Modelling the organisation of innovative activity using the NK-model

Koen Frenken

Department of Innovation Studies, University of Utrecht
P.O. Box 80125, NL-3508 TC, Utrecht, The Netherlands
phone: +31.30.253.7850, fax: +31.30.253.3939, e-mail: k.frenken@geog.uu.nl

Paper prepared for the Nelson-and-Winter Conference, Aalborg, 12-16 June 2001

Abstract

In this paper, I aim to provide both an introduction to and an extension of a research program in evolutionary economics based on Kauffman's (1993) NK-model of complex evolving systems. Complex systems, such as organisations and technological artefacts are characterised by conflicting constraints due to interdependencies between its constituting elements. These constraints can lead organisations or technologies to lock-in into local optima on a rugged fitness landscape. The NK-model research program can be considered further elaboration of the program of evolutionary economics as laid down by Simon (1969), Nelson and Winter (1982) and others. In particular, questions related to dynamic efficiency can be addressed by looking in what ways firms can search a fitness landscape intelligently.

The central question holds what modes of organisation can be distinguished in designing a complex system and how can their performance how search activity in complex systems is best organised. Three types of organisational modes are distinguished:

1. a centralised organisation in which innovations are assessed with reference to the system as a whole,
2. a decentralised organisation where innovations are assessed with reference to an individual agents only,
3. a network organisation of strategic alliances where innovations are assessed with reference to the network of co-designing agents.

Analytical results show that when system exceeds a critical size, networking may well prove more effective than centralised decision-making. For many parameter regimes though, simulation exercises will have to show the relative performance of alternative modes of organisation. Such a research program is outlines in the final part of the paper.

1. INTRODUCTION

Throughout the history of economic thought, evolutionary concepts have played an important role. For example, the static concept of equilibrium in perfectly competitive markets as described in neoclassical economics has been interpreted as the outcome of a dynamic selection process between competing firms (Alchian 1950). The resulting equilibrium reflects, given particular assumptions, an efficient allocation of resources with marginal costs equalling marginal benefits at the level of individual firms.

Nelson and Winter (1982) argued that marginal analysis of static efficiency only refers to markets with little innovation and contributes little to the understanding of technological innovation and economic development. One of the distinguishing facets of evolutionary models is its emphasis on technological innovation and the analysis of dynamic efficiency.

In this paper, I aim to provide both an introduction to and an extension of a research program in evolutionary economics based on Kauffman's (1993) NK-model of complex evolving systems. Complex systems, such as organisations and technological artefacts, are characterised by conflicting constraints due to interdependencies between the constituting elements. These constraints can lead organisations or technologies to lock-in into local optima on a rugged fitness landscape. The NK-model research program can be considered further elaboration of the program of evolutionary economics as laid down by Simon (1969), Nelson and Winter (1982) and others. In particular, questions related to dynamic efficiency can be addressed by looking in what ways firms can search a fitness landscape intelligently.

Apart from contributions that focus on formal properties of the NK-model (Frenken *et al.*1999), the use of the NK-model in evolutionary economics hitherto can be divided into studies addressing different systems of reference. First, there is a body of literature that aims to translate the properties of the NK-model to issues of firm strategy (Levinthal 1997; Gavetti and Levtinthal 2000; Rivkin 2000). These contributions are following up on earlier approaches in organsiation theory that concentrate on the local nature of organisational adaptation and the importance of imitation and heuristics in search (Cyert and March 1963; Simon 1969). Second, there are simulation studies that take the NK-model to represent a production technology (Auerswald *et al.* 2000; Kauffman *et al.* 2000). In this context, the complexity stems of interrelations between different tasks required to produce a particular output. Third, there is a number of works that simulate complex technological artefacts as evolving NK-systems (Valente 2000; Frenken 2001a,b). In these approaches, artefacts are represented by the choice of elements incorporated in its design. The complexity stems from interdependencies between different elements in a technological artefact (for example, between a car's engine, gearbox, brakes, springs and tires). In the review part of this paper, studies on all systems of reference will be discussed. In the modelling part of this paper, the system of reference will be a technological artefact, but the formal nature of the NK-model and its extension that follows allows one to consider the model equally as a model of strategy formation in firms or as a model of production technologies.

The central question holds what modes of organisation can be distinguished in designing a complex system and how can their performance of search activity in complex systems be compared ?

Three types of organisational modes are distinguished: a centralised organisation, a decentralised organisation and a network organisation:

(i) a centralised organisation in which innovations are assessed with reference to the system as a whole,

(ii) a decentralised organisation where innovations are assessed with reference to an individual agents only,

(iii) a network organisation of strategic alliances where innovations are assessed with reference to the network of co-designing agents.

The research concerning the performance of alternative organisational forms has also been addressed in a transaction-cost economics framework (Williamson 1991), but with reference to static efficiency. Here, the performance of alternative modes is assessed in terms of the relative performance to design a new technological artefact (*cf.* Marengo *et al.* 2000).

The paper is organised as follows. In section 2, I introduce the reader into the NK-model and its generalised version. In section 3 I outline a research program in which the NK-model is used to analyse the performance of different organisational modes with reference to innovation in a complex technological artefacts. A number of formal properties is derived for artefacts with minimum and maximum complexity. Finally, in concluding remarks are listed in section 4.

2. TWO NK-MODELS WITH REFERENCE TO TECHNOLOGICAL ARTEFACTS

2.1 Complex systems

Complex systems have been defined as systems containing elements that are interrelated within a particular structure (Simon 1969). The evolutionary properties of complex systems have been subject of research in theoretical biology (Kauffman 1989, 1993; Altenberg 1995, 1997). In biology, one distinguishes between the genotype and the phenotype of an organism. At the level of the genotype mutations occur that lead to new variants in a population. At the level of the phenotype, which is the ensemble of traits that make up an organism's fitness, natural selection operates in terms of differential rates of reproduction. Dependencies among genes imply a complex relation between an organism's genotype and an organism's phenotype.

Complexity means here that a mutation in one gene may not only change the functional contribution of the mutated gene to the entire phenotype, but can also affect the functional contributions of interrelated genes to the phenotype. A gene does not simply translate into a particular trait, but operates in conjunction with other genes by regulating other genes' state of activity. Due to these dependencies among genes, a mutation in a single gene may have both positive effects on some traits and negative effects on other traits, which jointly determine an organism's fitness. For this reason, the existent set of genes put structural constraints on the possible directions in further evolution: a certain phenotypic trait can only be improved by a particular mutation when the improvement in one trait outweighs the negative by-effects of this mutation with respect to other traits. In biology, this insight has led scholars to conclude that natural selection is not expected to lead to perfectly adapted organisms (Kauffman 1993: 3-26).

Technological artefacts have also been described as complex systems. For example, Simon (1969 [1996]: 4) defined artefacts as man-made systems which are made up of elements that collectively attain one or a number of goals.[1] The complexity in designing an artefact is caused by the

[1] Hughes' (1987: 51) concept of technological system includes, apart from technical components, organisations, scientific texts, patents, and laws. However, Hughes (1987: 55) acknowledges the usefulness of approaches that define systems solely in terms of the technical components embodied.

dependencies in the working of elements that make up an assembled technological system. Only some combinations between elements "fit" in the sense that they are *complementary*. Changing one of these elements by a new element may improve the working of this element, but the negative by-effects on other elements may imply a loss in overall functioning of the system as a whole.

