i$ change. Typically, a uniform distribution over [0, 1] is used but since the focus tends to be on ordering of the payoffs and as $N$

[2] This is at most $1/N$ when the uniform distribution is used for generating payoffs.

increases the payoff distribution always converges to normal (due to the Central Limit Theorem), the choice of the underlying distribution is not crucial. The mapping for a particular vector $\boldsymbol{x}$ is given by

$$\pi(x,N,K)=\frac{1}{N}\sum_{i=1}^{N}\pi_i(x_i;x^1_{j(i)}...x^K_{j(i)}),\, i\notin j(i)\,, \qquad [1]$$

where for any $i$ we obtain a vector of indexes $j(i)$ mapping from $N$ *to* $N^K$. None of the indexes of $j(i)$ can be equal to $i$. The notation $x^k_{j(i)}$ means that the index of $x$ is the $k^{th}$ element of the vector $j(i)$. To create an overall mapping, one needs to randomly generate $2^{K+1}N$ payoff values. The structure of the mapping $\pi(.)$ is often depicted as a matrix called the interaction or influence matrix. The rows and columns represent the individual decisions. The matrix has ones in all entries that are affected by a particular decision (typically ones are in the entries where the row is affected by the column). For instance, for $K = 0$, the interaction matrix is an $N \times N$ identity matrix and for $K = N-1$ it is $N \times N$ matrix of ones. In many of the initial model specifications, the independencies within the interaction matrix are either randomly distributed (Kauffman, 1993, 1995, Rivkin, 2000) – the elements of the vector $j(i)$ are generated randomly from the interval $[1, N]$ with $i \notin j(i)$ or; only first $K/2$ neighbors on each side of the focal decision are linked (Kauffman, 1993, 1995), $j(i) = \{i-K/2,\ldots, i-2, i-1, i+1, i+2\ldots i+K/2\}$ – the interaction matrix has ones across $K/2$ diagonals to the left and to the right of the main diagonal.[3] Figure 1 shows examples of interaction matrices with $N = 5$, $K = 2$ and two types of link distribution.

[Insert Figure 1 about here]

However, many modifications of the NK model exploit specific patterns in the distribution of linkages – aimed at modeling modularity (Ethiraj and Levinthal, 2004, Ethiraj *et al.*, 2008) or special hierarchical patterns (e.g. Rivkin and Siggelkow, 2007).

The "peaks" or local optima on the NK landscape are defined as a configuration of elements of the decision vector such that it is not possible to improve the decision's overall

---

[3] In such case, the space may "wrap around" at the boundaries when $i<K/2+1$ or $i>N-K/2$ as indicated in Figure 1.

payoff by performing a given type of search.[4] The *N*-dimensional landscape is typically depicted either as a cube (Kauffman, 1993, 1995 for $N = 3$) or only illustratively in two dimensions where the payoffs are shown as a rugged surface over this space (with many peaks and valleys) if *K* is high and a smooth one if *K* is low. Interestingly, Rivkin and Siggelkow (2007) utilize a visualization of a subspace of the actual NK space in which they transform the *N*-dimensional binary space into two-dimensional space in which each axis has $2^{N/2}$ points by simply ordering some of the neighboring points along each axis. Despite such visualization distorts the ruggedness of the space (by reducing the local neighborhood of each interior point from *N* to 4) and the visualized spaces appear more rugged than they actually are (and this distortion increases with *N*) it is obviously the best we can do in displaying an *N*-dimensional binary space in two dimensions. Figure 2 shows a visualization of a space with $N = 6$, $K = 2$ and randomly distributed linkages.

[Insert Figure 2 about here]

***Implementation of the NK Space***

The most challenging and computationally demanding task in the design of the NK model is the generation of the NK landscape. Three approaches to this problem have been developed:

a) ***Generate the NK space at the beginning of each simulation run***. This is the most natural and straightforward solution. The researcher first creates an NK landscape and then lets agent(s) search it. Having the matrix of all possible positions on the landscape (the *N*-dimensional binary space) and their corresponding payoffs (let's define this as a Reference matrix) available before the agent(s) starts searching is very practical – to obtain payoff for a given position, only a table lookup in the Reference matrix is necessary. The global payoff of the landscape is also straightforward since all that is needed is to select the maximum payoff and pick the appropriate entry in the corresponding position. However, this approach has a significant computational problem. The length of the reference matrix increases

[4] As we discuss below, the most common is a local myopic search which proceeds by an alteration of a single decision element.

exponentially with $N$. Its size is $2^N$ so for $N = 10$, it has only 1024 rows; for $N = 20$ it has already over 1 million rows. Searching the entire matrix each simulation step for each agent can be computationally daunting even for small $N$ like $N = 10$.

*b)* ***Generate the NK space "on-the-fly"***. One possible solution of such computational problem is to generate the Reference matrix only partially. The predominantly used approach (Kauffman, 1993, 1995; Altenberg, 1994, 1996) is to generate the Reference matrix gradually as agents move through the NK landscape. In other words, only positions that are visited or tested by agents are generated and stored in the Reference matrix. The idea is simple. When an agent needs to know the payoff for given decision vector it first looks in the existing Reference matrix. If the position is already stored it selects the appropriate payoff value. If it is not stored it generates the value anew (utilizing the mapping function which derives the new payoff based on the interaction matrix and the existing payoffs in the stored matrix) and appends the generated payoff at the end of the Reference matrix. This approach is dramatically faster than searching and storing the entire NK space in the Reference matrix. Unfortunately, the search speed slows down as the simulation time progresses since the Reference matrix gets larger as agent visits more new positions in the space. For larger $N$ (e.g. $N > 20$), the Reference matrix length increases rapidly since it is unlikely that agent revisits the existing points in the space – values for new positions are frequently generated and added to the Reference matrix.

***c) Generate Payoff Values "on-the-fly"***. This method is technically most advanced and involves neither storing of positions nor payoff values. Not even the revisited positions are stored and the Reference matrix does not exist within this algorithm. The idea was proposed by Altenberg (1994, 1996) and involves generating seeds of a random number generator according to the rules of the NK space (utilizing the correspondence mapping). When two positions of the space are supposed to have the same payoff contribution $\pi_i$ the algorithm generates the same integer for both positions. This in turns implies that when the random number generator is seeded with the same integer it will generate the same random number. Altenberg (1994, 1996) implemented the algorithm for a specific type of a random number generator in a specific programming environment that required $N \leq 32$. However, a more general version of this algorithm has been recently developed (Ganco and Agarwal, 2008, Ganco, 2008) that allows for any *N* and which can be implemented in any programming language.

The "Altenberg/Ganco" algorithm is inefficient for small values of *N* compared to the algorithms a) and b) but superior for larger *N* and long simulation runs. Importantly, this algorithm also allows endogenizing *N* and *K* though some caveats apply (see below). Additionally, a potential drawback of the algorithm is that it can generate errors. Depending on how large integers the programming environment allows it may lead to occasional Type II errors – two values that should be generated as different are generated as equal. Such occurrence should be a probability zero event given the structure of the payoff function. This problem is related to the range of possible integer values available in a given programming environment (that can be used as seeds). However, the typical 32-bit storage of integer data allows integers of up to $2^{31}$-1 which is sufficient for the purposes of an NK space. Even for very large *N* (*N*~100) the frequency of errors is less than 0.1%. Consequently, this issue does not have practical significance in most settings.

