Consumer Preferences and Technological Innovation in the Evolution of Markets*

Marco Valente †
mv@business.auc.dk

September 8, 2003

Abstract

The study of markets' evolution is almost universally focused on the role of the supply side of market, relegating the demand side to a relatively minor role, moreover based on exogenous factors. In this work I try to make a more equilibrate assessment providing a model where demand and supply co-evolve influencing each other's action. The model implements sets of consumers that determine their preferences under the influence of sellers' marketing strategies, and apply their preferences to make purchasing decisions based on products' qualities. Suppliers undertake decisions concerning the allocation of R&D and marketing budget, for product improvements and demand biasing respectively, based on the survey of consumers' behaviour.

Though stylised, the model is a useful tool for reasoning about the interplay of demand and supply. Initially the the paper focuses on the two sides of the market in isolation, to better describe the assumptions used. Such partial analysis already provide useful insights. It is shown that the dynamics of demand only can produce several market configurations, usually explained as the results of suppliers' behaviour. The analysis of the R&D policies of producers highlight the reasons for trade-off's between speed of returns and overall levels of technological improvements. Finally, the whole model tells a virtual history where firms face a "loose-loose" choice between focusing on short- or long-term results. In the first case they gain immediate financial benefits, but waste the possibility to keep up with general technological innovations. In the latter case they obtain the necessary technological competence for continuous technological leadership, but lack the funds to finance it under-performing competitors in the short term.

A hidden goal of this paper is to suggest a methodologically robust use of abstract computer simulations. Partial models (neglecting important elements of real world counter-parts) containing thousands of parameters (hindering the possibility of calibrating) are used as sand-boxes to investigate causal relationships between economic phenomena. The hard-proven explanations induced by the study of virtual markets are shown to shed light on relevant aspects, otherwise left to the sensibility of economic researchers. It is suggested that the lack of formal results in economics, as compared to hard sciences, does not depends on a sloppy methodological approach, but on the very nature of the object of research, concerning historical events of human behaviours.

*Paper prepared for a seminar in Ferrara, 8 September 2003. The paper if a broad summary of the author's PhD thesis, supervised by Prof. E.S.Andersen, discussed in January, 2000.

†Università dell'Aquila, Italy, and DRUID, University of Aalborg, Denmark.

Contents

1 Introduction

Most of the literature on market dynamics, and especially the works focusing innovation, center their interest on the supply side of markets. Demand is either neglected, or, even when it has an important role, is considered exogenous to the events taking place during the evolution of markets.

In this work it is suggested that part of the specification of demand are endogenously determined by the behaviour of the supply side, namely through marketing[1]. The first section proposes a general model of consumers, represented as bounded rational decision makers, whose preferences are influenced by producers.

Normally, models studying technological innovations assume homogeneous products because of the lack of a demand side adapt to treat heterogeneous products. Once this is no longer a constraint, as with the proposed model, we have the possibility to study product innovations, instead of limiting the analysis to cost-cutting process innovations. The third section of this work analyses technological research as the search of solutions for a complex problem, transferring some of the results from the abstract analysis of complexity to the specific issue of technological research in economically competitive environment.

The last section before the conclusions, joins the two models of demand and technological innovation to provide a model representing the evolution of markets. Such model of markets is extremely simplified in many respects, but, because of this simplicity, allows a detailed analysis of the interplay between demand and forms of technological innovation.

The presentation of the various models mentioned in this work will be very general, and lacking any formalism, in order to allow an easier appreciation of their scientific contents. For a more detailed account of the models see Valente (2000), or request the models, complete with documentation and users' manual, to the author.

2 Modelling Demand

Decisions of consumers are made, in the vast majority of cases, by people who are not perfectly informed on the technical aspect of the product they purchase. Moreover, there is no economic punishment in making the "wrong" decision which can either press a consumer to correct its decision in the future, or select him out of the market. Lastly, the majority of decisions by final consumers concern goods of relatively scarce importance in respect of the overall life of a person, who can therefore be assumed to not devote a large amount of time and attention. All the three considerations above suggests that the most appropriate way to model consumers' behaviour is a bounded rational one (Simon, 1982), rather than the optimal decision theory used in mainstream economics.

Though attractive on a theoretical side, the exact meaning of bounded rationality lacks a generally agreed upon positive definition on which a model can be built. In fact, most of the definitions of bounded rationality are based on negative arguments: they tell what

[1] In the following "marketing" is used in a very general sense as any activity by firms aimed at influencing actual or potential consumers, besides the direct trade of the good or service.

bounded rationality is not, typically perfect rationality. However, Gigerenzer (see, for example, Gigerenzer and Goldstein, 1996) presses for an interpretation of Simon's work that does not makes the opposition perfect vs. bounded rationality. His argument is that the perfect rationality describes a quality of a decision (i.e. the best), while the bounded rationality refers to a process of decision (the viable), namely the ones used by human beings. There can be cases, typically for simple problems, in which bounded rational and perfectly rational results coincide. This author proposes several ways to implement processes of decision making that aim both at being rational, providing good results, and bounded, feasible under normal conditions. In particular, Gigerenzer underlines the role of the context in which decisions are made (**?**), reaching the same conclusions of other authors: "*[...] there is a growing body of evidence that people's preferences depend on the context of choice, defined by the set of options under considerations*" (Shafir *et al.*, 1993, p.21).

The model presented below is based on such proposals. A bounded rational decision making algorithm is adapted to the problem of the consumer. While the original model leaves as exogenous the preferences of the decision maker, here we propose to represent them as depending on the context of consumption: marketing actions by sellers and the environment of other consumers.

Let's consider a product or service defined over a set of different dimensions, so that each single product can be considered as a vector of qualities, all supposed positive. For example, a product may be defined by: cheapness (i.e. the inverse of price); reliability; design; energy consumption; weight; etc. For the time being, let's assume that such dimensions are objectively given, as provided by measurements, engineers' specifications and the like. Later, we can relax this assumption allowing for consumers to have a subjective opinion of the values of each product[2].

The action of a consumer in relation of a given product can be separated in the following parts:

1. Be aware that the product is available;
2. Determine the initial preferences to be used to choose among the different products on offer;
3. Observe available products, and their values.
4. Choose a product, according to the preferences;
5. Revise your preferences according to the experience made with the product and to further information coming from the environment;
6. Buy other products (step 3. and 4.) for replacement and/or upgrading;

[2]We also assume that all characteristics are positive ones, so that a higher value is preferred to a lower one. The unit of measure of quality levels can be arbitrary, since we will compare only values of quality along the same characteristics.

Points 1. and 2. are, of course, events that happens only once in the life of a consumer. Point 3., 4. and 6. are, instead, regular activities repeatedly performed by a consumer. Finally, point 5. will be, for the moment ignored because this section places the focus on the functioning of demand only. Therefore, to limit the dynamics at work, we will "freeze" the supply side, in that producer are not able to modify any aspect of their product on offer. In section 4, when we will describe a model with entry of new producers and possibility of incumbents to modify their strategies and products via R&D, and we will add the possibility for consumers to modify their preferences. Let's now see how each of the remaining points points can be implemented to construct a model for demand.

2.1 Consumers Entry in the Market

The awareness of the availability of a novel product is potentially triggered by many events, like advertisement, professional information gathered in the working environment, general media, etc. However, a huge literature (and, most likely, our own personal experience) shows that the most effective means of diffusion of the awareness that a novel product exists is the direct experience reported by friends, relatives, colleagues, and, in general people with whom we have some sort of proximity.

Accepting this premise, the way to model the diffusion of a novel product is quite straightforward. Let's assume to have a universe of potential consumers distributed over a network of connections. Each potential consumer is linked to some other consumers. That is, agent i has n "friends", each of which shares with i only some of his friends. Let's imagine that one single agent buys a novel product, and starts showing it around to her friends. In short time, all her friends will appreciate the product and will become themselves consumers of the novel product. These "first tier" entrants will be able to present the new product only to *n-1* friends at most, since one of their friend is the one who actually introduce the product to them in the first place (may be less than n-1 because of some of them may be friends of the same first user). Assuming that from the first observation of a friend having a never-seen-before product to the actual purchase passes some time, we have that the number of consumers entering in the market (i.e. making their first purchase), has the usual s-shape form, shown in figure 1.

Let's see now how to represent the activities of first-time consumers of a novel product.