Where biological organisms change through blind mutations without the intervention of a "designer", technologies change through search activity by designers. The design of technological artefacts can thus be analysed as a choice problem between alternatives. From this perspective, the dependencies between elements imply that the choice of an element cannot be made independently from the choice of other elements due to interaction effects. The set of optimal choices for the elements with regard to element-specific output variables may prove sub-optimal when the effects of dependencies between elements are taken into account. For example, a type of suspension which is found optimal according to suspension tests, and a type of engine which is found optimal in engine energy-efficiency tests, may prove to be sub-optimal when put together in a car system. The engine may generate negative effects on the working of suspension, for example, caused by high vibration. Or, *vice versa*, the suspension may generate negative effects on the working of the engine, for example, caused by high resistance. These dependencies between a system's elements imply that the performance of elements can only be fully understood when one analyses the effects on the system as a whole.[2]

To find out the working of elements at the system level poses a "problem of complexity". The number of possible combinations between different variants of elements is an exponential function of the number of elements. Therefore, the difficulty in finding a good design is of a higher magnitude than finding a good element design. Simon (1969 [1996]: 194) explains combinatorial complexity of systems using the example of a working and a defective lock:

> *"Suppose the task is to open a safe whose lock has 10 dials, each with 100 possible settings, numbered from 0 to 99. How long will it take to open the safe by a blind trial-and-error search for the correct setting? Since there are 100^{10} possible settings, we may expect to examine about half of these, on the average, before finding the correct one – that is, 50 billion billion settings."*

The strategy of evaluating all possible combinations between elements is called *global* trial-and-error. Contrary to complex systems, as Simon (1969 [1996]: 194) goes on explaining, simple systems that are characterised by independence between its elements, can be optimised by *local* trial-and-error:

> *"Suppose, however, that the safe is defective, so that a click can be heard when any one dial is turned to the correct setting. Now each dial can be adjusted independently and does not need to be touched again while the others are being set. The total number of settings that have to be tried is only 10 x 50, or 500. The task of opening the safe has been altered, by the cues the clicks provide, from a practically impossible one to a trivial one."*

In the latter case, each element can be optimised locally, *i.e.* independently of the state of other elements. The problem of finding the right combination of all ten elements can be decomposed in

[2] This does not imply that "the whole is greater than the sum of its parts" as will become clear from the exposition of the NK-model. Confer Jervis (1997: 572) stating that "(i)f we are dealing with a system, the whole is different from not greater than the sum of its parts".

ten sub-problems, which can be solved independently. Then, the combinatorial complexity vanishes and the problem becomes feasible to handle. Simon's (1969) example of the lock illustrates that optimisation through local trial-and-error works well for non-complex systems. In the case of complex systems, one can generally find the optimal solution only through global trial-and-error (*cf.* Alexander 1964 [1994]: 21).

2.2 The NK-model

The anecdotal description of complexity by Simon (1969 [1996]) can be modelled analytically by Kauffman's (1993) NK-model. This model of complex systems has originally been developed as a model of biological evolution, but its formal structure allows for applications to artefacts as a complex system (Kauffman 1989, 1993; Kauffman and Macready 1995; Frenken *et al.* 1997, 1999a; Auerswald *et al.* 2000; Kauffman *et al.* 2000; Valente 2000).[3]

Kauffman (1993) describes a system by a string of N elements ($n=1,\ldots,N$). For each element n, there exist a number of dummy values called "alleles" that refer to the possible variants of this element. The different alleles of an element are labelled by integers "0", "1", "2", "3", etc. The number of alleles of element n is described as A_n. For example, a particular artefact as a vehicle technology can be described by the following three elements and their respective alleles:

$n=1$ ($A_1 = 3$):	*an engine element with three alleles*	*gasoline (0), electric (1) or steam (2)*
$n=2$ ($A_2 = 2$):	*a suspension element with two alleles*	*spring (0) or hydraulic (1)*
$n=3$ ($A_3 = 2$):	*a brake element with two alleles*	*block (0) or disc (1)*

The description of the elements in a system and their alleles is by no means given to the designer. Rather, it is a set of elements that a designer considers to be the relevant choice variables (Alexander 1964 [1994]: 20-21). Once a classification of elements and alleles has been established, each design can be labelled by a string of alleles. For example, following the classification above, a vehicle design with a steam engine, spring suspension, and block brakes is referred to as string "200".

Each string s is described by alleles $s_1s_2...s_N$ and is part of a possibility set S, for which holds:

$$s \in S \; ; \; s = s_1 s_2 ... s_N \; ; \; s_n \in \{0, 1, ..., A_n - 1\} \qquad (2.1)^4$$

In the following, we assume for simplicity that A=2 for all elements without loss of generality. Put another way, we limit the discussion to binary strings for which holds that alleles of all elements are either 0 or 1. The N-dimensional possibility space S is called the "design space" of a system, includes all possible combinations between the alleles of elements (Bradshaw, 1992;

[3] Birchenhall (1995), Windrum and Birchenhall (1998) and Cooper (2000) developed related evolutionary models of complex technological systems based on genetic algorithms as developed by Holland (1975, 1992). The important difference between the NK-model and genetic algorithms is that the latter include apart from mutation crossover as a variety generating mechanism. This mechanism holds that designs are split in two parts, which are randomly matched with parts of other designs analogous to sexual reproduction in biological evolution. The analogy in technological evolution is that particular solutions in one design can be introduced in another design and vice versa. Here, we limit the analysis to mutation as the variety generating mechanism.

[4] Note that, since the first allele is labelled "0", the description of alleles of an element ranges from 0 to A_n-1, while the number of alleles ranges from 1 to A_n.

Dennett, 1995).[5] The combinatorial nature of a design space implies that its size increases exponentially for linear increases in N. The size of the design space S is given by:

$$S = 2^N \tag{2.2}$$

The combinatorial nature of the design space of a system requires that elements are orthogonal to one another, *i.e.* that elements can be treated as dimensions. Therefore, one element of a system cannot correspond with an allele of another element in the same system. For example, the description of alleles of the engine element as gasoline ("0"), electric ("1"), and steam ("2") implies that the type of battery used in electric engines cannot count as another element *casu quo* dimension in the description of the vehicle as a system. The choice for a type of battery only constitutes a dimension for electric vehicles, and not for vehicle technologies in general. Put another way, the choice of battery takes place at a lower level in the decision tree:

→ what kind of engine ?
  → if electric engine: what kind of battery ?

Thus, a description of a combinatorial space should therefore be such that only elements at the same level in the hierarchical decision tree are taken into account, where each element itself may again be described as a system in its own right (Hughes 1987: 55; Metcalfe 1995: 36).

Kauffman (1993) developed a model of complex system called the NK-model, which has been generalised by Altenberg (1997). I will discuss both models below both as a model of biological organisms and as a model of technological artefacts.

Complexity in systems stems from dependencies between its constituting elements. The dependencies between elements in a complex system are called "epistatic relations" (Kauffman, 1993). An epistatic relation from one element to another element implies that when an element mutates (from 0 to 1 or from 1 to 0), the mutation affects both the functioning of the one element itself and the functioning of the element that it epistatically affects. Epistatic relations thus imply that a mutation in one element can affect the functioning of many other elements. The ensemble of epistatic relations within a technological system is called a technology's architecture (Henderson and Clark, 1990).

In the NK-model Kauffman (1993) restricted his analysis of complex systems to particular types of architectures that can be expressed by one parameter K, which stands for the number of elements that affect the functioning of each element. For example, the class of systems for which holds K=1 refers to systems with an architecture in which the functionality of each element depends on the choice of allele of the element itself and on the choice of the allele of one other element. This parameter can be considered an indicator of a system's complexity with K=0 being the least complex and K=N-1 the most complex architecture. Consider as an explanatory example, a system for which holds N=3 and K=1 with an architecture as specified in Fig. 1. The architecture specifies the following epistatic relations between three elements and their fitness w_n. The fitness value of the first element changes only when the element itself or the second element is mutated. The fitness value of the second element changes only when the element itself or the first element is mutated. And, the fitness value of the third element changes only the element itself or the first element is mutated.