***Nature of Search over the NK Landscape***

The payoff dynamics in the NK model results from interplay between the attributes of the NK landscape and properties of the search over this space. The parameters controlling the NK space typically define the environment and the parameters or rules of search define the agent behavior within the environment. Since the focus of the model is on the optimization process and path dependencies, how modeler defines the "boundedness" of the agent's search strongly determines the dynamics of the model. For instance, if the search is completely unbounded and the agent is fully rational, it can "see" all possible points on the landscape and the NK model becomes trivial – the agent immediately selects the global optimum. As a result, any non-trivial search over the NK landscape is always to some extent bounded and myopic.

The studies using the NK models usually look at how attributes of the search space affect performance while keeping the nature of the search fixed (e.g. Ethiraj and Levinthal, 2004; Ganco and Agarwal, 2008), fix space and look at changes in the nature of search (Kollman *et al.*, 2000) or vary both (Rivkin and Siggelkow, 2003).

## Agent-level Search

The most common type of search used in Strategy literature is a simple local search which represents an adaptive behavior of agent. One step of the local search is performed by alteration of a single random bit of the decision vector. The agent performing local search thus considers making a single step within its immediate neighborhood. If this change implies a higher payoff, the agent takes the move. If the payoff is lower or equal, the new vector is discarded and the agent stays at the original position. The agent thus exhibits "hill-climbing" adaptive behavior. Local search is the slowest, most bounded and most myopic type of simple search behavior than one can implement within the NK model.

Other possibilities include:

a) "greedy" local search or steepest ascent search simply accelerates local search – the agent selects in one step the best position from all surrounding positions – those that differ by one altered decision,

b) agent is assumed to make larger "jumps" – consider positions that differ by more than one decision element (it can look beyond immediate neighborhood). The search radius (number of altered decisions it can consider) is controlled by a parameter. The decisions to consider are typically generated at random and the number of different positions the agent can consider in one step may be controlled by another parameter (Siggelkow and Rivkin, 2003). Increasing the search radius or increasing the number of considered decisions not only accelerates search but makes agent more powerful – it can find better positions on the landscape by traversing local minima,

c) the stronger version of the previous type of search is used to model cognition (Gavetti and Levinthal, 2000; Gavetti, 2005; Sommer and Loch, 2004). Within this type of search, the agent is assumed to be less bounded along some dimensions of the space – which are called cognitive dimensions – has greater search radius over them or can consider more alternatives in one step. Since the agent needs to evaluate all decision elements to obtain the payoff, the models either assume that the agent creates unbiased expectation over the non-cognitive locally searched dimensions (Gavetti and Levinthal, 2000; Gavetti, 2005)[5] or simply take the starting solution in the non-cognitive dimensions as given (Sommer and Loch, 2004).

## Population Level Search

The NK space can be searched by many agents at the same time. As opposed to agents receiving immediate feedback when considering a move, agents may be completely "blind" and the adaptation can occur at the population level. Such dynamics directly corresponds to the biological evolution where the decision vector is analogous to the genome. The agents then typically change through genetic algorithm (GA) (Holland, 1992; Koza, 1992; Mitchell, 1996). GA's are much more frequently used outside of the Strategy domain likely due to the unclear mapping between the genetic recombination and decision-making process of social agents. Genetic algorithm typically includes a mutation (change in a random bit regardless of its payoff implications) and a recombination at a (usually randomly chosen) crossover point of the two decision vectors originating from two agents searching the same space. The payoff feedback then occurs at the population level through selection. The worst performers are selected out while the best performers replicate in the population (directly or through recombination).[6]

Within the Strategy literature, the agent-level adaptive search is often combined with selection (Levinthal, 1997; Levinthal, 2006, Sommer and Loch, 2004; etc.). The population level selection mechanisms then complement the adaptive search behavior of individual agents.

[5] This is accomplished by averaging payoffs over all possible solutions in the non-cognitive dimensions consistent with a given solution in the cognitive dimensions
[6] The class of GA's using crossovers and mutations combined with the population level dynamics is the most powerful non-cognitive type of search available for searching NK spaces.

### ***Implementation of the search through the NK space***

From the implementation perspective, implementing the search is typically straightforward since search is defined by a simple set of rules that can be represented through a single main loop. As an example, we can briefly describe the algorithm for a simple local search with non-interacting agents: in each step (let's assume period *t*) each agent first obtains the payoff of the decision stored for the agent in period *t-1*. Then agent randomly chooses one element of the decision vector and switches its value (from 1 to 0 or from 0 to 1) and obtains its new payoff (using one of NK space implementation algorithms discussed above). If the payoff of the new altered decision is strictly greater than the current payoff the agent moves to the new position. The new decision vector and the new payoff value are stored under time *t* in the agent data array. Using the rule of strict inequality is not absolutely necessary but it avoids oscillations between positions having equal payoff values which may occasionally occur.[7] If the new payoff is less than or equal to the current payoff the new decision is ignored and the existing decision and payoff are stored – the decision and the payoff values are simply copied from *t-1* into *t* cells of the agent array.

## NK Model as Used in the Strategy Literature

After briefly discussing the model and some of its implementation issues we proceed with the discussion of the model use within the Strategy literature. Due to the focus of our paper we concentrate on subtle modeling and sometimes technical issues that may not be explicitly spelled out in a typical Strategy journal publication.

### ***Relative versus Absolute Performance***

Performance - modeled as the payoff over the NK landscape - is the single outcome variable of the model. There are two ways how one can measure and report the performance in an NK model. One can measure either *Absolute Performance* - which is simply the value of the payoff mapping $\pi$ (equation [1]), or *Relative Performance* – which is the ratio between the payoff $\pi$ and the global optimum on a given landscape. The relative performance thus

[7] Weak inequality leads to "neutral mutation" which can have benefits under some conditions.

measures the percentage of the global max that the agent has achieved or, in other words, the ratio of the current agent's performance to the performance if it happens to discover the best possible solution on a given NK space. Since the values of the global payoff and the payoff discovered by a searching agent change non-linearly with changes in *N* and *K*, measuring and reporting absolute versus relative performance may in some cases lead to different predictions. The papers within Strategy literature tend to use both types of performance measures (e.g. Ethiraj and Levinthal, 2004; Ganco and Agarwal, 2008; Siggelkow and Rivkin, 2003). Such approach is reasonable as long as the results are interpreted as conditional on the given NK landscape. This tends to be the case in most of the simulation studies utilizing the NK model which assume static (Siggelkow and Rivkin, 2003), exogenously driven (Rivkin and Siggelkow, 2006) or environment driven by simple selection pressures (Levinthal, 1997). However, as researchers move toward interactive environments (Caldart and Ricart, 2008; Ganco and Agarwal, 2008; Lenox *et al.*, 2006, 2007) or attempt to endogenize the *N* and *K* parameters by linking them to agent search behavior the choice of the performance measure becomes important.

The modelers are typically interested in some statistic related to the performance measure – most frequently the mean and the variance. In Figure 3, we plot the absolute and relative mean performance in period 50 as well as the global max for an average of 1,000 landscapes for $N = 10$ and *K* between 0 and 9. The absolute mean performance has an inverted U-shaped relationship while the relative performance decreases with K. The value of the global max increases with *K*. Similar non-linear relationships exist if one looks at the standard error (Figure 4) or at the same relationship with fixed *K* and varied *N* or for the fixed $K/N$ (Ganco, 2009).