2.2 Forming first-time Preferences

Before discussing how preferences are formed, we need to agree on what preferences actually are. Preferences serve as a general criterion to determine which option is better among many on offer[3]. This problem is not trivial only if we assume that the options can be compared along many dimensions, as it is our case, where preferences must be used to choose one of the products defined over many dimensions. Assuming that, given one single dimension, the decision maker can easily determine what is better, problems arise

[3] We neglect aspects of decision related to differences in incomes and uses which may exogenously segment the universe of consumers. That is, we assume to consider only a set of products that may serve the needs of consumers in similar ways, and that are all accessible.

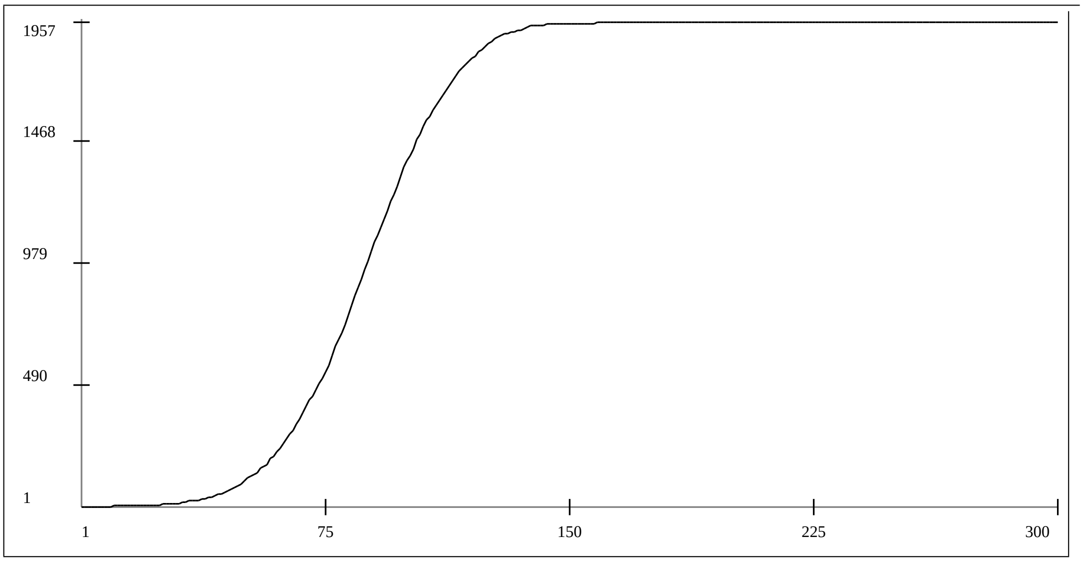

Figure 1: Number of consumers who made their first purchase of a novel product through time. The universe is made by 1957, which is given by the number of friends equal to 6, with no shared friends. Therefore, the only initial consumer brings 6 new consumers in the market, each of which bring 5 new consumers, who bring 4 new ones, etc.

when we have trade-offs, with one of the options dominating the others in respect of a dimension, but dominated on others.

In order to keep preferences separated from the actual decision mechanism implemented, the most general way representation is therefore a ranking of importance of the dimensions defining the product, say from the most important to the least important. Note that this definition of preferences does not imply one single method to actually perform the choice, but only states that, say, one consumer may consider the price more important than quality, and another the reverse[4].

Having agreed on what preferences are, in order to determine how they are formed we need to ask: How do people, who never managed a product, bar from one example in the hands of a friend, actually create their preferences? It is likely that there is a huge variety of examples, from people imitating others, or having meta-preferences of some sort (e.g. cultural or social predisposition, etc.). However, for our purposes we need to ignore realistic, but idiosyncratic mechanism, and concentrate on the effects on preferences from purely market forces, since these are the boundaries of our research. This, of course, does not mean to negate that other social forces influence preferences, but it is an obvious application of the *ceteris paribus* principle.

We can assume that when preferences are formed there are two forces at work. Firstly, first-time consumers of a product try to gather information relevant for preferences. Secondly, they try to obviate some probable inconsistencies. We can safely assume that the

[4]The separation of preferences from the decision mechanism is not in contrast with the huge evidence showing that preferences are constructed, and not just revealed, at decision time. As we will see, the formation and application of preferences will actually take place at the same time, and serves only as a logical distinction.

major source of information about the values of product comes directly or indirectly from the producers' themselves. And, to resolve conflicting messages, they rely on what other consumers do. Let's see more in detail the reasons for this approach.

Supporting the hypothesis that producers are the main source of information, there is the huge evidence of the resources spent in advertisement campaigns, and marketing in general. In effect, such activities are ways to push people to adopt a given set of preferences promoting one product, in respect of another undermining it. For example, making an ad spelling "our product is the cheapest" is meant to say that the cost is the most relevant aspect, implicitly suggesting that all the products on offer are equivalent on the other dimensions, and that the consumer should decide on the base of the price. On the opposite, a marketing campaign based, say, on the reliability of the product try to push through the consumers the idea that durability is worth a somewhat higher price[5].

Competing producers are likely to press through consumers reciprocally inconsistent messages: should I care more about prices, as suggested by one producer, or durability, as pressed by another producer? As an obvious method to solve these inconsistencies we can imagine that actual consumers (or the majority of them) applies a simple rule: trust the message that is more trusted by other fellow consumers. In other terms, weight the messages with the relative importance of producers determined by their market shares. Such method has been discussed by Smallwood and Conlisk (1979), who show under which circumstances imitating the majority behaviour actually identifies the best product. In our case there is no optimal choice to be made, but the same mechanism is applied to implement the management of information on products by consumers.

In conclusion, preferences are defined as a rank of importance of the dimensions over which the product is defined. In order to determine this ranking, first time consumers consider the marketing messages from producers applied to push up a given dimension, weighted with their market share. For example, if only two products are on offer, one marketed as cheap and another as reliable, the preferences of consumers will favor cheapness or reliability on the basis of how many consumers prefer one or the other criterion.

2.3 Reading Products' Values

In order to make an informed choice among different product, necessarily one has to be able to see in what they differ. That is, in our terminology, to read the values of products along different dimensions. As a metaphorical reference, consider this activity as going around visiting shops asking for the price and qualities of alternative offerings of a given product. For the large majority of modern consumers' product, this information can be obtained with varying degrees of accuracy, but it is never certain nor trivial to gather. Firstly, because sellers can, for example, apply different prices. Secondly, because the objective characteristics of products not always match perfectly with the subjective perception and actual experience of all consumers.

[5]We are clearly ignoring the presence of exogenous market segments in demand, with consumers having different needs. Of course, adding this aspect is perfectly compatible with our approach, but would create unnecessary complications.

To model this effect we can simply use a random function, so that when a consumer investigates what is the value of a certain product in respect of a certain dimension, the value considered is a draw from a normal random function. The mean value of the random function is the "true" value, common to all consumers, while the variance is supposed to vary across consumers according to their "experience": first time buyers will have a larger variance, assuming they are scarcely informed and can, therefore, make larger errors; consumers who buy a replacement or an upgrade of the product after many previous purchases will instead have much smaller variance, assuming they can be quite precise in estimating the correct values defining a product properties.

In other terms, we assume buyers to learn how to judge products by using the product itself. As we will see, we can make two different assumptions: in a first case we can assume that buyers eventually become perfectly able to judge products' values (i.e. reach 0 variance in their random distortion); in a second case we can assume that buyers do never reach such perfect capacity, and remain even in the long term with some degree of uncertainty over the correct value of product, though this degree become smaller in time. This second assumption reflects the nature of modern products, where the exact values of the technology, or its components, or production methods are largely ignored by consumers.

2.4 Procedural Choice under Uncertainty

Having determined the nature and formation of preferences, and how product information is gathered, we can finally determine how a purchasing decision is reached. Plenty of cognitive studies have analyzed how people actually behave when needing to choose one option over many with lack of precise information. We present here one of these proposal that is both well grounded on empirical evidence, and shown theoretically to produce correct results as compared to other choice methods (Gigerenzer and Selten, 2001).

A consumer is, as we have seen, endowed with preferences in the form of a ranking of product dimensions. The choice mechanism proposed takes the highest ranking dimension (the most important), and uses it to compare all products. The products that perform too poorly on this dimension are discarded, while the others are maintained as possible choices. Now the choice is restricted to all product that score sufficiently on the top dimension. The consumer then uses the second ranking dimension, again removing from the restricted set all products that do not score high enough. Applying the same rule in sequence using the dimension according to their ranking ends up with two possibilities: only one product remains in the choice set; or more than one. In the first case, the choice is obviously the only produce that survived the selection via preferences. In the second case, we make a random choice, with equal probabilities, of all the products in the remaining set.