[5] Note that the concept of design space corresponds to the concept of morphological space used by Foray and Grübler (1990).

	n=1	n=2	n=3
w_1	**x**	x	-
w_2	x	**x**	-
w_3	x	-	**x**

FIGURE 1. Example of an architecture of N=3-system with K=1

Following Kauffman (1993), we construct a fitness landscape by drawing randomly the value of the fitness of an element w_n from the uniform distribution between 0 and 1, each time it is itself mutated and each time an element with which it is epistatically related is mutated. System fitness is derived as the mean value of the fitness of elements:

$$W(s) = \frac{1}{N} \cdot \sum_{n=1}^{N} w_n\ (s_n) \tag{2.3}$$

A simulation of a fitness landscape is given in Fig. 2.

	w_1	w_2	w_3	W
000:	0.7	0.2	0.3	0.40
001:	0.7	0.2	0.6	0.50
010:	0.9	0.3	0.3	0.50
011:	0.9	0.3	0.6	0.60
100:	0.8	0.7	0.1	0.53
101:	0.8	0.7	0.4	0.63
110:	0.3	0.5	0.1	0.30
111:	0.3	0.5	0.4	0.40

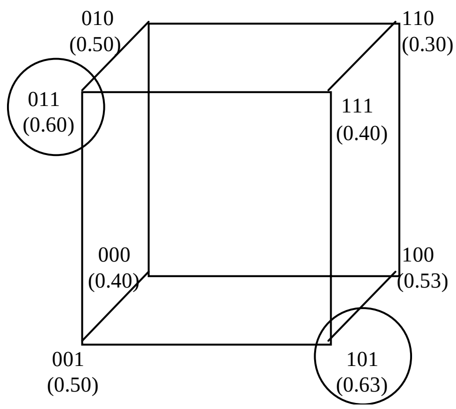

FIGURE 2. Simulation of a fitness landscape of a N=3-system with K=1

The circled strings are generally called local optima. For these strings it holds that there exist no neighbouring string with higher fitness, where neighbouring means a string that share all alleles except one. In the simulation in Fig. 2, this property holds for strings 011 and 101.

Simulation exercises done by Kauffman (1993) showed that, given *N*, the number of local optima increases for increases in K. The existence of multiple peaks characterise a "rugged landscape". Furthermore, the fitness of local optima decreases for increases in K. One can understand this outcome as the consequence of increasing *conflicting constraints* between the elements. The higher K, the more difficult it becomes to improve the fitness of one element without lowering fitness of other elements. Consequently, fitness remains low. The ultimate degree of complexity is K=N-1 which reflects the case that all elements influence all other elements. In that case, given an N-value, the number of local optima is highest and the fitness of local optima is on average lowest.[6]

For a string corresponding to a peak in the landscape it holds that all its *N* neighbouring strings have a lower fitness value. A peak in a fitness landscape implies that a search algorithm that searches by mutations in single elements gets stuck ("lock-in"). Once such an algorithm has found a peak, it can no longer escape the peak. The most well-known search algorithm that applies mutation in single elements is local trial-and-error which formally corresponds to natural selection in biological evolution.

Local trial-and-error generates a new string (trial) by randomly changing the allele of one element. The value of this element changes from 0 to 1 or vice versa. A trial thus implies that one moves along one dimension in the cube from one string to a neighbouring string and on the fitness landscape from one fitness value to a neighbouring fitness value. Trial-and-error proceeds by evaluating how system fitness W is affected by a mutation. If the trial turns out to increase W, firms continue the next mutation from this string, while a lower W induces a firm to return to the previous string, and continue the next mutation from there.[7]

In a fitness landscape, local trial-and-error thus implies that as long as there exist at least one neighbouring string that has a higher fitness, search can continue. Search will halt when a peak string is found that can no longer be improved by means of a mutation in one element. Search can thus be considered as an "adaptive walk" over a fitness landscape towards a "peak", and search will only halt when a peak is reached. Following the metaphor of the fitness landscape, search in complex technological systems can be considered a process of "hill-climbing" (Kauffman 1993).

Local trial-and-error in problem-solving on a fitness landscape is formally equivalent to mutation and selection in biological evolution (Simon 1969 [1996]: 45-47). To understand this analogy consider again Fig. 4. Imagine all organisms in a biological population to have the same genotype described by string 101 and that a new variation is generated in one organism by a mutation in the first gene leading to 001. In this simulation, the mutation increases the organism's fitness from 0.40 to 0.70. As W(001) > W(101), organisms with genotype 001 will reproduce at a faster rate at the cost of organisms with genotype 101. After a sufficiently long period of selection, the population will be completely dominated by organisms that have the new string of genes 001, and

[6] Since the development of the NK-model by Kauffman and Levin (1987) and Kauffman (1989, 1993), many formal and statistical properties of NK-systems have been researched. For a good overview, see Altenberg (1997).
[7] Note that trial-and-error is close to, yet a simplified representation of Popper's (1959, 1963) description of scientific development as series of conjectures (*trials*) and refutations (*errors*). As long as a new conjecture does not yield higher "fitness" (*e.g.* goodness of fit with empirical data), the existing conjecture will be maintained. Only when a new conjecture leads to a higher fitness, the previous one is abandoned for the newly found one.

the next mutation will occur in an organism with genotype 001. If the fitness of string 001 would have been lower than string 101, organisms with genotype 001 would have reproduced at a slower rate than organisms with genotype 101, and string 001 would have disappeared from the population after a sufficient amount of time. In that case, the next mutation would have occurred again in an organism with genotype 101.[8]

Consider now a problem-solving designer that starts from design 101, and randomly picks the first element to be changed from 1 to 0 thus moving from 101 to 001 in the fitness landscape. Since the technology's functionality is improved, the designer accepts the trial, and the next mutation will start from the newly found design 100. If the fitness of design 100 would have been lower than design 000, the designer would have rejected the trial and the next mutation would then have occurred again in design 000. The concept of trial-and-error learning can thus be considered to be analogous to natural selection in biological evolution (Simon 1969 [1996]: 45-47).

To understand more generally the relationship between the complexity of an architecture and the properties of its fitness landscape, Kauffman simulated a large number of fitness landscapes for different values of K and N (Kauffman 1993: 33-67; Altenberg 1995, 1997). By tuning the parameter K from its lowest to highest value and comparing the properties of the fitness landscapes he found that:

1. The number of local optima increases exponentially with K. This means that for systems with higher complexity K, it becomes increasingly more likely that local trial-and-error leads to a local optimum rather than the global optimum.

2. The mean fitness of local optima is highest for systems with a positive but low complexity (around K=3 for systems containing eight or more elements). For higher N-values, the mean fitness of local optima for K=N-1 systems tends to 0.5, which is the expected value of a random draw from a uniform distribution between 0.0 and 1.0. Also note that K=N-1-systems of infinite size have exactly 0.5 fitness, as the mean of the fitness values of elements corresponds to the expected value.

3. The average correlation between local optima as measured by the number of alleles that two local optima have in common, decreases for increases in K. This means that the higher the complexity of a system, the more randomly spread the local optima are in design space.

4. The higher the fitness value of a local optimum in a given fitness landscape, the larger on average its "basin of attraction",[9] which is the number of strings from which local trial-and-error can end up in a local optimum. This implies that the probability of finding a local optimum with a high fitness value is higher then the probability to find a local optimum with a low fitness values. Since the variance of fitness values of local optima decreases for increasing K, the variance in the size of their basins of attraction also decreases for increasing K.

[8] Selection in terms of genotypes' shares in the population can be modelled using Fisher's equation, which states that reproduction rates of genotypes are proportionate to their fitness relative to the mean (Fisher 1930; Frenken 2001a).

[9] Note that basin of attraction metaphorically refers to basins that end up in minima. The term originally comes from minimisation problems, but can also be used with refence to maximisation problems. One can visualise basins of attraction by turning the fitness landscapes up-side-down, which renders the fitness peaks the basins.

What do these properties mean when the NK-model is taken as a model of technological design ? The first property holds that the more complex a technological system in terms of the number of interdependencies between its elements, the more difficult is becomes to find the optimal solution. The number of loal optima increases exponentially with K and the probability to end up in a sub-optimal solution also increases with K. Local trial-and-error thus only leads to a good solution of the design when the system under design is simple.

The second property holds that the highest fitness of a design is expected to be found for a system with an architecture consisting of a positive but moderate number of epistatic links. The reason why the simplest systems do not have optima with highest fitness is that no complementarities or "synergistic specificity" (Schilling 2000) can exist between the functioning of elements when epistatic relations are absent. And, the reason why very complex systems do not have local optima with highest fitness is that the many epistatic relations impose too many conflicting constraints *casu quo* trade-offs between elements. The highest average fitness of local optima of moderate complex systems compared to the average fitness of very simple and very complex systems suggests that successful technological systems contain some, but only a moderate number of epistatic relations between its elements.

The third property holds that the more complex a system, the more different local optima are in terms of their alleles. Very complex systems have local optima that are, on average, maximally dissimilar in terms of the alleles of elements, while less complex systems have local optima that are, on average, more similar. Combining this property with the second property, the NK-model suggests that successful designs have quite some alleles of elements in common.