Reporting relative versus absolute performance implies same ordering conditional on *N* and *K* (i.e. statements like "type A agent performs best when the environment is complex" are valid based on either relative or absolute performance). Reporting relative performance also provides the information on whether the agent has reached the best possible solution in a given space (global maximum) or, in terms of the percentage, how far is it from this global peak. It thus gives the information of how well the agent performs relative to how well it *could* perform. But relative performance provides no information relevant for comparing agents' performance across landscapes. In such case it would always imply that solving simple problems is better. Consequently, as one moves to more interactive settings where agents solving different types of problems interact and compete together in both theoretical (Caldart and Ricart, 2008; Ganco and Agarwal, 2008; Lenox, *et al.*, 2006, 2007) and empirical settings (Fleming and Sorenson, 2001), the relative performance cannot be used. The same applies to environments where *K* is endogenized and depends on agent's behavior.

For instance, Ganco and Agarwal (2008) operationalize diversifying entrants as having higher *K* relative to entrepreneurial startups. Comparing these two groups of agents is possible only through absolute performance. The fact that some agents face smoother (lower *K*) spaces and are closer to their global maxima (higher relative performance) does not necessarily imply that their optima are better in absolute terms. Similarly, if one endogenizes interdependence within the environment the agents would be driven to search perfectly smooth spaces if relative performance is used despite more complex spaces provide more opportunities.

[Insert Figure 3 about here]

Similar differences exist in the interpretation of variance between the relative and absolute performance (we report S.E. below in Figure 4). The variance of the absolute performance decreases monotonically with *K* as opposed to the variance of the relative performance that has inverted U-shaped relationship and increases rapidly for low *K*. The variance of the absolute performance always exceeds the one of the relative performance. The variance of the relative performance can be seen as *within-landscape* variance (driven by variation in terms of how agents search a given space) and the variance of the absolute performance contains both – *within-* and *across-landscape* variance components (driven in addition by variation in properties of the NK spaces across independent runs). For absolute performance, as K increases from small values, the large decrease in initially large across-landscape variance dominates over increasing within-landscape variance. For large K, the landscapes are very similar and across-landscape variance approaches zero.

Interestingly, using relative vs. absolute performance variance offers very different qualitative interpretations. Does solving more complex problems imply more varied (e.g. Taylor and Greve, 2006) or less varied outcomes? Our brief analysis suggests that it depends on the frame of reference – whether one compares the performance of the focal problem to all problems or only problems of the same type.

[Insert Figure 4 about here]

In summary, the relative performance allows parsimonious comparisons between agents who take their environment as given and allows judging their performance relative to the best case scenario. The absolute performance in turn allows comparisons of agents that can influence their environment or compete across different environments.

***Simulation Mean as the Outcome Variable***

Very closely related is the question of whether focusing on the mean payoff of many agents across simulation runs is an informative and empirically relevant performance measure. In some settings, one may be interested in the performance of extreme (best) performers or perhaps empirically one can observe only the performance of some truncated distribution (e.g. firms that enter the industry and start producing). Consequently, in some settings, looking at the performance of extreme performers may be more informative and important than looking at the mean.

If one is indeed interested in the extreme performers, the notion of variance explicitly enters the calculation of performance. Following the discussion in March (1991), population variance is an important determinant of relative performance within a population of firms when the number of firms increases. March (1991) shows that when comparing two distributions (two firms competing together and the performance of each is drawn from independent distributions), the variance is irrelevant for the relative ranking of their performance. However, as the number of firms increases, the probability of being the best performer (or any rank above the median) increases with variance. If the lower mean is sufficiently offset by the increased variance, the likelihood of being the best performer increases. This effect increases with the number of firms in the population. Within the context of the NK model, this means that if model attributes affect variance across agents comparing means may yield different predictions than comparing the performance of the top performers. If one models a problem that invites such kind of ambiguity (i.e. whether some extreme type of measure is more informative than the mean) the ideal solution may be to obtain predictions which are robust to whether one looks at the mean or at the best performers. In other cases, the context of the model has to suggest which prediction is more relevant or if both should be considered.

For instance, Ganco and Agarwal (2008) model the effects of learning and turbulence on performance of two types of firms – diversifying entrants and entrepreneurial startups. They find that, though looking at the best versus the mean performers yields different quantitative predictions (period when one group outperforms the other), the qualitative effect of learning and turbulence remain the same whether one looks at the best or the mean performers.

To elaborate more on the logic of when the model specification may be robust to different types of the outcome variable, March (1991) notes that when the mechanics of the model affects the agent variance and one is interested in the population dynamics, it is important to incorporate the variance explicitly (e.g. if learning affects variance in the model of Ganco and Agarwal, 2008 it would be necessary to incorporate the variance explicitly). If the variance is determined by the factors which are exogenous from the perspective of the modeling analysis the qualitative patterns are likely to be robust to whether one looks at the mean performers or agents at the tails of the performance distribution.

***Temporary vs. Convergent Performance***

Another important issue related to measuring performance in the NK models is *when* to measure the performance. The consensus in the Strategy literature appears to be to focus on convergent performance (Rivkin, 2000, Rivkin and Siggelkow, 2003) though papers that show graphs of the simulation runs implicitly report both – short-term and long-term performance (e.g. Ethiraj and Levinthal, 2004, Gavetti, 2005; Gavetti and Levinthal, 2000).

The implied definition of the convergent performance (since one cannot let the model run for infinite number of periods with $t \rightarrow \infty$) within the context of the NK models is based on the behavior of agents. Convergent performance is such performance which is stable over time – agents found their local peaks and do not make additional moves. How the behavioral stability is technically measured is again somewhat ambiguous. Ideally, one would like to compare whether each agent found a local peak which in practice requires the knowledge of all local peaks in the given space – which requires additional calculations and may be computationally infeasible. More practical solution is to follow the agents for a certain number of periods – if they do not move over the last $l$ periods one can infer that the performance has stabilized and converged. But shorter $l$ will lead to shorter convergent periods and the tendency to stabilize will also crucially depend on $N$ and $K$ (affected by the variance).

There is another less technical and more theoretical problem with measuring convergent performance. Consider a situation as depicted in Figure 5 (which is a frequent outcome in NK models). Despite the agent 1 achieves superior performance when convergent performance is considered, in agent 2 performs better the short-term. Measuring convergent performance implicitly assumes that agents operate in a vacuum and face no real time constraints. When we add selection to the above model and assume that selection pressures are strong at the beginning of the simulation run (e.g. around the shock events, growth stage of an industry, etc.) the eventually inferior agent will be more likely to survive. In such environment, the convergent performance is not informative of the survival chances of the two agents. The differences between short- and long-term payoffs may also have population level implications. Any given agent 2 may be more likely to survive but in the long-term, it is more likely that the best performing agent will be of the type 1. Those agents 1 that happen to survive the initial selection pressures will outperform agents of the type 2.

Consequently, having purely static analysis based on convergent performance is incomplete and a good research design should address how the interplay between the short- and long-term performances possibly affects predictions of the model.

[Insert Figure 5 about here]

***Establishing Significance***

The notion of statistical significance plays a central role in establishing validity of empirical results. Within the context of a simulation study, the notion of statistical significance similarly plays an important role that is directly tied to model design. To establish validity of simulation predictions, the researcher needs to establish that two payoffs that are proposed as different conditional on given parameter values are indeed different in a statistical sense. It appears that all is needed to achieve such goal is to increase the number of simulation runs. However, the researcher's ability to increase the number of runs sufficiently rests on a number of factors including the size of the NK space, the efficiency of the code and the model design. Meeting these demands is far from trivial in many NK models. If a researcher makes wrong choices in the simulation implementation it may take excessive amounts of CPU time before it converges to state where meaningful statistically significant differences can be observed even with powerful hardware platforms.