The method is clearly a non-compensatory one: that is, one product cannot compensate its weaknesses in some dimensions using strengths in other dimensions. However, as already said, the effects of this choice algorithm is strikingly similar to what people actually do, as it is possible to investigate with experiments. And, surprisingly, it is shown to be a very good algorithm to solve problems with scarce information, providing similar, or even better, results than sophisticated statistical inductive algorithms.

For our purposes we can enjoy a degree of freedom. What is the threshold to determine the products “sufficiently” good to not be removed from the choice set? As we will see, we can test several options in a continuous range. One extreme case is that consumers have no tolerance, so that only the best scoring products on the dimensions used remain as valid alternatives. In this case, only products having maximum *and* identical values on the dimension used survive the selection with one single dimension. Otherwise, we can determine a range of values, as a percentage from the top score, that the consumer consider equivalent to the maximum value. In this case, we have in general more than one product surviving the selection with one dimension, and several selection cycles, using less and less important dimensions, are necessary in order to reach a final decision.

2.5 Testing the Demand Model

We have now enough elements to build a demand model for a novel product: timing of entry; preferences formation; and decisional algorithm. The goal of the exercises presented below is to investigate the aggregate behaviour of a group of such consumers. Since the basic elements considered represent assumptions embedded in the code of a simulation model, we will be able to trace back the motivations of interesting configurations to the elementary components of the model, that is the consumers' behaviour.

To perform this exercise we need to set up an environment where the consumers act as described, that is, a supply side of a market. In order to facilitate our research of motivations we neglect the realism of the environment, and instead we choose an environment in which we can easily control the effects on the object of our study. Therefore, the firms in our supply side will be extremely simplified, so that it will be easy to determine the supply-motivated causes of the model results from the demand-motivated ones.

2.6 A stylized supply side

We consider in our model a fixed number of firms, each offering a single product. As we have seen, a product is defined as a vector of values, one for each dimension. We have also seen that the preferences are determined by the marketing pressure from firms that try to press consumers to adopt some of the dimension as a main criterion of decision. We represent this as another vector with the same number of elements as the product dimensions. This second vector, that we can refer to as the marketing vector, expresses in each element the importance that a firm would like to give to that dimension.

For the time being, we refrain from allowing firms to modify either the products' values or the marketing ones, so that we can test how certain configurations determine the demand behaviour.

2.7 A Model for Monopoly

Let's begin with an exercise considering the most obvious case: that of a “natural” monopoly, and let's see if our demand is able to represent the expected configuration.

We implement a set of firms whose products are all identical, but one firm that is superior in all the characteristic. We also assume that consumer can reach the perfect capacity of reading the true products' quality values, though making some errors in the beginning of their consumers' career.

Of course, our expected result is that all consumers converge in choosing the dominating product. In fact, this is what happens, as it is shown in figure 2.

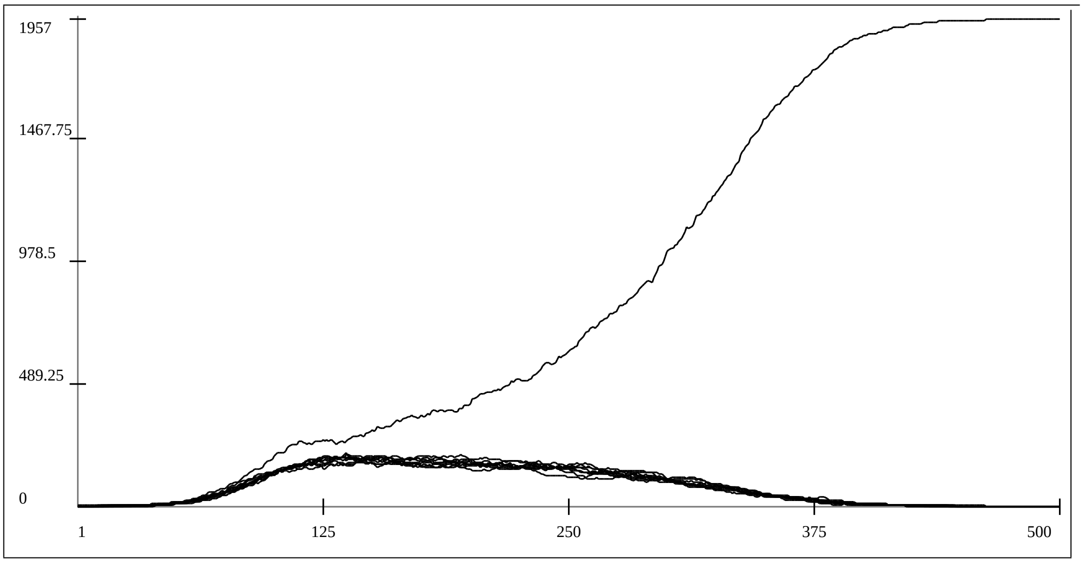

Figure 2: Number of consumers that choose a product through time, having the possibility to make a perfect reading of products' value after a period of learning. The final "monopolist" is the dominating product.

The final configuration of this exercise is quite obvious: given the decisional algorithm endowed in the artificial consumers, they are able to recognize the monopolist, when the reach the perfect capacity to read the true product's values. The pattern leading to this final outcome, however, reminds us that, until such perfect capacity is obtained by all the consumers, even dominated products enjoy a positive market share. This is due to the fact that, being biased by inexperience, our consumers may consider as superior one of the actually dominated product.

In other exercises it is possible to show that if the consumers maintain a certain degree of bias, that is error in judging products' values, than even inferior products enjoy indefinitely positive, though small, market shares. That is, if we limit the capacity of consumers to perfectly read quality levels, all products may find some consumers who could choose their product, despite their inferiority. Such phenomenon is not only a theoretical possibility, but has been noted also in the marketing literature, where it is suggested that some producers may purposefully aim at targeting the "minimum market share" formed by the most distracted consumers (Karnani, 1983).

However obvious, our model reminds us that the possibility of reading correctly the qualities of product is crucial in order to determine the configuration of a market. In general, the model trivially shows that more informed are consumers, the better is for the superior product, *as long that the superiority is actually important.*

2.8 Apparent Superiority: Ignorance is Good for Bill G.

Figure 3 reports the results of the same model above where we made a small change to the configuration used in the previous exercise: we used the tolerance parameter, indicating that the superiority of the "monopolist" is actual not that important, so that all products should be judged, by experienced consumers, as identical, and therefore having identical the market shares.

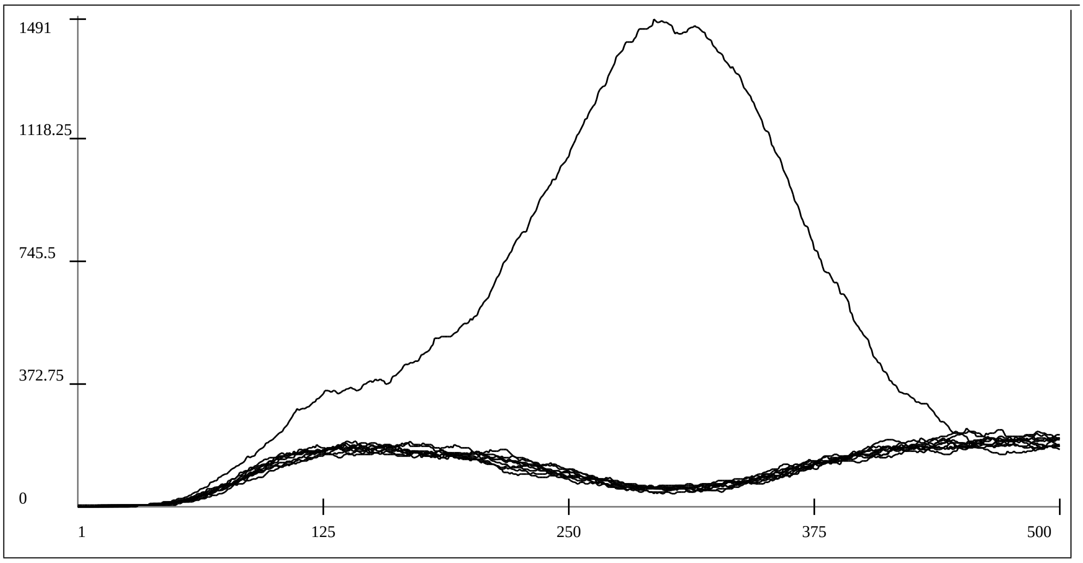

Figure 3: Number of consumers that choose a product through time, having the possibility to make a perfect reading of products' value after a period of learning. The tolerance level is such that the superiority of the "monopolist" does not actually matter to consumers' perfectly informed.