The fourth property indicates that the higher the fitness of a local optimum in a given landscape, the larger its basin of attraction. This implies that better than average local optima have a higher probability to be found by local trial-and-error from a randomly chosen starting string. This property gets lost when the system gets more complex (higher K). Systems with high complexity are characterised by a low variance in mean fitness of local optima, and therefore similar sizes of basins of attraction. Combining this property with the third property, the NK-model suggests that different designers of systems with higher complexity are expected to come up with a larger variety in solutions.[10]

2.3 The generalised NK-model

The NK-model is not a general model with respect to the representation of the genotype (the set of elements) in relation to the phenotype (the set of functions). The NK-model is based on the idea that each element in a system performs an "own" sub-function within the system with regard to the attainment of one overall function on which selection operates (Kauffman 1993: 37). Each element n is conceived to have a particular fitness value w_n that reflects its functional contribution to the system as a whole. The fitness of the system as a whole is derived as the average of the fitness of individual elements.

A generalised model of complex systems described by N elements ($n=1,\ldots,N$) and F functions ($f=1,\ldots,F$) has been developed in the context of biology by Altenberg (1995, 1997) and will be used here as a model of technological design. In biological systems, for which this generalised NK-model was conceived, an organism's N genes are the system's elements and an organism's F

[10] Confer the concept of complex product systems (CoPS) introduced by Hobday (1998).

traits are the selection criteria. The string of genes constitutes an organism's genotype and the set of traits constitutes an organism's phenotype.

The genotype of an organism is the level at which mutations take place, which are transmitted in its offspring. The phenotype is the level at which natural selection operates in terms of its relative fitness to the environment. A mutation in a single gene may affect one or several traits in the phenotype and a single trait may be affected by a mutation in one or several genes in the genotype. The number of traits that is affected by a particular gene in the genotype is referred to as a gene's "pleiotropy". The number of genes that affects a trait in the phenotype is referred to as a trait's "polygeny". The structure of epistatic relations between genes and traits is represented in a "genotype-phenotype map".

Analogously, a technological system can be described in terms of its *N* elements and the *F* functions it performs.[11] The string of alleles of elements describes the "genotype" of a technological system, and the list of functions describes the "phenotype" of this system. Typical functions are speed, weight, comfort, safety, *et cetera*.[12] The architecture of a technology is represented by an "element-function matrix" of size FxN with:

$$M = [m_{fn}]\,,\; f = 1,...,F\;\;, n = 1,...,N \tag{2.4}$$

As in the NK-model, an epistatic relation is represented by x when function *f* is affected by element *n* and by – when function *f* is not affected by the element *n*. As explained by Altenberg (1997), an NK-system is a special case of the generalised element-function matrices. For NK-systems, it holds that the number of functions *F* equals the number of elements *N*. In the element-function matrix, this implies that the diagonal is always characterised by presence of a relation between element and function. Furthermore, the K-value in the NK-model implies that each function is affected by the same number of elements. Thus, in the NK-model the polygeny of each function is assumed to be equal to K (and pleiotropy being on average equal to K). Dropping these two restrictions provides a general model of complex systems. An example of an element-function matrix is given Fig. 3.

	n=1	n=2	n=3
w_1	x	x	-
w_2	-	x	x

FIGURE 3. Example of a generalised element-function matrix

The way in which fitness landscape are constructed for generalised element-function matrices follows the same logic as the original NK-model discussed in the previous section [10]. For each element that is mutated, all functions that are affected by this element are assigned a new,

[11] The distinction between elements and functions is similar to the description of technology as a system and its environment (Simon, 1969) and to the description of technology by a set of technical characteristics that are the object of design and service characteristics on which users select (Saviotti and Metcalfe, 1984).
[12] This perspective on fitness differs from the NK-model applied to process technology where fitness is expressed by a single cost criterion only (Auerswald *et al.*, 2000; Kauffman *et al.*, 2000).

randomly drawn value from the uniform distribution between 0 and 1. Total fitness is again derived as the mean of the fitness values of all functions:

$$W(s) = \frac{1}{\mathrm{F}} \cdot \sum_{f=1}^{\mathrm{F}} \mathrm{w}_{\mathrm{f}}(s) \qquad (2.5)$$

A simulation of the fitness landscape example of the element-function matrix given in Fig. 3 is given in Fig. 4 for all possible combinations between two alleles of three elements.

The concept of fitness landscape and of local optimum remains the same in Alternberg's generalised NK-model. The only difference compared to the original NK-model holds that the number of elements that does equal the number of function *per se*. Altenberg's (1997) model can therefore be considered as an important generalisation of the original NK-model of complex systems by Kauffman (1993).[13]

	w_1	w_2	W
000:	0.8	0.9	0.85
001:	0.8	0.6	0.70
010:	0.4	0.3	0.35
011:	0.4	0.2	0.30
100:	0.2	0.9	0.55
101:	0.2	0.6	0.40
110:	0.9	0.3	0.60
111:	0.9	0.2	0.55

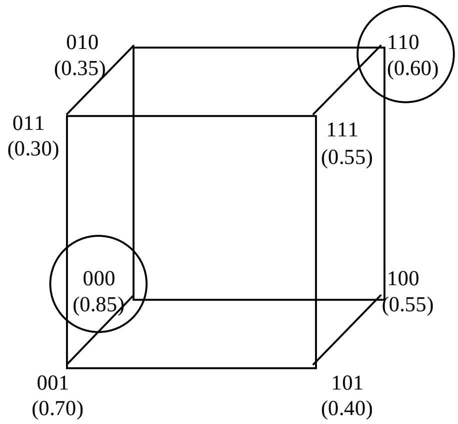

FIGURE 4. Simulation of fitness landscape of the matrix in Fig. 3

In the simulation in Fig. 4, a mutation in the first allele generates a random change in w_1, a mutation in the second allele a random change in both w_1 and w_2, and a mutation in the third allele a random change in w_2. Since only the second element affects both functions, the existence of local optima can only be related to different alleles of the first element (here 000 and 110). For each allele 0 or 1 of the second element, there exist an optimal set of other alleles. The simulation in Figure 3-1-3 shows a fitness landscape in which the *trade-off* between optimising the first function in the local optimum 110 ($w_1 = 0.9$, $w_2 = 0.3$) and optimising the second function in the global optimum 000 ($w_1 = 0.8$, $w_2 = 0.9$). Hill-climbing on this fitness landscape by means of a one-element mutation strategy can end up in either optimum, depending on starting string and the sequence of mutations that follow.

[13] An important line of research initiated by Altenberg (1995) using his generalised NK-model concerns the understanding of systems that evolve through the addition of new elements in the system. This line of research may well yield interesting possibilities for research in the various systems of reference in economics (artefacts, production processes, organisational startegies), the size of which generally increase by means of the addition of new elements.

2.4 Previous research on the NK-model in evolutionary economics

The NK-model has motivated a number of researchers to specify more detailed models in the context of evolutionary economics. As stated earlier, three systems of reference have been used. First, the set of strategies that determine its performance have been modelled as a complex evolving system. The complexity stems from the interdependence of different strategies in their success. Very often, the success of a choice for a particular strategy depends on the choice of strategy along other dimensions (Levinthal 1997; Gavetti and Levtinthal 2000; Rivkin 2000). The evolution of the set of firm strategies can then be understood as a hill-climbing process in which a firm's management occasionally experiments with a new strategy along one dimension and accepts the experiment when proven successful.

Second, a number of models have used the NK framework to represent a production technology (Auerswald *et al.* 2000; Kauffman *et al.* 2000). In this context, the elements of a system stand for particular tasks and the alleles for possible ways to carry out the task. The design space then stand for all possible combinations to organise a production process to produce a given output and the fitness values stand for the efficiency of each combination. This use of the NK-model follows the neoclassical framework in that the design space stands for a possibility space to produce a given homogeneous output and each possible technology is evaluated on its efficiency. Of course, this model differs from neoclassical production theory in that the interdependencies between the different tasks in which a production process is organised are taken into account. Furthermore, firms in this model are not assumed to know at beforehand all fitness values of all possible technologies. These values can only be known by means of hill-climbing the fitness landscape. Simulation research has shown that hill-climbing by local optimisation of individual tasks leads a firm to improve its efficiency according to a classical learning curve reflecting learning-by-doing in production (Arrow 1962). This can be considered an important results that suggests that the shape of learning curves can be explained by the complexities among tasks in a production process (which in turn may relate to complexities among elements in the artefact that is being produced such as an aircraft).