Reflecting this problem, if a graphical output of the NK simulation shows jagged mean performance curves the confidence bands around these curves are likely to be thick which can make qualitative judgments or comparisons across multiple curves with statistical significance problematic.[8] The best practice thus should be to display error bars or

[8] A rule of the thumb in an NK model simulation is that a jagged mean performance curve suggests large standard error around the mean. However, a smooth curve is a necessary but not a sufficient condition for achieving small standard error. Two even very jagged mean performance curves with large standard errors may be statistically distinct – if the difference is sufficiently large.

significance bands or discuss at which level has been the significance achieved - as it is a common practice in hard sciences.

It is also worth noting, that earlier researchers utilizing the NK modeling had to make significant compromises due to the computational constraints. However, as available computational power increases such compromises should be avoided. For instance, one way that has been used to cut computational costs is to average over *n* agents but over only some fraction of *n* NK landscapes when calculating mean performance. If one uses the implementation algorithms a) or b) mentioned above, the NK space has to be ideally generated in each simulation run and this is computationally very costly. To minimize these costs a simple solution appears to be to generate a single landscape for multiple otherwise independent agents (in a setting where there is no selection or other interaction between the agents justifying joint NK landscape). For instance, instead of generating 10,000 agents where each has a uniquely drawn NK space one may generate 10,000 agents with each 100 agents sharing the same space.

To illustrate the problem with such specification, we calculate and show 10 performance means where each mean is calculated as the average of the performance of 10,000 agents. We use the model with $K = 3$, $N = 10$ and interdependencies randomly distributed within the interaction matrix. In Figure 6, we use the first method where for each mean each of the 10,000 agents draws its independent landscape. In Figure 7, we use the second method where every 100 agents share the same landscape. As we see, the means in Figure 6 are much closer together (and appear as a thick single line) than on Figure 7 implying greater variance of the mean estimate using the second method.[9]

The computational gains are offset by losses in the variance of the mean estimates and what is more troubling it appears that these losses cannot be analytically quantified.[10] Then the only way to obtain correct confidence intervals would be to numerically generate

[9] The effect is very similar when one measures the relative performance.
[10] In other words, the standard formula for calculating the variance of the mean $Var(\pi)=\sigma^2/n$ does not apply.

them through a similar exercise like we just did which, however, defeats the entire purpose of computational savings.

[Insert Figures 6 and 7 about here]

## Model Extensions Not Currently Exploited within the Strategy Literature

Many useful and important extensions of the NK model have been developed in the recent Strategy literature. These include, for instance, analysis of the interaction patterns that give rise to NK spaces with certain characteristics like modularity (Ethiraj and Levitnhal, 2004; Ethiraj *et al.*, 2008) or various hierarchical patterns (Siggelkow and Rivkin, 2007), modeling cognitive search processes through cognition representations (Gavetti and Levinthal, 2000; Gavetti, 2005; Sommer and Loch, 2004) or vicarious learning (Rivkin, 2000; 2005). Despite these advances, there are still several very interesting and important extensions developed in other fields that may prove useful in both theoretical and empirical settings within Strategy.

### *Modeling the Varying Effect of Linked Decision Elements*

The original specification of the NK model has to be seen as relatively abstract and crude representation of interdependent decision-making process. However, in situations that require more realistic approximation of decision-making – presumably when modeling micro-level agent behavior – more precise specifications may prove to be valuable. One of the simplifying assumptions in the original specification arises from the nature of effect that interactions have on the performance contributions of the focal decision. As we mentioned above, for instance, when there exists a coupling between the decision elements $x_1$ and $x_3$ then a change in decision $x_3$ will cause the payoff contribution of the element $x_1$ to be redrawn from the underlying distribution. When the focal decision element $x_1$ is coupled to many ($K$) other elements, its payoff will be redrawn whenever any of the coupled decision elements change. This specification then implies that change in even a single linked decision element has very dramatic effect on the performance contribution of the focal decision element – i.e. it will completely change it. Furthermore, the payoff contribution of the focal decision element will be redrawn whenever any of the coupled decisions change – the magnitude of the payoff change is independent of the number of coupled decisions that change.

As a way of providing a more realistic specification, Milhiser and Solow (2005) have extended the original framework by developing a bounded NK model (Milhiser and Solow, 2005). In the bounded NK model, the performance contribution of the focal decision element is not redrawn from the entire underlying distribution but from a bounded interval with the bounds depending on the extent of change in the coupled decision elements.

Using our previous notation, the performance contribution $\pi_i$ of the decision element $i$ is: $\pi_i = \pi_i(x_i; x^1_{j(i)} \ldots x^K_{j(i)}), i \notin j(i)$. In the original model, whenever any of the coupled $K$ decision elements change, the value of $\pi_i$ is redrawn. Milhiser and Solow (2005) instead assume that the value will be redrawn from the interval with lower and upper bounds (using underlying uniform [0,1] distribution):

$$l_i = \pi_i - \frac{\delta_i}{K}\pi_i\,,\ u_i = \pi_i + \frac{\delta_i}{K}(1-\pi_i)$$

where $\pi_i$ is the performance contribution prior to the change and $\delta_i$ denotes the number of coupled decisions that have been altered. $\delta_i$ does not include the change in the focal $i^{th}$ decision element. If element $x_i$ changes the value is always redrawn from the entire interval.

The logic of this specification is straightforward – if none of the coupled decision elements change ($\delta_i = 0$) the new performance contribution will be identical to the existing one – the value will be drawn from the interval $[\pi_i, \pi_i]$. If all $K$ decisions change ($\delta_i = K$), the value for the new performance contribution will be drawn from the entire interval [0, 1]. The size of the interval in between these two extreme values will be proportional to the number of decision elements that changed.

Nevertheless, adding more realism into the NK model assumptions not only reduces its parsimony but also creates additional challenges that researchers need to be aware of.

First, the algorithms #1 and #3 cannot be used for the bounded NK model. The bounded NK space is not exogenous (compared to the original NK model) to the moves of the individual agent in a sense that path dependencies in the model determine the shape of the space. Consequently, it needs to be progressively generated during the simulation run and only a modified algorithm #2 is feasible. Second, the specification of the bounded NK model implies that two identical decisions may be generated as having two different payoffs depending on how agent arrived at them. The NK space is "subjective" to the agent and in most settings two agents cannot search the same space. Space searched by an agent making more distant jumps will be more rugged than an initially same space searched by an agent performing incremental moves.

The above specification can be thus problematic in some settings. However, it may also create opportunities. Milhiser and Sollow (2005) use this specification for the study of team decision-making dynamics but we can imagine that this more fine-grained specification of the interactions may provide interesting insights when built into the models of cognition (Gavetti and Levinthal, 2000) or imitation (Rivkin, 2000). We may also speculate that such specification may be useful when one is interested in modeling the micro dynamics of turbulence (e.g. Ganco and Agarwal, 2008). In the turbulent environment, the extent of change will be self-reinforcing but at different rates depending on the amount of existing turbulence which may lead to phase transitions and non-trivial dynamics.

***Modeling the Varying Importance of Decision Elements***

The original NK model assumes not only that the interdependencies are randomly distributed within the interaction matrix but also that all decision elements have the equal importance in determining the overall system payoff. The Strategy literature has extensively exploited relaxing the first assumption by modeling modularity (Ethiraj and Levinthal, 2004; Ethiraj *et al.*, 2008) and patterned interactions in general (Siggelkow and Rivkin, 2007). However, relaxing the second assumption may also prove useful especially in settings where interaction patterns complement unequal distribution of power or decision-making importance. For instance, such setting may include the decision-making dynamics within the top management teams or dynamics within design structures where component interactions do not only follow specific patterns but also have unequal importance for the overall payoffs.