The final configuration is, obviously, such that consumers choose randomly among all the products. The unexpected thing is the pattern followed by the market to reach such configuration. Why does the (apparently) superior product dominates the market for long time, if its superiority is not relevant? The reason lies in the (temporary) incapacity of consumers of perfectly read products' quality values. As long as consumer do some mistakes, it is likely that the (truly) small advantage that is considered not relevant, appears as a bigger advantage, such that the slightly superior product appears as largely superior. Under this circumstances, the producer of the superior product gains from the ignorance of consumers: this firm can correctly claim a superiority, but only perfectly informed consumers can understand that it is not so relevant.

This exercise, though limited in its realism by all the simplifying assumptions used, suggest one more reasons for some of the dilemmas considered by empirical economists. As in the case of the products by Microsoft, a large consensus agrees that the dominance of this firm is not motivated by an actual technical superiority. Our exercise suggests that, along network externalities, switching costs, and other factor, we may also consider the technical incapacity of consumers to determine the exact technical validity of Microsoft products.

2.9 Different Market Segmentations

The previous exercises ignored the potential trade-off implicit in the use of multiple product dimensions. As a first exercise to test this aspect we may produce a trivial market segmentation by allowing competing firms to be "specialised" in different dimensions of the product. Assuming that buyers can reach perfect capacity to read products' quality values, we obtain that each firm serves buyers that value mostly the dimension in which they are specialised. Therefore, we could obtain as many market segments as many dimensions are assumed for the product, that is, having a "niche" for each dimension[6]. However, the model can show a more elaborated market segmentation, which can suggest interesting dynamics on real markets.

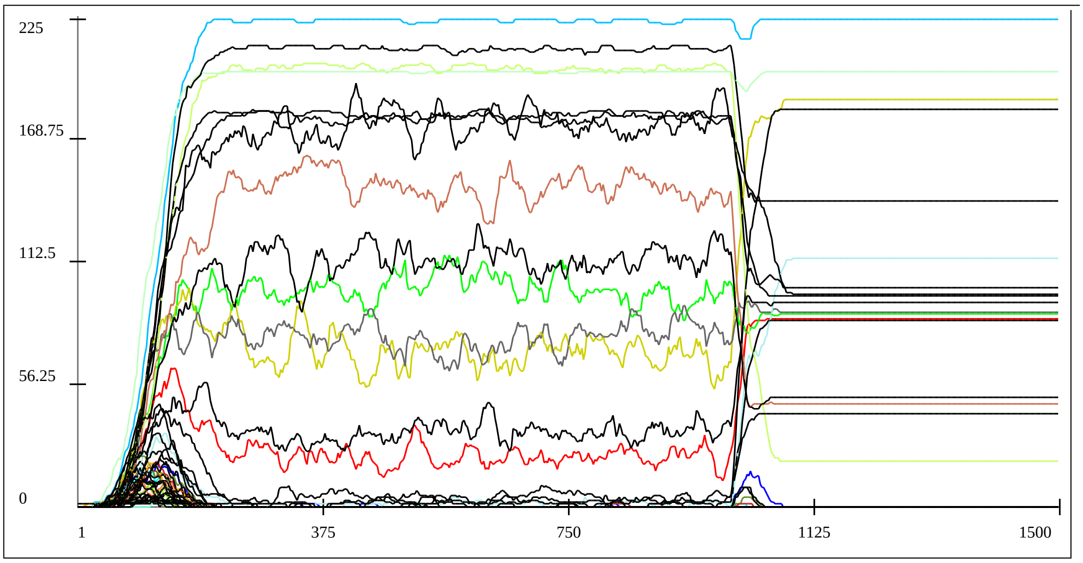

Figure 4: Number of consumers that choose a product through time, having the possibility to make a errors when reading products' value. The tolerance level 0 up to time 1000 and instead consider equivalent quality levels within 3% of the best level afterward.

In figure 4 we report the results produced when the product is defined over 10 dimensions and using 50 firms. Both quality values and marketing stratetegies' values are drawn randomly. Consumers are not allowed to reach perfect capacity to read products quality values, so that they continue for ever to make a small errors. Until time step 1000, consumers adopt a "zero tolerance" approach. That is, they choose the best product according to their (possibly erroneous) observations. Therefore, they keep on using only the very first dimension as the only criterion used to judge products. However, the market allows for rather high market shares for quite a number of firms, well above 10, since more than one firm can serve the same "niche" of consumers sharing the same most important dimension. The type of market segmentation appearing until this point is a segmentation

[6]Such trivial results are not shown to save space.

built on consumers' error: if consumers could correctly read products' quality levels, they would eliminate all but one firm from each niche. In this set up, each firm may well serve several niches, since having relatively high quality levels (not necessarily the highest) on many dimensions can appeal to many customers who (possibly mistakenly) look for the best product.

This type of segmentation can afford more firms than product dimensions. Moreover, each producer can serve several different "niches". Each defined by consumers with (erroneously) "sharp" tastes. Let's see how the configuration would be completely different with an apparently small change.

In time step 1000 we change the tolerance level[7] , from no tolerance up to a small value. This means that now consumers consider equivalent quality levels even less then maximum, provided they are high enough. Therefore, consumers start using also the second, third and so on ranking dimensions, as defined in their preferences. To further clarify the effect of tolerance, we also bring to 0 the error of consumers, so that they always correctly choose the best product according to their currently used dimension (provided that no other product score similarly).

With this modification the configuration changes dramatically, though consumers' preferences and products' values are unchanged. Firstly, and more obviously, the series stop the random oscillation, since there are no more random factors in their choices. Secondly, we have still a large number of "niches", that is, firms making positive shares, higher than the number of dimensions. However, the nature of these niches is very different from the those before the change. In fact, without tolerance consumers used only one dimension to make their choice. Several consumes using the same dimension could end up choosing different products because of their random errors in identifying the best product in respect of the chosen dimension. Instead, without errors and using a small tolerance margin, consumers operate a selection of available products. They end up making a choice using the secondary dimensions in their preferences. It is well possible that different consumers make their final choice using the same dimension but choose different products (without errors) this can happen because they may have operated a different selection before making their choice. In practice, each consumer is a niche on itself, having a pattern of selection of products that make it unique.

The different types of niches are very important in order, for example, to determine the best strategy for firms. That is, suppose a firm has the possibility to improve its product on only one dimension, which one should it choose? If consumers are segmented with the first type of consumers (no tolerance, consumers ignore all characteristics but one), the firm should aim at be the best on the characteristic that the highest number of consumer are mostly concerned. Instead, if the market is of the second type (consumers observe many characteristics), the choice is much more difficult: it could be a waste to improve the most popular characteristics, if the consumers have a large tolerance. Rather, the

[7]The switch of parameters during a simulation run is used to ensure that we do not modify consumers' preferences. As mentioned above, in fact, the preferences are defined once and for all at the time the consumer enters the market, making the first-time purchase. Modifying a parameter in the middle of a simulation we ensure that the set of consumers' maintain the same preferences.

firm should select the characteristic, that, though valued as less important by consumer, is determinant in their actual choices.

The demand model described in this chapter does not claim to represent any real world counterpart, since it is highly incomplete (lacking, for example, a dynamics of the supply side and considerations among income stratification of consumers. However, exactly because of these simplifications it allows to study in detail the demand behaviour, showing how this can generate several realistic market configurations. The results obtained can be used in the study of real world markets by identifying the mechanisms governing the demand side of markets, with potentially fruitful applications for both interpretative or normative purposes.

The next step in developing an organic representation of market dynamics consists in developing a model of supply. In the next section we will focus on one specific aspect of supply: that of technological innovations affecting the products' characteristics.

3 Modelling Technological Innovation

Many studies on technological innovation focus on the aspect of process innovation by focusing on the variable "productivity" as capturing the results of the innovative activities. One of the motivations is that trying to model product innovations (that is, innovations affecting the qualities of products or services, and not only their production costs) one is forced to consider competition among heterogeneous products, notoriously a difficult task. In the previous chapter we have discussed competition among heterogenous products, so we are free now to analyse technological innovation affecting both process and product innovation[8].

A company involved in a technological research is by definition involved in an uncertain task, given the difficulty in predicting costs and yields of research. However, there are always different options concerning the strategy to use for carrying on a research project. For example, one can try to re-design completely the product in question, or trying only to improve some aspect of the current product. Though uncertainty reigns, one can safely predict that a complete re-design will cost more (and take longer) than a minor adjustment. In this chapter we present a formal model that captures the central questions concerning the economic aspects of technological innovation.