Third, there is a number of works that takes complex technological artefacts as the system of reference, as we do here (Valente 2000; Frenken 2001a,b). In these approaches, complexity stems from interdependencies between different elements in a technological system. The generalised NK-model can be considered in this context as a model in which the distinction between technical and service characteristics has been operationalised (Frenken 2001a; *cf.* Saviotti and Metcalfe 1984). The alleles of the N elements stand for technical characteristics and the fitness values of the F functions. This model has been further elaborated to include the classification of Henderson and Clark (1990) of four types of innovation (Frenken 2001a). These innovation types concern modular innovation, architectural innovation, incremental innovation and radical innovation. In the NK framework, modular innovation corresponds to a mutation in an element while architectural innovation corresponds to a change in the set of epistatic relation between elements and functions. An incremental innovation concerns an improvement in the system without a mutation in an element and without a change in epistatic relations while a radical innovation refers to the co-occurrence of a mutation in an element and a change in epistatic relations. Furthermore, as explained in Frenken (2001a), the NK-model of complex technological artefacts can formally define a technological paradigm as a set of designs that share the same alleles for high pleiotropy elements. The paradigmatic nature of these designs is explained by the fact that mutations in elements with high pleiotropy seldom increase a system's fitness (Altenberg 1995). Due to the large number of functions that is affected by a high-pleiotropy element, the positive effect of a mutation on some elements is generally outweighed by the negative effect on other elements.

What is important to recognise, whatever system one takes as reference, is that the concept of hill-climbing towards peaks in a fitness landscape pre-assumes learning through random mutation in single elements. However, humans can use many alternative search algorithms to search a fitness landscape. For example, there are algorithms that mutate more than one element at the same time. Such an algorithm would be able to escape a string that is a local optimum with respect to an algorithm that mutates only in single elements at the time. Thus, a string that is a local optimum with respect to mutations in single elements need not be a local optimum with respect to mutations in several elements at the same time (Frenken *et al.* 1999; Kauffman *et al.* 2000).

Varying the number of elements that is allowed to mutate at the same time is not the only dimension in which search strategies can differ. Gavetti and Levinthal (2000) modelled another type of heuristic that agents can apply in hill-climbing a NK fitness landscape. The agents that try to optimise a complex system by means of hill-climbing use a heuristic that specifies a subset of the N elements of the system. Mutations are allowed only in this subset of elements. In other words, it is assumed that agents do not take into account all elements of a system, but restrict mutation to take place in a subset of elements. This specification of search is consonant with the concept of "bounded rationality" that holds that designers are cognitively restricted in the number of dimensions that they can take into account in the development of a new design (Simon 1955, 1969; Allen 1994: 9; Metcalfe 1994: 935). Instead, designers apply "mental maps" that frame the high-dimensional system as a system with a lower number of dimensions. The mental map consists of the subset of elements in which mutation can take place. In this way, they reduce the size of the design space as to speed up the search process to a local optimum.[14]

Another line of research has been explored by Rivkin (2000). His central question holds how one can assess the performance of imitative strategies. The possibility of imitation in complex fitness landscapes implies that designers that find local optima with relative low fitness compared to the fitness of local optima found by competitors, could improve their position by imitation of the string of alleles corresponding to local optima with relatively high fitness (Alchian 1950; Nelson and Winter 1982). Imitation of relatively good solutions by designers that previously found relatively bad solutions "speeds up" the selection process. From the NK-model, it can be derived that for an imitation to be successful, it is crucial that a firm succeeds in copying *exactly* all alleles of a local optimum that is occupied by a more successful firm. Any error in copying alleles can result in a drastic fall in fitness due to epistiatic relations between elements. And, the loss in fitness due to a copying error increases as a function of the complexity parameter K of the system being copied. Therefore, one expects imitation strategies to be the more successful the less complex the system being imitated.

Finally, a heuristic has been formulated that makes use of information on the selection environment (Frenken 2001a). In short, this heuristic refers to the general case in which the fitness of a system is not determined by the average of the fitness values of its *F* functions but by a weighted sum of the fitness values of its *F* functions. Then, some functions are weighted more

[14] Yet another example of a search heuristic, which is related to the former one but which has not been explored so far, is the application of a mental map regarding the matrix of epistatic relations between elements and functions. This mental map does not refer to a subset of the elements that are candidate for mutation as in Gavetti and Levinthal (2000), but concerns a subset of epistatic relations between elements and functions that are candidate for architectural innovation.

than others, so it is expected to payoff when search concentrates on the more important functions at the costs of sacrificing fitness values of less important functions. By confining early search to elements that affect the more important functions one is secured that the final solutions performs best in the most weighted functions. In this way, search time can be reduced as the number of relevant elements to be searched is reduced, while the fitness of the final solution is still expected to be quite high.

Summarising, the NK-model and generalised NK-model of complex fitness landscapes have motivated a promising research program concerning the evaluation of alternative types of heuristics. Elsewhere, I termed the meta-space of possible search strategies the set of "innovation rules" and attempted to specify a number of core dimensions of this meta-space (Frenken 2001b). The set of dimensions of innovation rules span a meta-space of possible innovation rules in which agents have to choose. The NK-model and the generalised NK-model provide one with formal frameworks to assess what types of innovation rules perform best in terms of the fitness of the local optimum that is reached and the time required to reach this local optimum.

3. MODELLING THE ORGANISATION OF INNOVATIVE ACTIVITY

The remainder of this paper will concentrate on a second research program that can be based on the (generalised) NK-model. The central topic of such a program does not concern the way an agent searches a landscape as above, but how this search activity is organised. In short, the organisation of innovative activity deals with the incentive structure for participating agents that jointly develop a new technology. The problem how to coordinate innovative activity is central as the large majority of new technologies in modern society is developed by many actors working on different parts of the artefact to be designed, or more generally, on different dimensions of the problem to be solved (Kauffman and Macready 1995; Marengo *et al.* 2000).

In the NK-model, one can represent of a spectrum of degrees of coordination in innovative activity. At the one end of the spectrum coordination is fully decentralised and at the other end of the spectrum coordination is fully centralised. The latter case of fully centralised control over the choice of allele for each element corresponds with the perspective taken above and in the majority of NK-models. In these models, there is a single agent that decides for all elements whether these are mutated or not, and evaluates each mutation on the basis of its effect on the fitness W(s) of the system *as a whole*. In the former case of fully decentralised search, there are N agents ($n=1,\ldots N$) who each can decide to mutate its one element or not and evaluate a mutation only with respect to the fitness $w_n(s_n)$ of its own element (Kauffman and Macready 1995). Decentralisation thus implies "anarchy" in that each agent is autonomous in deciding whether to mutate its own element without having any control of other elements.[15] Note that in the context of collaborative design, agents can be individuals in a team, teams in an firm, and firms in an inter-firm network. In all cases, the question is how the incentive structure is organised with regard to the pay-off structure of mutations in different elements in a complex technological system.

In the following, we will use the original (non-generalised) NK-model as the model of reference. As the number of functions is equal to the number of elements in the original NK-model, it allows us to define the pay-off for an individual agent by the fitness value of one individual element. After the exposition, some remarks follow on the conceptualisation of decentralised search in fitness landscapes where the number of elements (firms) does not equal the number of functions.

[15] Confer the conceptual model presented in Frenken (2000) of interdependencies between producers, users and governments in the development of new aircraft programmes.