As a way of relaxing the second assumption, Solow (1999) developed a simple modification of the NK model by adding weights to the mapping [1]. The modified NKW model has the form:

$$\pi(x,N,K,w)=\sum_{i=1}^{N} w_i\pi_i(x_i;x^1_{j(i)}\ldots x^K_{j(i)}),\, i\notin j(i), \sum_{i=1}^{N} w_i=1.$$

If $w_i = 1/N$ for all $i$ the model reduces to the standard NK model.

This apparently simple change in the specification has some interesting implications. For instance, Solow (1999) shows that when weights are distributed unequally – e.g. there is one decision element with disproportionately more weight than all others – the lock-in problems are attenuated. At the extreme ($w_i > 0.9$, for some $i$), the lock-in problems disappear completely and performance increases in $K$. This interesting finding suggests that by redesigning the systems to reallocate more importance to one of the decision elements, the designer may create a smoother space that can be optimized more easily. This may generate novel insights in the study of innovations and technology management (Fleming and Sorenson, 2001; Frenken, 2000) and may be particularly interesting when combined with the notions of modularity (Ethiraj and Levinthal, 2004) and patterned interactions (Siggelkow and Rivkin, 2007). For instance, Siggelkow and Rivkin (2007) discuss how hierarchical design within the interaction matrix helps to mitigate the lock-in problems by allowing to concentrate on the core components first (defined as those that have many linkages to other components). It is possible that such mechanism is reinforced – eliminating the lock-in problems - when the core component has also the most weight or alternatively breaks down if some of the peripheral component has the most weight.

***Modeling Different Functional Forms of the Payoff Function***

The original specification of the NK model (Kauffman, 1993, 1995) as well as all of its uses within the Strategy field specify the payoff mapping using the average of the individual payoff components. Solow *et al.* (2000) show that the choice of the (weighted) average is non-trivial and discuss implications of other functional forms like the min(), max() and geometric mean. They find that, for instance, using max() in the mapping instead of the mean() leads to increasing payoff with increasing *K*. Such finding is intuitive since increasing degree of interdependence creates more opportunities (increase in global maximum, see Figure 3) which improves the payoff of best performers.

Such modification in the specification can have implications on the type of questions that the model can answer. For instance, using the max() instead of the mean() may be more appropriate in the "winner-takes-all" settings where agents maximize a portfolio of projects and only the best projects count (e.g. innovation races).

***More Speculative Extensions: Modeling Endogeneity***

The original specification of NK model implies that agents solve problems that are exogenous and the underlying characteristic given by the parameters *N*, *K* as well as the positions of interactions over the simulation run remain fixed. This specification makes the model not only computationally tractable but also parsimonious and transparent. However, in some settings (e.g. modeling industry evolution, technological innovations) it may be desirable to model how attributes of the search space change over time – e.g. technological innovations may be changing in both interdependence and scope as industry evolves, firms may be getting more internally complex as they grow. How to model these processes within the framework of an NK model remains a significant challenge.

The problem, however, is not in the computational feasibility but rather in the theoretical specification. The main issue is how one incorporates the changes in *N* and *K* into the existing payoff structure of the model. The simplest solution (Altenberg, 1994) is to assume that after "rewiring" (rewiring is defined as changes in the location or density of linkages – the interdependencies are rewired if the locations of ones in the interaction matrix change or if new ones appear or the existing ones disappear) - or after changing the number of decision elements (N) - the algorithm simply *redraws the payoff contribution of all the decision elements that are linked to the focal decision element* (the new random interdependencies are generated if a new element is added). For instance, let's assume that the interactions among the decision elements in an NK model have the form:

$$L = \begin{array}{c|cccc} & 1 & 2 & 3 & 4 \\ \hline 1 & 1 & & 1 & 1 \\ 2 & 1 & 1 & & 1 \\ 3 & 1 & & 1 & 1 \\ 4 & & 1 & 1 & 1 \end{array}$$

We assume that the rows are affected by columns – i.e. the payoff contribution of the first component (row 1) is affected by the choice of itself (column 1), the third (column 3) and the fourth (column 4) decision element. Now, assume that the system is rewired and the interaction between the first and third elements (row 1, column 3) disappears. The new system has the form:

$$L'_{\mathrm{Re}wired} = \begin{array}{c|cccc} & 1 & 2 & 3 & 4 \\ \hline 1 & 1 & & & 1 \\ 2 & 1 & 1 & & 1 \\ 3 & 1 & & 1 & 1 \\ 4 & & 1 & 1 & 1 \end{array}$$

The Altenberg's (1994) specification now implies that the performance contribution of the first element $\pi_1$ will be redrawn from the underlying distribution. Similarly, if one adds an element (fifth element is added), the new system can look like the following:

$$L'_{AddN} = \begin{array}{c|ccccc} & 1 & 2 & 3 & 4 & 5 \\ \hline 1 & 1 & & 1 & 1 & 1 \\ 2 & 1 & 1 & & 1 & \\ 3 & 1 & & 1 & 1 & 1 \\ 4 & & 1 & 1 & 1 & \\ 5 & 1 & 1 & & & 1 \end{array}$$

The fifth column and row were generated. Since we randomly assigned that the new element interacts with decisions 1 and 3, the payoff contributions of decisions 1 and 3 has been redrawn.

The algorithm #3 can be, in fact, used for such implementation of endogeneity without any modifications. Altenberg (1994) used his algorithm with endogenous *N* and *K* to model selective genome growth (which would correspond to the endogenous growth in the size and interdependence of the decision vector). He assumed that when new elements are added and new up to *K* linkages are generated the overall system performance is calculated. Subsequently, only those new elements are kept that increase overall payoff – i.e. the agent is assumed to perform local (greedy) search over genome (decision vector) expansions. One of the interesting findings is that as the genome (decision vector) grows it is optimal to add elements with decreasing number of linkages – the NK space is becoming smoother and the additional elements more modular as the vector grows. The intuition behind this result is that the probability that the performance after adding a new element exceeds the existing performance decreases with the number of linkages of the new element. In a more coupled system, the performance improvements brought by changing one element are hampered by negative effects in coupled decisions (which are redrawn) and consequently, the performance will regress to the mean of the underlying distribution. As simulation progresses only increasingly less coupled elements have the chance of increasing the payoff further – at the extreme uncoupled element can draw from an entire distribution without affecting any other payoff and has the greatest likelihood of contributing to the overall payoff.

Interestingly, Altenberg's (1994) selective genome growth model did find some application within the Strategy and Technology Management literature domains. For instance, Frenken (2000) has applied the model to the study of modularity and dominant design. He maintains that the selective growth explains why complex and highly coupled components of technological innovations are settled first and then innovation typically proceeds with more modular and peripheral components surrounding the core innovation.

Despite some applicability of this specification, the main problem likely preventing its broader usefulness is related to the random redrawing after the changes in *N* and *K*. More specifically, the model implies that it is optimal to add increasingly less coupled elements to the system. If given the choice the agent would myopically choose only new components that are uncoupled to the old system (since coupling causes redrawing of some of the existing payoffs, which is on average detrimental). This process will push the overall coupling of the system below the optimal level (see Figure 3). The myopic process as defined in Altenberg (1994) is not sufficient to ensure convergence to the optimum (Kauffman, 1995). The reason is that the optimality of the potential incorporation of changes in *N* and *K* is completely driven by the statistical properties of the NK space (and basically by the difference in the payoff distributions before and after the change). In other words, there is no feedback to payoff from the agent search behavior that would interact with the potential changes in *N* and *K*. The type of agent's search and its search behavior are irrelevant for making the decision about changes in *N* and *K*. This is in stark contrast to the 'philosophy' of the original NK model specification where the fundamental property of its dynamics is the interplay between the attributes of the NK space and attributes of the agent search behavior.