3.1 Complexity in Problem Solving

We can abstract the problem of technological research, considering it as an example of a generalized search of solutions for a complex problem. The basic definition problem solving is: having an evaluation criterion and a set of potential solutions, find the solution with the best score in respect of the chosen criterion. How difficult is to solve a problem is partly a subjective matter, but we can express it in a formal way by representing the

[8]Morever, treating products as vectors of qualities we can include also markets for services, which are already the most important sector in modern economies (and still growing), and which can hardly be considered homogeneous.

difficulty of a problem (i.e. in finding the solution) against a benchmark solution strategy. Such benchmark is the most simple solution strategy: try all potential solutions and pick the best one. We can define a problem as the most complex if this is the only solving strategy. A problem is simpler than another if we can devise solving strategies that permit to test a smaller portion of the set of solutions. In general, we can assign a measure of complexity of a problem as the number of potential solutions to test (or the percentage over the total solutions), in order to reach with certainty the best solution[9].

To explain this definition of complexity we can use the example proposed by Herbert Simon: the "ticking safe". Suppose you have a safe with N wheels, each having A positions. You can open the safe only of you set all wheels in one specific position. If you don't have the combination and the safe is well-built, you should test all the A^N possible combinations of the wheels positions to be sure to open the safe. If A=100 and N=3 it means 1 million combinations. Spending 1 second to adjust the wheels and testing the door, you would spend around 11 days to be sure of opening the door. However, if the safe is poorly-built you may discover that the wheels produce a "tick" when set in the correct position, disregarding the position of the other wheels. Now you have reduced the problem to the analysis of 3x100 positions, or 5 minutes: quite an easy robbery.

Why is the ticking safe a simpler problem than the non-ticking one? The reason lies in the interdependencies among the components of the problem. The non-ticking safe has highly interdependent wheels. The "criterion" we use in this problem to test the solution has only two conditions: door close/door open. A variation of the score is appreciated only if *all* the wheels, at the same time, are set in the correct position. Instead, the ticking safe provides an additional score value: the tick. Such event takes place independently for each wheel, and therefore we can decompose the overall problem (open the safe) in many *independent* sub-problems (make the wheels tick).

3.2 NK Fitness Landscapes

The study of the complexity is therefore the study of the interactions among modules composing sub-problems, whose solutions provide the overall solution. In the last decade a number of researchers have used various formalisms to study the properties of complex systems. One of the most popular is the NK system, by Stuart Kauffman (Kauffman, 1993). In this model, a problem is represented as a set binary strings of a given length N, which gives a set of 2^N possible solutions. To each string it is assigned a "score", called fitness because the model was originally meant to represent the evolution of biological organisms, with the strings representing the genetic code of competing animals. The fitness of a string is determined as the average of the contribution from each element of the string. The crucial aspect is the way such individual contributions are determined. If these contributions differ depending only on the state of the element concerned, then the system represents a simple problem. Figure 5 shows a simple NK system with N=3. The table on the left shows the fitness of each of the 8 strings computed as the average of the

[9] Without loss of generality, given our goals, we can ignore the problems with multiple best solutions.

contributions. Notice that the contributions of each element depend only on the state of the element itself. The "cube" on the right represents the paths available when exploring the solutions with the strategy of testing neighboring strings where only one element has a different state. If one moves always "uphill" going on strings with higher fitness, one finds always the way to the global maximum. In other terms, the problem is fully modular, or decomposable. One can solve the problems of each element independently from the state of the other modules, and, in so doing, finding the best solution for the overall problem.

	c1	c2	c3	f
000:	0.2	0.6	0.8	0.53
001:	0.2	0.6	0.5	0.43
010:	0.2	0.9	0.8	0.63
011:	0.2	0.9	0.5	0.53
100:	0.7	0.6	0.8	0.70
101:	0.7	0.6	0.5	0.60
110:	0.7	0.9	0.8	0.80
111:	0.7	0.9	0.5	0.70

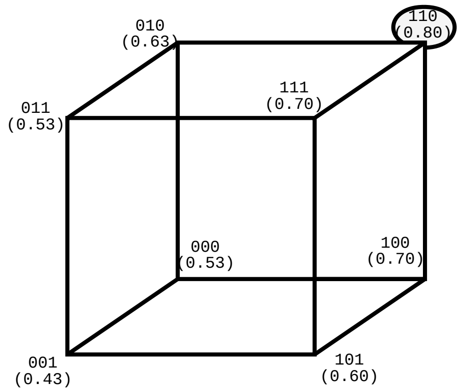

Figure 5: Fitness landscape without interdependencies. The contribution of each element does not depend on other elements. A unique optimum can be reached by any starting point.

Alternatively, it is possible to build a NK system with different contributions for each combination of the element concerned and those of interdependent elements. Figure 6 shows another NK model where the contributions of each element differs anytime one of the elements changes. The result is that a "walk" in the space of the solutions does not guarantee to provide the best solution. That is, the solution to each sub-problem depends on the state of the other elements. One cannot be sure that the solutions of the sub-problems generate the solution for the whole problem. In fact, this will depend on the initial point. The reason lies in the interdependencies: the solution of a sub-problem depends on the state of the other sub-problems.

The example above uses the maximum interdependency, where each element depends on every other element. In general, one can build systems where each element depends only on a number of other element. That is, the contributions of each element depend on the state of only some of the other elements. Indicating with K the number of interdependencies between the elements of the string, we have a simple problem if $K = 0$, a maximally complex problem when $K = N - 1$, and intermediate complexity when $0 < K < N - 1$.

3.3 Solving Strategies

The NK systems (called also *fitness landscapes*) provide a very useful formalism. In fact, they represent a problem in terms of a finite set of solutions, though the dimensions of this set can be easily expanded to be infinite for all practical uses. For example, for $N = 50$ the number of strings runs in the millions of billions: even the fastest computer would

	c1	c2	c3	f
000:	0.6	0.3	0.5	0.47
001:	0.1	0.5	0.9	0.50
010:	0.7	0.9	0.5	0.70
011:	0.3	0.5	0.8	0.53
100:	0.9	0.9	0.7	0.83
101:	0.7	0.2	0.3	0.40
110:	0.6	0.7	0.6	0.63
111:	0.4	0.8	0.1	0.43

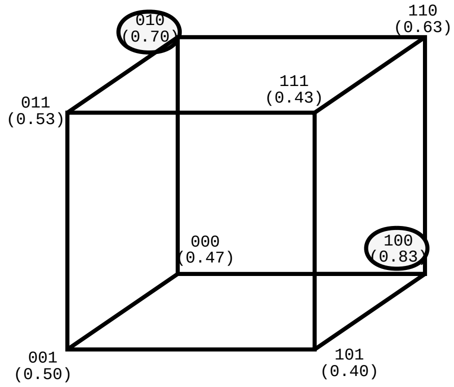

Figure 6: Fitness landscape with interdependencies. The contribution of each element depends on other elements. The solution depends on the starting point.

spend eons trying to elaborate such numbers. However, with few technical tricks it is possible to build programs that implement a few thousands, or hundreds of thousands, of strings with the same properties that they would have in a full fitness landscape of such dimensions. Notice also that the NK models do not depend on the type of score, or fitness; all you need to ensure is that the contributions vary when the element's state, or one of the linked elements, changes. In fact, generally the NK are built using random values for the contributions. This ensures that the results obtained do not depend on a particular pattern of the fitness levels, but only on the interdependency structure.

The NK system is used for two purposes. On the one hand we can investigate complex systems as such, trying to find detailed measures of complexity. On the other hand we can test different solution strategies to match such complexity, evaluating their capacity to match different degrees of complexity (Frenken *et al.*, 1999).

For example, let's consider a problem with $N = 40$ where the elements of the strings (to be considered as the basic building blocks) are each linked to other 4 elements ($K = 4$). We can study the efficiency of three different strategies. All the strategies apply the same routine: measure the fitness of a string; modify the current string to produce a new one; move to the new string if it provides a higher fitness, otherwise remain on the previous one. The difference between the strategies consists in the method of production of a new string starting from the current one. A first strategy, that we can label as "overmodular", changes the state of one single bit in the current string. A second one, call it "optimal", modifies from 1 to 5 elements. The last one, call it "integral", mutates from 1 to 10 elements to generate a new string. We run a simulation where each strategy is represented by 100 artificial agents, and we measure the average performance for each population of agents adopting the same strategy.