3.1 Centralised *versus* decentralised search

To explain the difference between centralised and decentralised search in the NK-model, consider the example of a fitness landscape of *N=3* and *K=2* in Fig. 5. When search is centralised and takes place by means of local hill-climbing, there are to optima: strings 010 and 100. For these strings, it holds that any mutation in one element would lower the total fitness of the system as a whole *W*. By contrast, when search is decentralised there is only one optimum: string 010. Only for this string it holds that any mutation in one element would lower the fitness value of the individual elements w_n. Consequently, once the three agents have found string 010, no single agent has an incentive to mutate its own element.[16] However, 100 would be optimal from the user's point of view, as the average fitness W of all elements is highest for this design. However, though optimal for users, design 100 will not be accepted by the agent controlling the second element, since this firm can improve its individual fitness by mutating from 0 to 1 moving from design 100 ($w_2 = 0.5$) to design 110 ($w_2 = 0.9$).

	w_1	w_2	w_3	W
000:	0.5	0.1	0.7	0.43
001:	0.2	0.2	0.8	0.40
010:	0.7	0.8	0.6	0.70
011:	0.6	0.5	0.3	0.47
100:	0.9	0.5	0.8	0.73
101:	0.2	0.3	0.4	0.30
110:	0.5	0.9	0.4	0.60
111:	0.4	0.8	0.1	0.43

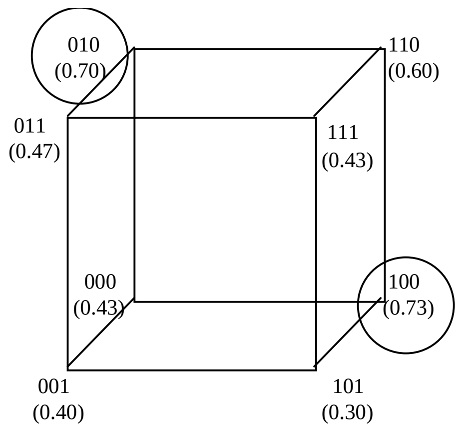

FIGURE 5. Simulation of fitness landscape of system with N=3 and K=2

As shown in simulations by Kauffman and Macready (1995), decentralised control generally does not optimise a complex system. In many cases, the strings corresponding to optima when search would be centralised, do not correspond to optima when search is decentralised. The reason that fewer optima exist for fully decentralised search compared to fully centralised search is that for strings corresponding to optima in centralised search, there is generally at least one agent that can improve its own fitness by mutation of its element. Moreover, decentralised search runs the risk of finding no optimum at all when it holds for all strings that at least one agent can improve by

[16] This equilibrium is generally called a Nash-equilibrium To verify whether design 010 is indeed a Nash-equilibrium, one can look at the payoffs for each firm and check whether each firm cannot improve its payoff by mutation. Payoffs are $w_1(010) = 0.7$ for the firm responsible for the first element (FIRM1), $w_2(010) = 0.8$ for the firm responsible for the second element (FIRM2), and $w_3(010) = 0.6$ for the firm responsible for the third element (FIRM3). A mutation by FIRM 1 would lead to design 110 and payoff $w_1(110) = 0.5$, a mutation by FIRM 2 would lead to design 000 and payoff $w_2(000) = 0.1$, and a mutation by FIRM 3 would lead to design 011 and payoff $w_3(011) = 0.5$.

mutation. And, since fewer optima exist for fully decentralised search, it generally takes more mutations to find this optimum than in the case of centralised search.

Another result found by Kauffman and Macready (1995) holds that, although there are fewer optima in decentralised search, the average fitness of these optima are higher than the optima found by centralised search. This result is understandable since optima in decentralised search have to meet the hard criterion that all N agents cannot improve their fitness by means of mutation of their element. Thus, although decentralised search takes longer to find a local optimum (also called an equilibrium) and even runs a serious risk of finding no local optimum at all, the expected fitness of the local optima in decentralised search is higher than the expected fitness of local optima for centralised search.

Kauffman and Macready (1995) argued that both fully centralised and fully decentralised search suffer from serious deficiencies in optimising complex systems by means of local hill-climbing. Under fully centralised search, an agent generally ends up in poor local optima as many strings correspond to solutions with low fitness.[17] Under decentralised search, the collective search behaviour by agents generally leads to better optima, but the search process generally takes much longer compared to centralised search. The central research question to be addressed in this context holds whether alternative forms of coordination can be specified that can overcome the problems of decentralised control while avoiding the high search costs of exhaustive search under centralised control.

Kauffman and Macready (1995) studied a form of coordination that is intermediate between fully centralised and fully decentralised coordination. This intermediate form of coordination refers to the case in which there are several agents, each of which has exclusive control over more than one but less than all elements. The number of elements controlled by each agents is denoted by *P* with *2 • P • ½N*.[18] With regards to its block of *P* elements, an agent can mutate individual elements.[19] Each mutation of one element is evaluated on the effect on all *P* elements controlled by a single agent thus ignoring the effects a mutation might have on the other elements controlled by other agents.[20]

The partitioning of elements over agents is called "patching" and each block of elements that is controlled by a single agent is called a "patch". By means of tuning the *K* parameter, Kauffman and Macready (1995) found that patching leads to better optima compared to centralised search when *K* exceeds a critical value. Furthermore, these authors found that the optimal patch size is only a fraction of the size N of the system. This result indicates that a major reduction in search time and search costs can be achieved by patching since the optimal patch size is considerably smaller than the size of the system.

[17] Note again that we assume one-element mutations only. Importantly, centralised search has the advantage over decentralised search that one is not restricted to mutations in single elements only. One can also mutate several at the same time, which allows one to escape poor local optima. However, the search time required to find a local optimum increases exponentially with the number of elements that is allowed to mutate at the same time. And, the global optimum generally can only be found by exhaustive search, i.e. by evaluating all 2^N possible designs (Frenken *et al.* 1999).
[18] When *P = 2* there is a maximum of N/2 firms all controlling two elements, and when *P = ½N* there is a minimum of two firms both controlling half of all elements. *P=1* corresponds to fully decentralised coordination of innovation with N firms and *P=N* corresponds to fully centralised coordination with one firm.
[19] Again, mutations in more than one element at the time are not allowed in this model.
[20] Note that the patch model is formally equivalent to Kauffman's (1993) *NKC*-model.

3.2 Networking in NK-models

An alternative form of coordination between agents, which has not yet been addressed in NK-models, takes the form of networks between agents. In this view, each agent controls again only one element, but patches of several elements emerge from collaborations between agents. When *Q* agents decide to form a network of size *Q* involving *Q* elements of the system, this means here that a mutation by one agent in one element is accepted or rejected depending on its effect on the fitness values of all participants in the network. For example, one can assume that a mutation by one agent in the network is accepted only if no participant in the network is worse off (win-win-win). Using different rules of acceptance of an innovation of an agent in a network of agents, one can analyse the performance of different network rules for different parameter settings *N, K*, and *Q*.

In this study, we restrict ourselves to analytical results that can be deduced from a particular class of landscapes being fully random NK landscapes of maximum complexity (*K=N-1*). For landscapes with maximum complexity it holds that any mutation in an element affects the fitness values of all other elements. Therefore, each mutation will yield *N* new fitness values that are randomly drawn from a uniform distribution [0,1] and consequently there is exist no correlation whatsoever between fitness values of a string and its neighbouring strings. This property renders NK landscapes amenable for analytical analysis. Other landscapes will have to be researched by simulation.

The analysis on fully random NK landscapes will concentrate on the difference between centralised, decentralised, and network organisations. For all three modes of organisation one can derive the expected number of equilibria.

For centralised organisation, it holds that the number of local optima can be derived by looking at the probability that a string is superior to all its one-mutant neighbours. This probability being 1/N+1, the total expected number of local optima adds up to (Kauffman 1993):

$$L(centralised) = 2^N / N + 1 \tag{3.1}$$

For fully decentralised organisation, the number of local optima called - equilibria in this context - is derived by looking at the probability that a string is a Nash-equilibrium. Given that each of the N agents can increase its fitness by mutating its element with probability half, the probability that no single agent can increase its fitness by mutation (i.e. that the current string is a Nash-equilibrium) is 2^{-N}. The total number of expected equilibria thus holds:

$$L(decentralised) = 2^N \cdot 2^{-N} = 1 \tag{3.2}$$

Thus, whatever the size of the system N, on average there is only one equilibrium in a decentralised organisation. This means that in many cases, there will be no equilibrium at all. This is the case when for all strings there exist always at least one agent that can improve its fitness. In such a case, there can never be consensus on the technological design and decision-making will not result in an outcome.