The consequence of this specification problem is that the endogenous NK model is a poor fit for modeling conscientious decisions of human actors. Regardless of the circumstances (e.g. degree of turbulence in the environment - Ganco and Agarwal, 2008) and expected outcomes of the search behavior conditional on the given type of NK space, the coupling will converge to zero. On the other hand, in real-life settings, for instance, inventors design the system structure and its coupling in anticipation of its performance as exploited by local search. The endogenous NK model lacks a feedback mechanism that would incorporate expectations about the results of subsequent search behavior into making the decision about changes in *N* and *K*.

***Modeling Larger Systems***

One of the valid criticisms of the original NK model specification is the static nature of its environment. The NK landscape stands still while the agent searches for the optimum. Despite such specification may be perfectly valid when, for instance, looking at the intra-firm dynamics (Rivkin and Siggelkow, 2003) or effects of interdependence patterns on agent behavior and performance (Ethiraj and Levinthal, 2004; Rivkin and Siggelkow, 2007) or, in general, at micro-level agent optimization process (Kollman *et al.*, 2000). However, once one is interested in studying more dynamic settings where agents face not only uncertain but also dynamic and uncertain environments (Page, 2007) modifications to the NK model are necessary. In fact, in dynamic settings, the simulation methodology like the NK model may be superior to empirical one as a way to qualitatively disentangle the population, agent-level and environmental effects.

There have been a number of studies including several in Strategy addressing this problem through developing dynamic NK environment. In general, there are two categories of solutions – the first assumes exogenous shocks in the environment (Levinthal, 1997, Siggelkow and Rivkin, 2005) and the second creates changes by modeling entire system by embedding the NK landscapes into more broader dynamics (Ganco and Agarwal, 2008; Lenox *et al.*, 2006, 2007).

The first category of models is methodologically relatively straightforward. At exogenously specified periods the NK space is distorted by either completely redrawing payoff contributions of some of the decision elements or the new 'after-shock' payoff is computed as a weighted average between the existing payoff and a random draw (Gavetti, 2005). The weights can then be used as one of the parameters controlling how correlated are the new and old NK landscapes. The endogenous NK model using the algorithm #3 can also serve as a powerful shock generator since rewiring of the interaction matrix will cause some payoff contributions to be redrawn.

The models in the second category attempt to model the entire dynamic system. The specification that is closest to the original NK model is the NKC model also developed by

Kauffman (1995). The NKC model is a relatively straightforward extension of the NK model. The additional parameter C specifies the extent to which individual "sub-landscapes" are tied together, i.e. the extent of coupling across the NK spaces. The important feature of the model is that by creating the couplings across landscapes the individual landscapes become dynamic. The shape of the focal NK landscape deforms (payoffs are redrawn) as agents change coupled decisions on other landscapes. In the NKC model specification, the payoff of each decision element within the decision vector of each agent $x_i$, is affected by its own value, *K* other decisions within the agent's decision vector and *C* decisions of each (or some) of the other agents in the environment. The important implication of the NKC model is that the shape of the NK space is no longer exogenous to the moves of agents. However, the shape of an NK space is still exogenous to the agent searching but it is endogenous to the moves of agents on coupled spaces. For instance, an agent moving on space A will deform the space B which will cause the agent searching the space to make a move which in turn deforms the space A and cause the agent searching space A to move again. These feedback loops fuel the dynamism and turbulence of the environment. The shocks to the NK spaces are endogenous arising from the agent moves. Although in the original NKC specification multiple agents generally cannot search the same space, one could design a setting where this is feasible (the effects through coupling would arise not from individual but collective behavior). The NKC models have been recently applied to the study of industry evolution (Ganco and Agarwal, 2008) and corporate strategy (Caldart and Ricart, 2008). However, an additional challenge associated with this specification is the necessity to establish empirical validity for the additional parameter *C* (inter-agent interdependencies). For instance, in the model by Ganco and Agarwal (2008) the agents searching individual landscapes represent firms. They maintain that the *C* parameter that determines the amount of linkages across the NK spaces faced by individual firms is given by structural characteristics of the given industry like the technological intensity or economies of scale.

Similarly, Lenox *et al.* (2006, 2007) implement an alternative general framework around the NK model. They link the individual agents searching the same NK landscape (the specification could be easily extended to multiple NK landscapes) together by assuming that the converged payoffs from the NK model feed as cost shifting parameters into a game that takes place among the agents. The game (Cournot) captures the notion of market interaction between the individual agents solving the NK space. The NK space here represents the search for product design over a technology space whereas the *Cournot* game captures the notion of market interaction. The model by Lenox *et al.* (2006, 2007) thus combines the analytical neoclassical model with the NK model. However, the applicability of this specification hinges on the assumption that the impact of the interaction dynamics does not change the shape of the production function but only shifts it in terms of costs. Consequently, it appears that such specification is more applicable in contexts where the focus is on process rather than product innovations – e.g. manufacturing-oriented industries. In support of this conjecture, Lenox *et al.* (2008) find support for their model when looking at a wide cross-section of industries.

A third alternative specification was developed by Chang and Harrington (2000). They abandon the original specification of the NK model but utilize its idea of a search over rugged landscape. They build a complete model in the tradition of neoclassical economics using the primitives like consumers with utility functions and firms with revenue and costs functions. Their model is thus internally logically consistent and avoids the issues of combining the NK space with the neoclassical primitives - they instead build the 'NK-like' space using the neoclassical primitives. More specifically, Chang and Harrington (2000) assume that the market is populated by consumers with complementary preferences. Each consumer is assumed to have a set of preferred product attributes and his/her utility decreases with the distance of a purchased product attributes from the ideal set. The market is populated by a mass of agents that have their preferences distributed according to a given distribution function. The firms cannot customize the products for each consumer and thus have to optimize over the product attribute space to maximize their profit. However, the firm profit function emerges as 'rugged' with local optima. Due to strong complementarities in consumer preferences for different product attributes, firms would have to ideally optimize by concurrently changing multiple product attributes. However, Chang and Harrington (2000) assume that the firms are myopic and search the space only locally by altering one product attribute at a time. This set of assumptions gives rise to lock-ins and the dynamics familiar from the traditional NK models.

The strength of the model by Chang and Harrington (2000) is in its internal consistency but abandoning the NK structure has its costs. For instance, the notion of interdependence and complexity which has very concrete form in the original NK model specification becomes only implicit. The interdependence cannot be easily tuned or measured since it arises from multiple factors like preference parameters for different attributes, distributional assumptions, strength of complementarities, etc. Despite of that, the model is important contribution by showing how one can design an NK-like model in a dynamic setting without necessarily relying on the original specification.

## Conclusion

The purpose of our paper was to discuss recent methodological advances in the field of Strategy related to the use of the NK model. In general, we found that researchers within our field use the NK models very creatively – they not only apply it to questions within Strategy but rather develop important modifications and extensions that may prove useful beyond the Strategic Management field. We have also reviewed a number of additional extensions not currently exploited that may bear fruit when applied to the right questions and possibly further extended.

We are also excited about the NK model as a powerful tool that has already created lasting contribution in the field and beyond. We are convinced that the importance of the NK model and similar agent-based simulation tools will only increase in the future - especially when applied to relevant questions and when modelers will be aware of its limitations. The interdisciplinary nature of simulation research also poses unique demands. We tried to facilitate interdisciplinary exchange of information and provide ”boundaryspanning" by our deep and narrow focus on the NK model methodology. We hope our paper provided a meaningful contribution to this discussion.