After few steps we observe that the agents using the overmodular strategy get stuck in local optima. That is, they have found strings that have all neighboring strings (that differ only for one bit) with a lower fitness, and therefore these agents cannot further improve their fitness. Instead, all agents using the optimal strategy are able to reach the string with the highest fitness after less than 10000 attempts (about 0.00000001% of the total

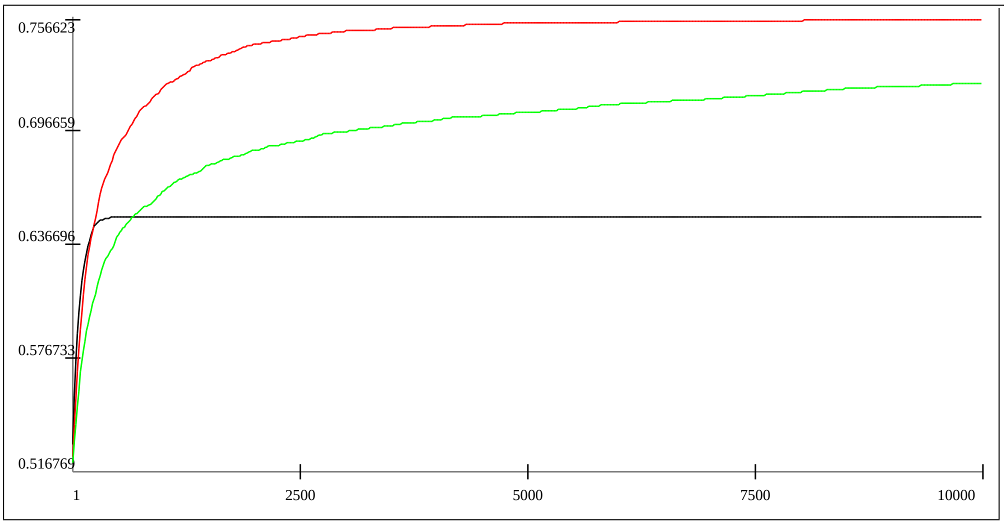

Figure 7: Average fitness of three populations using an overmodular strategy (black line), and optimal one (red line) and an integral strategy (green line). Data produced with N=40, K=4, with 100 agents per population.

number of strings). Agents using integral strategies keep on improving their fitness, that is, finding strings with higher fitness, but are quite slow at reaching the global maximum.

The overmodular agents adopts a strategy that ignores some of the interdependencies among the elements of the strings. Therefore, they cannot find a monotonous path upwards from a string and the highest fitness one (i.e. the *solution* of the problem represented by the landscape). They either need to make steps in strings with lower fitness, or must accept to get stuck in local peaks. Conversely integral agents adopts a strategy that takes into account non-existing interdependencies. The result is that they do reach the global optimum, eventually, but the journey is extremely long. The intermediate strategy is optimal because it is the fastest strategy that provides with certainty the global maximum. This is because this strategy takes into account all and only the interdependencies actually existing in the problem.

In other terms, for each problem it is possible to provide an objective measure of complexity as the minimal number of interdependencies one has to take into account in order to ensure to find a global maximum. Getting right the interdependency structure of the problem (i.e. what depends on what) is crucial to obtain the optimal solution. However, one may question the need of absolute optimality, balancing the trade-off between costs of having a sub-optimal solution and the gain of spending a fraction of the search resources requested by the optimality. Indeed, it has been shown that overmodular strategies accounting for some (but not all) interdependencies provide in general local optima very close to the global optimum.

3.4 Technological Innovation as a Complex Problem

Ignoring interdependencies implies the decomposition of a large problem into smaller sub-problems. Choosing a finer decomposition speeds up the research of a solution because of two reasons. On the one hand, smaller sub-problems are easier to solve. On the the hand, the independence of each sub-problems allows the use of parallelism in the research of solution.

In real world problems, such as those faced by R&D projects, researchers face two kinds of difficulties. The obvious one is to find solutions for the sub-problems. The second, and most important, is to determine the boundaries of the sub-problems. That is, to determine the decomposition of a the main problem into sub-problems. These decisions are quite difficult because one cannot be fully aware of the interdependencies structure of a problem until the problem is solved. Therefore, a good deal of guessing is required in order to determine the boundaries of problems. The observation of the technical development of a product shows that a firm, or a market, agrees upon an *architecture* setting the overall design of the product. Within this architecture each team of a research lab, or each firm in the market, can safely develop their component with the confidence that their improvements will be compatible with those of the other components. In fact, an architecture is defined by specifying the standards of communication between components. In other terms, an architecture consists in the set of interfaces linking the components.

The problem is that one has to determine the interfaces among the components before knowing with certainty the potential developments of each component. Historically, it has been evident that after some time of components' developments some of the interfaces become bottlenecks because of their wrong specification. When the bottlenecks are too many, then an overhaul of the architecture is called for, and a new generation of the technological development can start again.

We can represent this way of carry on technological innovations is like trying to solve a problem imposing artificially overmodular research strategy. The research within single components will be fast, but will necessarily meet the limits imposed by the excessive modularization, calling for a new definition of the decomposition allowing another wave of components' innovations.

From what we have learned using the NK model, we know that setting a very modular strategy will limit too much the overall performance that can be reached. But it will also provide quickly relatively good results. Conversely, setting an integral strategy allows for very high quality final results, but will costs much more in time of research. In the next chapter we will see what this implies for the economic life of companies.

4 Markets and the Co-evolution of Supply and Demand

Having studied some aspects of demand and supply separately we can now try to put the two sides of the markets together in order to represent a general, and simplified, market, so that we can study which market configuration emerges and the reasons for such events.

Before describing the model, a warning note is required. We do not aim at building a realistic model, but one that can provide us with insights. Real world markets have so many peculiar aspects that, even if we were able to build a good approximation of one of them, we would face the same problems of generalizations and understanding that we have looking at real data. What we will do is to implement some aspect of reality and intentionally neglect many others. When interpreting the results we will discuss sufficient conditions for the emergence of a configuration, though we need to keep in mind that, in many real-world cases, there may be others limiting or reinforcing the ones derived from our model. In other terms we will be able to say "event A can cause the event B by means of the mechanism Π". Whether B can be enhanced or hindered by mechanisms other than Π is left to the analysis of any specify case study.

Concerning demand we can straightly use the model described in section 2. That is, we will represent a model for a market of a new product, where early adopters slowly bring new consumers to buy the product for the first time. Any first-time consumer will initially suffer a random bias in reading the objective values of product characteristics, that will later wane while the consumers become more and more expert on the product.

4.1 Modelling R&D strategies

Concerning the supply side we focus on the role of technological research in bringing product innovation. For this reasons we ignore production decisions, related with physical and financial capital. We assume that all firms are endowed with unlimited production capacity, and that all of them gain a fixed amount on each piece sold which is entirely devoted to R&D. Such limitations of the model will allow us to concentrate on two aspects of firms' behaviour: R&D strategies and marketing.

We want to include in a model the results we had when discussing R&D. We saw that a fitness landscape can be used to represent the technological problem of developing a complex product. We also said that a firm faces the problem of choosing a modularization and performing the research in the resulting sub-problems. As a result we obtained a trade-off between length of duration of a research project and quality of the results guaranteed by the project itself. Rather than implementing a full fitness landscape and the firms acting on it, we put in the model a much simpler structure representing the same properties as the ones derived by a full-fledge model with a NK system.

We have seen in the demand model that a product is composed by a number of product characteristics. We assume that such characteristics corresponds to different modules on which firms can make research in order to improve one or more these aspects[10]. The money spent on R&D produce an improvement in the quality of the product. The richer a firm, because of higher sales, the more it can spend on R&D. However, the relation between R&D resources and quality improvements depends on the complexity of the task, that is, on the interdependency structure of the technological problem.

The interdependency structure is represented by a simple mechanism: the increment in quality in any one characteristic depends, other things being equal, on the quality levels in respect of the other characteristics. That is, the same amount spent on improving a characteristic gains a higher quality level if all the other characteristics are themselves at high quality levels. In the limit, if all the other characteristics score very poorly, the

[10]In practice, it is like saying that a fitness landscape has many fitness functions, so that each string is associated to many values, one for each characteristic.

increase of quality is infinitesimal even for large amounts of R&D money. The underlying idea is that, if you want to improve the speed of the car you produce, it is useless to come out with an engine of a Porsche if the rest of the car is like a Fiat 500.

Firms's technological decisions concerns how to allocate their R&D budget. The model implements this decision as a function of the profile of the customers' preferences as perceived by the firm. In the demand model we have seen that when a customer interviews a firm concerning the possible purchase of a product, it asks for the quality level of the product concerning one characteristic. Therefore firms can keep track of the characteristics that are more important for the customers, and they concentrate their R&D resources on these characteristics, diminishing the share of R&D budget for the least requested characteristics.