In the case of networking, the specific rule that agents follow when deciding to collaborate affects the number of equilibria. Given that networking in technological development is motivated by coordinating parallel changes in a system, one can assume that agents will mutate when a simultaneous mutation of their partner(s) does not decrease an agent's fitness and *vice versa*. In other words, networking occurs when a simultaneous mutation by all partners generates a win-win-win situation when a simultaneous mutation renders no one worse off than before. An equilibrium (*i.e.* a local optimum) is then defined by the condition that there exist no single network between agents that leads the participating agents to a win-win-win situation. For example, in Fig, 5 there are four strings that meet this criterion (010, 011, 100, and 110).

The number of equilibria for which holds that no network collaboration yields a win-win-win situation for the participating agents thus depends on the size of the networks that are allowed to form. Obviously, the larger the size of a network, the harder the constraint of a win-win-win situation. We will analyse the probabilities for networks of a given size *Q*.

For bilateral networks (*Q=2*) we have (N(N-1))/2 possible collaborations. The probability that a simultaneous mutation by both partners increase both fitness values equals 2^{-2} = ¼. Consequently, the probability that a simultaneous mutation fails equals (1 - 2^{-2}) = ¾. The probability that any possible collaboration fails, *i.e.* that the set of agents in equilibrium holds:

$$L(bilateral) = 2^N \cdot 0.75^{(N(N-1))/2} \qquad (3.3)$$

Generally, for networks of size *Q* we have probability 2^{-Q} that a simultaneous mutation of all *Q* partners will lead to a win-win-win situation and thus probability (1 - 2^{-Q}) that the collaboration will fail. Thus, the larger the size of a network, the smaller the probability that a win-win-win situation can be reached. To derive the number of equilibria for network of given size, one also needs to determine the number of possible networks of size *Q* for a given *N* which we will call *R* and which equals:

$$R(Q) = N \cdot \prod_{q=1}^{Q-1} ((N-q)/(q+1)) \qquad (3.4)$$

This parabolic function first increases and then decreases indicating that the number of possible networks for a given *N* and varying Q first increases with Q and then decreases with Q.

The expected number of equilibria as a function of the number of participants in a network *Q* becomes:

$$L(Q) = 2^N \cdot (1-2^{-Q})^{N \cdot \prod_{q=1}^{Q-1}((N-q)/(q+1))} \qquad (3.5)$$

In Fig. 6 the expected number of equilibria has been plotted for different values of *N* and *Q*. From the results the following proposition can be derived. For each size of the network, the expected number of equilibria has a maximum at some value of *N*. When varying N for a given Q, the expected number of equilibria first rises and then falls (which can be understood from the fact that when network size continues to increases, the number of possible networks that can be formed falls).

One can now judge the performance of a network in terms of the expected number of equilibria it generates compared to the alternative fully centralised and fully decentralised organisational forms. As long as the expected number of equilibria of a particular network *Q* for a particular value of *N* exceeds the expected number of equilibria of centralised search, the average fitness value W of these equilibria can be expected to be below the fitness values W of equilibria found by centralised search. However, when the expected number of equilibria of a particular network *Q* for a particular value of *N* is *in between* the expected number of equilibria of centralised and decentralised search, the average fitness value W of these equilibria can be expected to be above the fitness values W of equilibria found by centralised search. In the latter case, the network organisation is expected to perform better.

From the results in Fig. 6 it is clear that for each value of the network size Q, there is a region of values for N where the number of optima lies in between the number of equilibria of centralised and decentralised. Furthermore, the smallest possible networks size (*Q=2*) generates less local optima than centralised search and more than decentralised search only for N=4. Systems smaller than N=4 are best optimised by centralised search. Thus, the network organisation only performs well when the size of a complex system exceeds a critical size. This suggests that centralised decision-making can indeed be considered more effective for relative small systems, while networking becomes only effective once a critical system size is exceeded. Moreover, the preferable network organisation depends on the system size *N* with preferable network size increasing with increases in *N*.

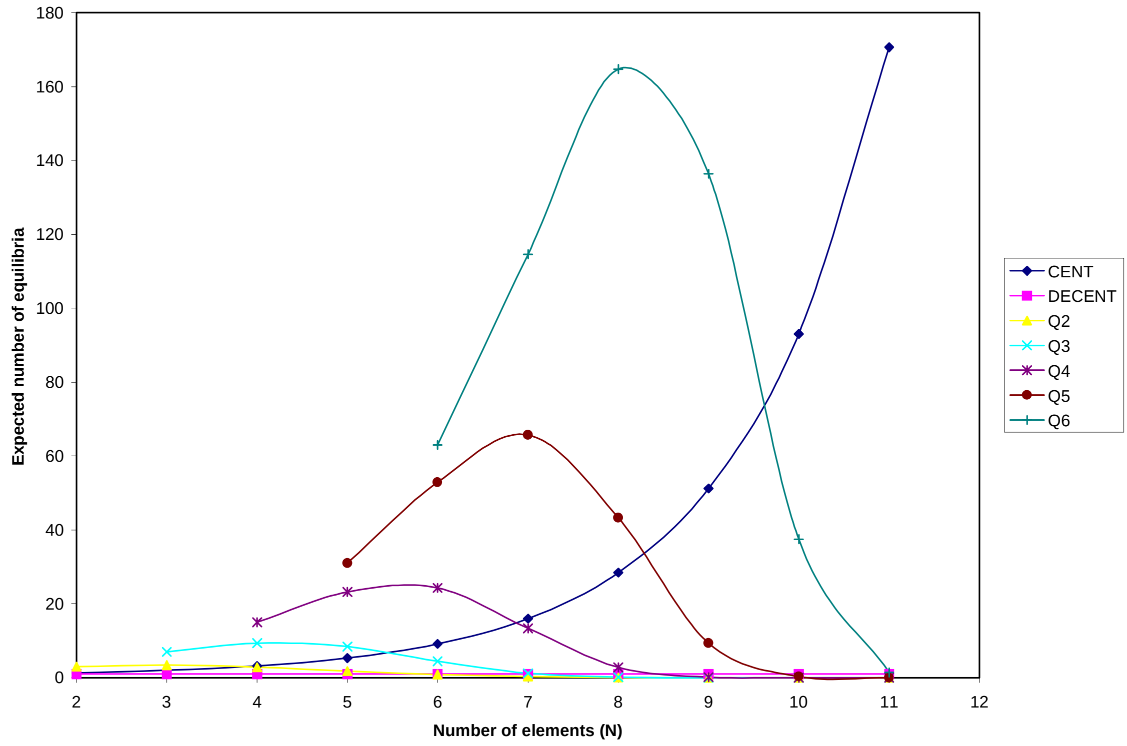

FIGURE 6. Expected number of equilibria on fully random fitness landscapes (K = N – 1) plotted against different values of N for fully centralised organisation (CENT), fully decentralised organisation (DECENT) and different values of network size Q (Q2 … Q6)

As the number of equilibria is critically dependent on the rules that allow firms to network, alternative rules for collaboration can be formulated and assessed in terms of the expected number of equilibria. For example, if one would put a restriction on the formation of possible networks, for example, that particular agents are not allowed to network, then the number of possible networks is less than as above. In that case, the curvatures of the expected number of equilibria of netwoprk organisation may cover a larger region where networking is favourable compared to the alternative modes. Further simulation research of this kind can be straightforwardly be implemented in the model set up as described here.

Further note that the discussion has been limited to fully random fitness landscapes ($K=N-1$). Further simulation research should focus on alternative fitness landscape of lower complexity. In that case, the expected number of equilibria of all modes of organisation will fall but the precise curvatures cannot be predicted *ex ante*. And, using simulation exercises, the average fitness value of equilibria can also be determined more precisely.

4. TOWARDS A RESEARCH PROGRAM

The formalisation of network organisations in searching a complex fitness landscape as specified above provides one with a rich framework in which many research questions can be addressed in simulation models. In this study, starting from the original NK-model, a model has been developed which allows one to determine what type of organisational modes yields a favourable number of equilibria.

The generalised NK-model as discussed earlier in this paper makes it possible to analyse systems with a number of functions that does not equal the number of elements. This generalised model can be expected to have substantial empirical relevance too, when studying networks in which pay-offs are not distributed individually but to teams of agents responsible for a particular function of the system.