**References**


Altenberg, L. 1994. Evolving better representations through selective genome growth, in: J.D. Schaffer, H.P. Schwefel, H. Kitano (eds.), *Proceedings of the IEEE World Congress on Computational Intelligence* (Piscataway NJ: IEEE), pp. 182-187.

Altenberg, L. 1996. NK Fitness Landscapes. In *Handbook of Evolutionary Computation*, ed. T. Back, D. Fogel, and Z. Michalewicz. Institute of Physics Press.

Caldart, A.A. and Ricart, J.E. 2008. Corporate Strategy: An agent-based approach. *European Management Review*, forthcoming

Davis, J. P., Bingham, C. B., & Eisenhardt, K.M. 2007. Developing theory through simulation methods. *Academy of Management Review*, forthcoming.

Davis, J.P., Eisenhardt, K.M., & Bingham, C.B. 2008. Complexity theory, market dynamism and the strategy of simple rules. Working paper, MIT, Boston, MA.

Ethiraj, S.K., and Levinthal, D. 2004. Modularity and innovation in complex systems. *Management Science*, 50(2): 159-173.

Ethiraj, S. K., D. A. Levinthal, and R. R. Roy. 2008. The dual role of modularity: Innovation and imitation, *Management Science*.

Fleming, L., & Sorenson, O. 2001. Technology as a complex adaptive system: evidence from patent data. *Research Policy*, 30: 1019-1039.

Fleming, L., & Sorenson, O. 2004. Science as a map in technological search. *Strategic Management Journal*, 25: 909-928.

Frenken, K., 2001a. *Understanding Product Innovation using Complex Systems Theory. Doctoral dissertation*. University of Amsterdam and University of Grenoble.

Frenken, K., 2001b. *Modelling the organisation of innovative activity using the NK-model*. Working paper.

Frenken, K., 2006. Technological innovation and complexity theory. *Economics of Innovation and New Technology*, 15(2): 137-155.

Ganco, M. 2008. Inventive Process as a Recombinant Search over Complex Landscape: Evidence from the Disk Drives Industry. Working paper. University of Illinois.

Ganco, M. and Agarwal, R. 2008. Performance Differentials between Diversifying Entrants and Entrepreneurial Start-ups: A Complexity Approach. Forthcoming in the *Academy of Management Review*.

Gavetti, G. 2005. Cognition and hierarchy: rethinking the microfoundations of capabilities' development. *Organization Science*, 16(6): 599-617.

Gavetti, G., & Levinthal, D., 2000. Looking forward and backward: cognitive and experiential search. *Administration Science Quarterly,* 45, 113–137.

Harrington, J.E. and M.H. Chung. 2000. Centralization vs. Decentralization in a Multi-Unit Organization: A Computational Model of a Retail Chain as a Multi-Agent Adaptive System, *Management Science*, Vol. 46, November 2000, 1427-1440.

Harrington, J.E. and M.H. Chung. 2006. Agent-Based Models of Organizations, Handbook of Computational Economics, II: Agent-Based Computational Economics, L. Tesfatsion and K. Judd, eds., Elsevier, 2006.

John H. Holland. 1992. Genetic Algorithms. *Scientific American*. July:66--72.

Kauffman, S. A. 1993. The origins of order: self-organization and selection in evolution. Oxford University Press, Oxford.

Kauffman, S.A. 1995. At home in the universe. Oxford University Press.

Kollman, K., J.H. Miller and S. Page, 2000. Decentralization and the search for policy solutions. *Journal of Law, Economics and Organization*. 46(1): 102-128.
Koza, J. R. 1992. Genetic Programming, MIT Press.
Lenox, M., S. Rockart, and A. Lewin. 2006. Interdependency, Competition, and the Distribution of Firm and Industry Profits. *Management Science*. 52(5): 757-772.
Lenox, M., S. Rockart, and A. Lewin. 2007. "Interdependency, Competition, and Industry Dynamics.*Management Science*. 53(4): 599-615.
Lenox, M., S. Rockart, and A. Lewin. 2007. Does Interdependency Affect Industry Profitability?" Working Paper.
Levinthal, D.A. 1997. Adaptation on rugged landscapes. *Management Science*, 43: 934-950.
March, J.G. 1991. Exploration and exploitation in organizational learning. *Organization Science*, 2(1): 71-87.
Millhiser, W. and Solow, D. 2005. On the Impact of Worker Interactions on How to Divide a Team, Working paper, Case Western Reserve University.
Mitchell, M. 1996. An Introduction to Genetic Algorithms. MIT Press, 1996.
Page, S.E. 2007. Uncertainty, Difficulty, and Complexity. Journal of Theoretical Politics.
Rivkin, J.W. 2000. Imitation of complex strategies. *Management Science*, 46: 824-844.
Rivkin, J.W., 2001. Reproducing knowledge: replication without imitation at moderate complexity. *Organization Science*, 12, 274-293.
Rivkin, J.W., and Siggelkow, N. 2003. Balancing search and stability: interdependencies among elements of organizational design. *Management Science*, 49(3): 290-311.
Rivkin, J.W., and Siggelkow, N. 2007. Patterned interactions in complex systems: implications for exploration. *Management Science*, 53(7): 1068-1085.
Solow, D., Burnetas, A., Tsai, M. and Greenspan, N. S., 1999. Understanding and Attenuating the Complexity Catastrophe in Kauffman's NK Model of Genome Evolution, Complexity, 5(1), 53-66.
Solow, D., Burnetas, A. Tsai, M. and Greenspan, N. 2000. On the Expected Performance of Systems with Complex Interactions Among Components, *Complex Systems*, 12(4).
Solow, D., Vairaktarakis, G. Piderit, S.K. and Tsai, M. 2002. Managerial Insights into the Effects of Interactions on Replacing Members of a Team. *Management Science*, 48(8), 1060-1073
Solow, D. Piderit, S., Burnetas, A. and Leenawong, Ch., 2005. Mathematical Models for Studying the Value of Motivational Leadership in Teams. *Computational and Mathematical Organization Theory*, 11: 5-36.
Sommer, S.C. and C.H. Loch, 2004. Selectionism and Learning in Projects with Complexity and Unforeseeable Uncertainty, *Management Science*, 50(10): 1334–1347.
Sorenson, O. 2002. Interorganizational Complexity and Computation. In J.A.C. Baum (ed.) "Companion to Organizations". Oxford: Blackwell Publishers: 664-685.
Sorenson, O., Rivkin, J., & Fleming, L., 2006. Complexity, networks and knowledge flow. *Research Policy*, 35, 994-1017.
Taylor, A., & Greve, H.R., 2006. Superman or the fantastic four? Knowledge combination and experience in innovative teams. *Academy of Management Journal*, 49, 723-740.