We have seen that customers preferences depends also on the marketing strategies of firms. Also these strategies are determined by firms in the same way as the R&D decisions: the more frequently a characteristic is considered by customers, the more important it is for the marketing of firms.

All the decisions of firms, concerning marketing and R&D strategies, are slowly updated through time. That is, at each time step a firm modifies only partially its strategies on the base of the information received in that time step, on the assumption that customers' messages are subject to a good deal of erratic fluctuations. This "stickiness" in the reaction of firms captures also the intuition that a firm needs time in order to modifies its internal structure: to switch from a R&D specialization to another needs time, and cannot be done in one day.

4.2 Modelling the Supply Side

The model consider 10 initial firms. Moreover, entry of novel firms and exit of failing is considered. Exit mechanism is quite obvious: incumbent firms that fail to make a sales for long time are removed from the set of potential sellers. A new firm enters with a probability of 1/3 at each time step, and are initialized as follows. Concerning their internal strategy structure (i.e. allocation of R&D budget and marketing strategy) they are set to "neutral", that is, having no *a priori* bias in favor of some characteristic. Concerning the quality levels for each characteristics of their products the entry mechanism replicates a mechanism of imitation in the industry. At any given time the industry records the best quality level for each characteristic, and the best *imitable* quality level. The imitable quality levels lags behind the highest level. If the highest level remains constant (because no firms improves it), after some time the imitable level reaches the highest one. New entrants draw initial quality levels for the characteristics of their product from random functions topped by the imitable levels. Therefore, if a new firms enters the market only few time step after a large improvement of quality in a characteristic, it cannot reach such high level. Instead, new firms entering after many periods that the best quality level has not been increased can hope to draw values very close to it.

Now we have a model ready to be explored: a demand side faces a supply side made of a varying number of firms, each of which develops endogenously its marketing and

R&D strategy. At industry level (that is, valid for all firms) the technological complexity of the product is defined as having a certain degree of interdependency. Moreover, a process of imitation is allowed for new entrants' only, while incumbents only can introduce innovations (i.e. quality improvements). This rigid (and quite unrealistic) division between the which actor can perform innovation and imitation serves the purpose of making model results more intelligible, while loosing somewhat in realism. We will return on this later.

As mentioned commenting the results of the model containing the demand only, our simulated market represents a market for a novel product. New customers are tempted by fellow customers, and develop through use a knowledge of the product allowing them to make more informed choices through time. The supply side develops endogenously R&D strategies to improve products' qualities in a complex technological space. We can now run the model study the results.

4.3 Simulated Market Dynamics

One of the most typical phenomenon observed in markets for new products is the strong increase in the number of firms selling the new product in the first phase of the product introduction, followed by a rapid fall in the same number, when most of the incumbent firms fails or leave the market (e.g. Klepper, 1996). Figure 8 shows the series of the number of firms and an inverse index of concentration (see figure's comments).

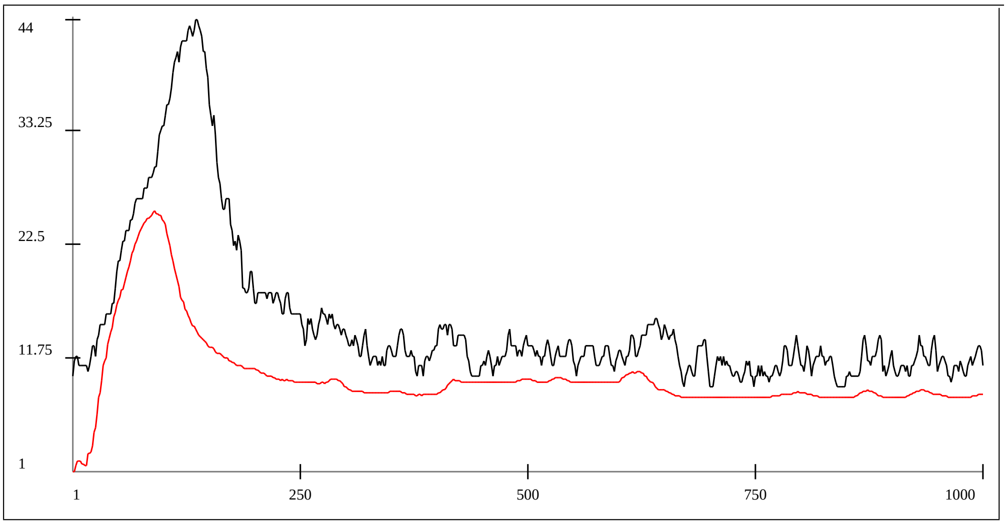

Figure 8: Number of firms (black lines) and inverse Herfindal index (red line) through time. At time 1 only one customer is present in the market, while at time 150 the market reaches saturation (almost 2000 agents). The inverse Herfindal index measures the equivalent number of same-size firms that would provide a concentration equal to the one shown in the market.

The reasons for the shake-out are the same that caused the emergence of a monopoly in figure 2: sharpening judgment capacities of customers allow them to identify the best firms in respect of each characteristic. Until most of customers are poorly informed about the product each firm, however poor is its product, can hope to convince someone that it has the best product. With more expert customers this is impossible, and the market selects one firm in each "niche", formed by customers with the same most important characteristic in their preferences.

This segmentation has an important consequence concerning the nature of R&D strategy of firms. After the shake-out the surviving firms enjoy a quasi-monopoly in their market niche. They can do this because they concentrate their R&D spending on one single characteristic, that defining the niche itself. The firm that is initially slightly bigger than the others can exploit the virtuous cycle of having more R&D money, improving its product more than competitors, and thus selling more generating even more R&D money.

However, in the long run this behaviour (the very reason that brought them to become the leaders of a niche) becomes their fatal error. Concentrating all the innovative efforts on one single characteristic they cannot improve the quality of the others, and therefore find more and more difficulties to further improve the quality of their product. Their short term strategy is necessary to become the market leaders, but limits the long-term capacity to further improve the product.

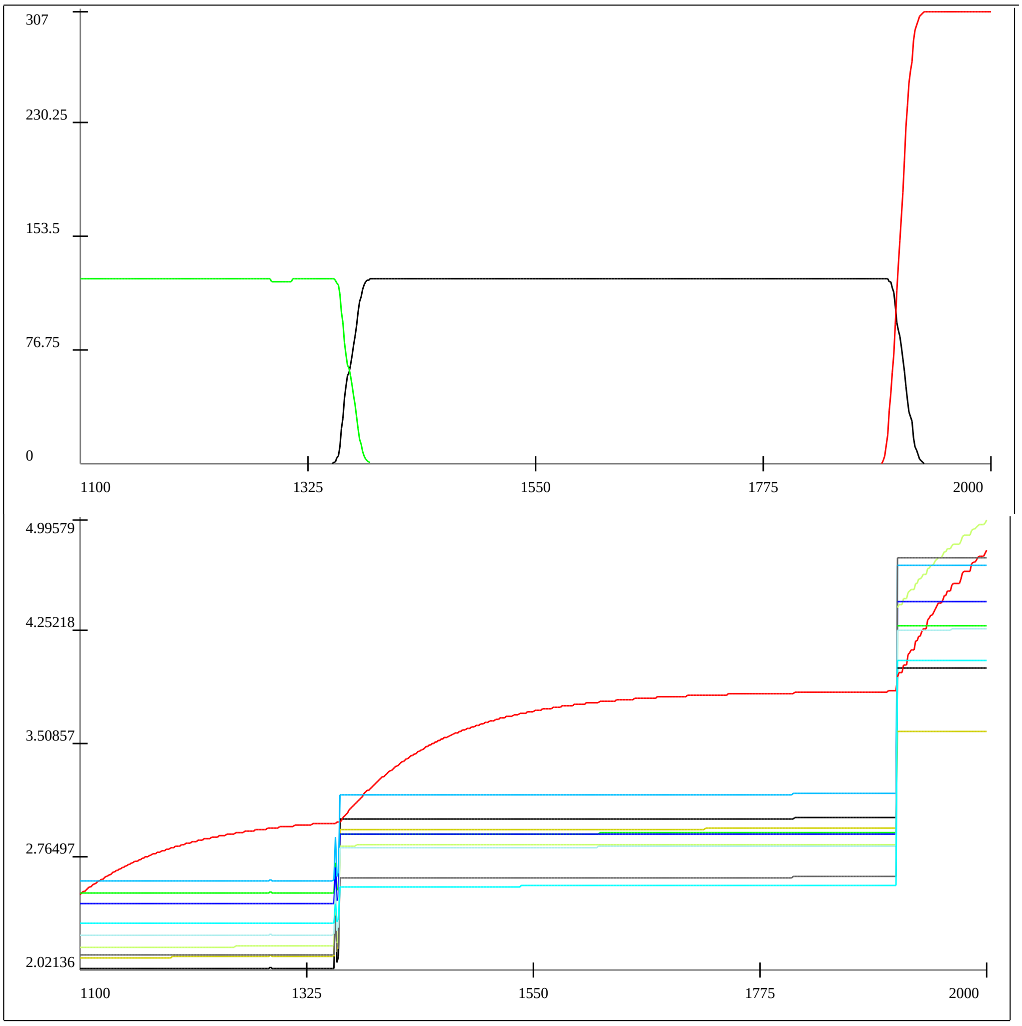

Figure 9: Example of niche dynamics. The top graph shows the number of consumers served three firms dominating the niche. The bottom graph shows the average quality levels of the firms serving the niche.