In a future exercise of analysing alternative modes of organisation in generalised fitness landscapes, further theorising is required regarding the question how the fitness values of *F* functions are translated into *N* pay-off values for individual agents controlling an element. As the number of functions that is affected can differ per element, the “position” of an agent controlling an element matters. Agents that control an element that affects many functions (high pleiotropy) have proportional more possibilities to collaborate.

Simulation exercises are required to understand the performance of alternative modes of organisation for parameter regimes that cannot be handled analytically. Once these properties of generalised NK-models are known, empirical examples of network collaboration between agents (individuals in a team, teams in a department, departments in firms, firms in an industry) can be researched in relation to the simulation outcomes and may provide new research questions for further specification of the model. It is hoped that the full complexity of technological alliances can then be appreciated both in simulation models and empirical studies.

References

Alchian, A.A. (1950) 'Uncertainty, evolution, and economic theory', *Journal of Political Economy* 58, pp. 211-221.

Alexander, C. (1964) *Notes on the Synthesis of Form* (Cambridge MA & London: Harvard University Press, 13th edition, 1994).

Allen, P.M. (1994) 'Evolutionary complex systems: models of technological change', pp. 1-18 in: Leydesdorff, L., Van den Besselaar, P. (eds.) *Evolutionary Economics and Chaos Theory: New Directions in Technology Studies* (London, etc.: Pinter).

Altenberg, L. (1995) 'Genome growth and the evolution of the genotype-phenotype map' pp. 205-259, in: Banzhaf, W., Eckman, F.H. (eds.) *Evolution and Biocomputation* (Berlin & Heidelberg: Springer-Verlag).

Altenberg, L. (1997) 'NK fitness landscapes', in: Back, T., Fogel, D. Michalewicz, Z. (eds.) *The Handbook of Evolutionary Computation* (Oxford etc.: Oxford University Press).

Arrow, K.J. (1962) 'The economic implications of learning by doing', *Review of Economic Studies* 29, pp. 155-173.

Auerswald P., Kauffman S., Lobo J., Shell K. (2000) 'The production recipes approach to modeling technological innovation: an application to learning-by-doing', *Journal of Economic Dynamics and Control* 24, pp. 389-450.

Birchenhall, C. (1995) 'Modular technical change and genetic algorithms', *Computational Economics* 8, pp. 233-253.

Bradshaw, G. (1992) 'The airplane and the logic of invention', pp. 239-250 in: Giere, R. (ed.) *Cognitive Models of Science* (Minneapolis: University of Minnesota Press).

Cooper, B. (2000) 'Modelling research and development: how *do* firms solve design problems?' *Journal of Evolutionary Economics* 10, pp. 395-414.

Cyert, R.M., March, J.G. (1963) *A Behavourial Theory of the Firm* (Englewood-Cliffs: Prentice-Hall).

Dennett, D. (1995) *Darwin's Dangerous Idea* (London: Penguin, paperback edition, 1996).

Fisher, R.A. (1930) *The Genetical Theory of Natural Selection* (Oxford: Clarendon Press).

Foray, D., Grübler, A. (1990) 'Morphological analysis, diffusion and lock-out of technologies: ferrous casting in France and the FRG', *Research Policy* 19, pp. 535-550.

Frenken, K. (2000) 'A complexity approach to innovation networks. The case of the aircraft industry (1909-1997)', *Research Policy* 29, pp. 257-272.

Frenken, K. (2001a) *Understanding Product Innovation using Complex Systems Theory*, Ph.D. Thesis, University of Amsterdam.

Frenken, K. (2001b) 'Innovation rules: the NK-model as a framework for memetics', *Journal of Memetics*, forthcoming.

Frenken, K., Leydesdorff, L., Saviotti, P.P. (1997) 'Self-organisation and the economics of technical change', in: *Actes de 46me Congrès Annuel de l'Association Française de Science Economique* (Paris: Association Française de Science Economique).

Frenken, K., Marengo, L., Valente, M. (1999) ‘Interdependencies, near-decomposability and adaptation’, pp. 145-165 in: Brenner, T. (ed.) *Computational Techniques for Modelling Learning in Economics* (Boston etc.: Kluwer Academic Publishers).

Gavetti, G., Levinthal, D. (2000) ‘Looking forward and looking backward: cognitive and experiential search’, *Administrative Science Quarterly* 45 (1), pp. 113-137.

Henderson, R., Clark, K. (1990) ‘Architectural innovation’, *Administrative Science Quarterly* 35, pp. 9-30.

Hobday, M. (1998) ‘Product complexity, innovation and industrial organisation’, *Research Policy* 27, pp. 689-710.

Holland, J.H. (1975) *Adaptation in Natural and Artificial Systems* (Ann Arbor: University of Michigan Press).

Holland, J.H. (1992) ‘Genetic algorithms’, *Scientific American* 267 (July), pp. 66-72.

Hughes, T.P. (1987) ‘The evolution of large technological systems’, pp. 51-82 in: Bijker, W. Hughes, T.P., Pinch, T.J. (eds.) *The Social Construction of Technological Systems* (Cambridge, Mass.: MIT Press).

Jervis, R. (1997) ‘Complexity and the analysis of political and social life’, *Political Science Quarterly* 112, pp. 569-594.

Kauffman, S.A. (1989) ‘Adaptation on rugged fitness landscapes’, pp. 527-618 in: Stein, D.L. (ed.) *Lectures in the Sciences of Complexity* (Redwood City, California: Addison-Wesley).

Kauffman, S.A. (1993) *The Origins of Order. Self-Organization and Selection in Evolution* (New York & Oxford: Oxford University Press).

Kauffman, S.A., Levin, S. (1987) ‘Towards a general theory of adaptive walks on rugged landscapes’, *Journal of Theoretical Biology* 128, pp. 11-45.

Kauffman, S.A., Lobo, J., Macready, W.G. (2000) ‘Optimal search on a technology landscape’, *Journal of Economic Behavior and Organization* 43, pp. 141-166.

Kauffman, S.A., Macready, W.G. (1995) ‘Technological evolution and adaptive organizations’, *Complexity* 1, pp. 26-43.

Levinthal, D.A. (1997) ‘Adaptation on rugged landscapes’, *Management Science* 43, pp. 934-950.

Marengo, L., Pasquali, C., Valente, M., Cefis, E. (2000). ‘Decentralisation and market mechanisms in collective problem-solving’, *Paper presented at the Eighth International Schumpeter Society Conference*, Manchester, 28 June - 1 July.

Metcalfe, J.S. (1995) ‘Technology systems and technology policy in an evolutionary framework’, *Cambridge Journal of Economics* 19, pp. 25-46.

Nelson, R.R., Winter, S.G. (1982) *An Evolutionary Theory of Economic Change* (Cambridge MA & London: Belknap Press of Harvard University Press).

Popper, K.R. (1959) *The Logic of Scientific Discovery* (New York: Basic Books).

Popper, K.R. (1963) *Conjectures and Refutations* (London: Routledge & Kegan Paul).

Rivkin, J.W. (2000) ‘Imitation of complex strategies’, *Management Science* 46, pp. 824-844.

Saviotti, P.P., Metcalfe, J.S. (1984) ‘A theoretical approach to the construction of technological output indicators’, *Research Policy* 13, pp. 141-151.

Schilling, M.A. (2000) ‘Toward a general modular systems theory and its application to interfirm product modularity’, *Academy of Management Review* 25, pp. 312-334.

Simon, H.A. (1955) ‘A behavioral model of rational choice’, *Quarterly Journal of Economics* 69, pp. 99-118.

Simon, H.A. (1969) *The Sciences of the Artificial* (Cambridge MA. & London: MIT Press, 3rd edition, 1996).

Valente, M. (2000) *Evolutionary Economics and Computer Simulations. A model for the evolution of markets*. PhD thesis, Aalborg University, www.busines.auc.dk/~mv.

Williamson, O.E., 1991, Comparative economic organization: the analysis of discrete structural alternatives, *Administrative Science Quarterly*, 36, pp.269-296.

Windrum, P., Birchenhall, C. (1998) ‘Is product life-cycle theory a special case? Dominant designs and the emergence of market niches through coevolutionary learning’, *Structural Change and Economic Dynamics* 9, pp. 109-134.