Figure 1

NK interaction matrix
$N = 5$, $K = 2$, randomly distributed links

$$\begin{pmatrix} 1 & & 1 & & 1 \\ & 1 & 1 & 1 & \\ 1 & 1 & 1 & & \\ & 1 & & 1 & \\ 1 & & & 1 & 1 \end{pmatrix}$$

NK interaction matrix
$N = 5$, $K = 2$, neighbors linked, with wrapping

$$\begin{pmatrix} 1 & 1 & & & 1 \\ 1 & 1 & 1 & & \\ & 1 & 1 & 1 & \\ & & 1 & 1 & 1 \\ 1 & & & 1 & 1 \end{pmatrix}$$

Figure 2
Visualization of the NK landscape, $N = 6$, $K = 2$

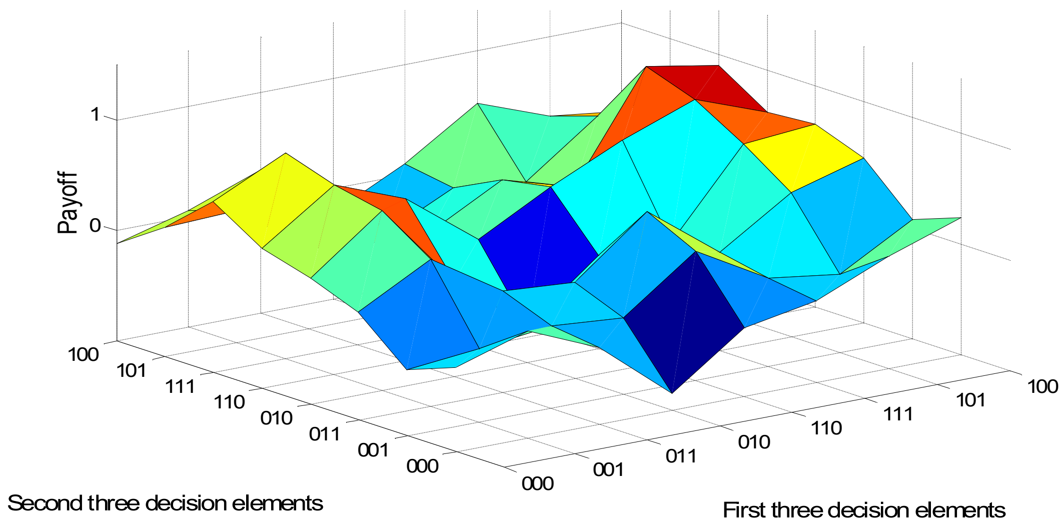

Figure 3

NK Model Performance for $N = 10$ as a function of $K$ at period 50

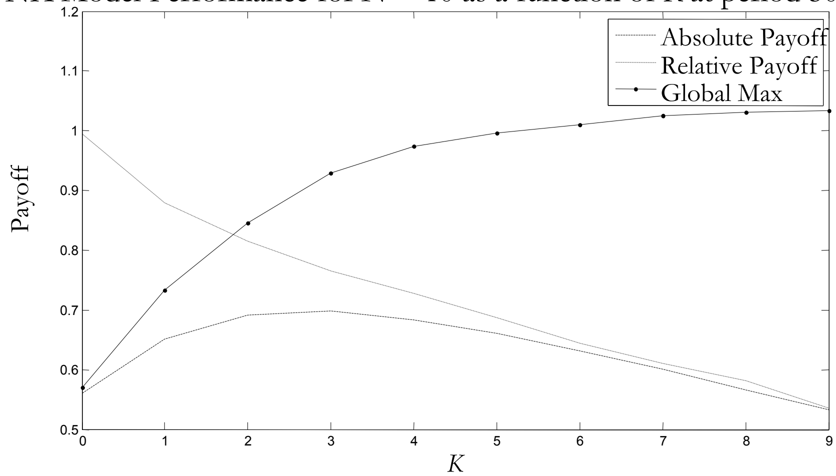


Figure 4

NK Model Standard Error of Performance for $N = 10$ as a function of $K$ at period 50

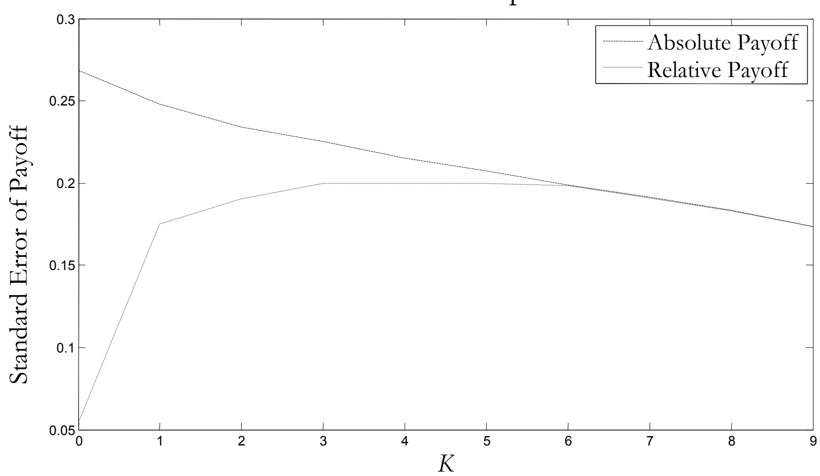

Figure 5

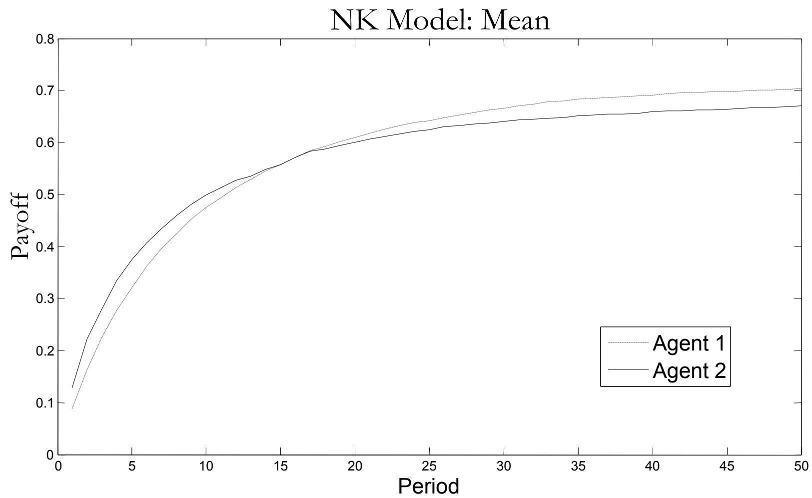
NK Model: Mean
Payoff
Period
Agent 1
Agent 2
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0
5
10
15
20
25
30
35
40
45
50

Figure 6

NK Model: Mean performance, 10 Samples, each 10,000 runs
Correct implementation method – independent landscapes

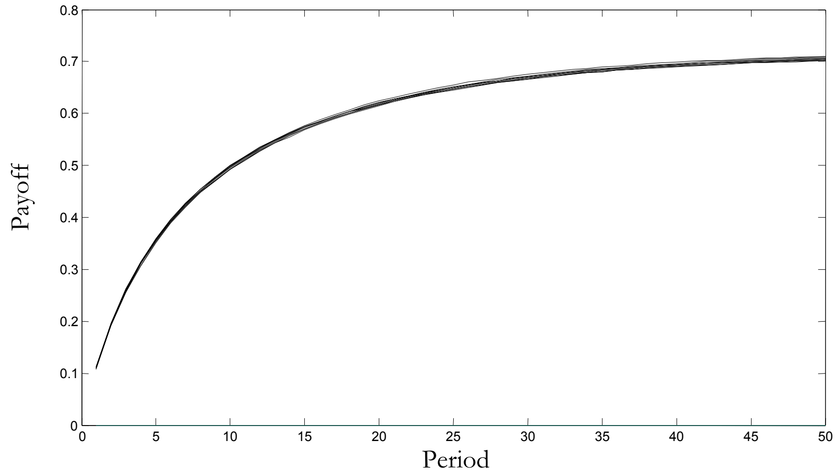


Figure 7

NK Model: Mean performance, 10 Samples, each 10,000 runs
Incorrect implementation method – multiple agents share landscape