The end of an incumbent firm (leader in a niche) happens as soon as the high quality levels of its product becomes easy to imitate. New firms can therefore offer a product that can imitate relatively high quality levels from all the market leaders of the market, that is, on all the characteristics. Hence, in a single niche, a new entrant has similar quality level for the main characteristic in the niche, and higher quality levels than the incumbent for the other characteristics. Though the new entrant has initially lower sales, its smaller R&D spending is much more effective than the incumbent, and therefore can quickly overtake the dying incumbent. Again, among the many entrant with similar qualities in their product (because they enter in the same period, imitating the same quality levels), only the ones that focus on one characteristic have the hope of surviving the competition, assuming by necessity the narrow R&D approach of the incumbents, and that dooms them to be removed in the future.

Figure 9 shows the dynamics of sales and average product qualities in a niche. Over the period considered (from time 1100 to time 2000) three firms serve the niche. The last firm captures also the customers of another niche, so that its sales are sensibly higher than the other two. The quality levels of the characteristics not relevant for the niche remain constant during the domination of a firm, because all the R&D is devoted to improve only the characteristic defining the niche. Unless, of course, a firm occupies two niches, like the third firm, in which case there are two characteristics improving their qualities. The improving qualities, however, are capped by the maximum improvement that can be reached given the interdependencies. Though the amount of R&D resources increases or remain constant, the resulting quality improvements diminish because of lack of improvements on the other characteristics. The lack of improvements allow start-ups to catch up with the quality levels of incumbents, so that when higher quality levels become easy to imitate the incumbent looses to a new entrant, which shows higher quality levels on each characteristic.

The results presented suggest technological innovations follow a pattern with sequences of periods with relatively stable "dominant designs", where products undergo incremental innovations, followed by short and chaotic periods of adjustments when a radical innovation is introduced. This vision is compatible with the evidence from the literature (see, e.g., Abernathy and Utterback, 1978). Of course, the model does not claim to provide a unique explanation for this phenomenon. However, it suggests that a crucial role is played by the interplay between the competitive pressure of economic conditions and the technological landscape unfolded by the research carried on by the economic actors. Moreover, the model bases the introduction of radical innovations on the possibility of merging technologies from different lines of products. This consideration brings to the mind the emergence of mobile phones out of the innovations in computing and communication, or the possible advancements in the unification of mobile phones and PDA's in a single product.

5 Conclusions

The aim of this work is to discuss the effects of the interactions between demand and supply in shaping the patterns of market dynamics. Concerning the demand side of the markets the analysis focuses on the endogenous portions of preferences influenced by sellers' marketing strategies. For the supply side the analysis consider the influence of competition on the nature of R&D carried on by suppliers.

The demand is represented by a set of bounded rational consumers, who enter in the market, generate their preferences, and regularly purchase one of the available products. A partial model of demand allows the analysis of the effects of demand-only dynamics in shaping market configurations.

Supply is composed by producers who can influence consumers' preferences by means of marketing, and spend part of their revenues in increasing products' qualities through R&D. The relation between R&D spending and quality of innovations is based on the literature concerning complex problem solving.

We have discussed a partial model where only demand is active, and we were able to individuate the contribution of demand to the formation of specific market configurations. We then review a strand of the literature on complex problem solving to individuate the likely solutions a that different research strategies generate as compared to problems of different complexity.

The final model for market dynamics includes demand and supply that endogenously determine their properties as resulting by the repeated trade. The model is purposefully lacking many features of real-world markets (i.e. production processes are not considered), so that we can easily determine the relation linking observed phenomena to the assumed behaviour of the agents.

The simulation results show several well-known phenomena, observed in many markets. For example, the increasing number of firms serving a market in the initial phase, and the subsequent shakeout drastically reducing the number of incumbents. Such event is motivated, in the present model, by the initially low capacity of consumers to recognise technically superior products. As soon as this capacity is increased, producers of dominated products fail to draw consumers.

Another phenomenon observed in the results is the discontinuous pattern of innovations; long periods of incremental changes are suddenly interrupted by a radical change, which deeply modifies the products on the market. This phenomenon is motivated by the R&D strategies of firms. In the short term they must focus their research in order to obtain quickly improvements, as required by a tough competition. However, this very choice condemn doom the possibility of providing further quality increments in the long term, since this can be obtained only by broadening the landscape of research. When the technological limits of a focused R&D strategy are met, a radical innovation suddenly becomes available. Such innovation can in fact collects all innovations from different technological specializations, and therefore breaks the limits of the old approach. A new cycle starts again, with competition forcing innovators to repeat the same "wrong" short-term

R&D strategy.

Though the present state of the model is still too abstract and lacking too many features, it is possible to envision several possible applications to be used for practical purposes. For example, the model may be used to test the effects of different anti-trust policies, or to allow a firm to study alternative scenarios consequential to different strategies. For such applications we would need to adapt initializations of the model reflecting some specific case.

Finally, given the controversial appreciation of simulation models, a short methodological comment. The models presented serve the purpose to organize coherently the different parts of a theory, since the programming language implementing the model requires explicit definition of functional forms and numerical values. The results obtained with simulation runs offer the possibility to analyse some histories of markets. Such histories concern virtual markets, but allow an enormously higher degree of detail in respect of real ones. Moreover, simulation models, if correctly implemented, allow gradual extensions. Already in this work we have developed two subjects in isolation and finally merged the two in order to form a more complex one. One may easily extend the model exploiting the existing structure and adding new modules to deal with some of the many aspects of the reality not considered. For example, all the problems concerning investments, capital, and, in general, the management of production processes are here ignored, but could be plugged in the present structure.

How the knowledge gained from virtual markets can be applied to real problems depends on the sensibility and knowledge of the specific market of the researcher. The purpose of this kinds of models is giving answers to questions like: *what happens if ...?*. The model can provide robust answers, but the appropriateness of the question depends only on the researcher's wisdom.

References

ABERNATHY, W. and UTTERBACK, J. (1978), "Patterns of Industrial Innovation", *Technology Review*, **80**.

FRENKEN, K., MARENGO, L. and VALENTE, M. (1999), "Interdependencies and Adaptation", in T. BRENNER, ed., "Computational Techniques to Model Learning in Economics", pp. 145–165, Kluwer Academics.

GIGERENZER, G. and GOLDSTEIN, D. (1996), "Reasoning the Fast and Frugal Way: Models of Bounded Rationality", *Psychological Review*, **103**(4), pp. 650–69.

GIGERENZER, G. and SELTEN, R., eds. (2001), *Bounded Rationality: The Adaptive Toolbox*, MIT Press, Cambridge, MA.

KARNANI, A. (1983), "Minimum Market Share", *Marketing Science*, **2**(1), pp. 75–94.

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

KLEPPER, S. (1996), "Entry, Exit, Growth and Innovation over the Product Life Cycle", *American Economic Review*, **86**(3), pp. 562–83.

SHAFIR, E., SIMONSON, I. and TVERSKY, A. (1993), "Reason-based Choice", *Cognition*, **49**, pp. 11–36.

SIMON, H. A. (1982), *Models of Bounded Rationality*, MIT Press, Cambridge.

SMALLWOOD, A. and CONLISK, J. (1979), "Product Quality in Markets where Consumers are Imperfectly Informed", *Quarterly Journal of Economics*, **93**(1), pp. 1–23.

VALENTE, M. (2000), *Evolutionary Economics and Computer Simulations - A Model for the Evolution of Markets*, Ph.D. thesis, University of Aalborg, Denmark.