A Model of Bounded Rational Consumers with Endogenous Preferences*

Marco Valente
Università dell'Aquila
mv@business.auc.dk

December 17, 2003

Abstract

This paper intends to be a contribution towards an evolutionary theory of consumer. Standard economic theory assumes perfectly rational agents acting upon exogenously given preferences. Here it is proposed a bounded rational representation of consumers, which seems even more justified for consumers than for producers since they are not subject to competitive pressures and there cannot exist any "punishment" for inefficient consumers.

The proposed representation acknowledges that the context in which consumers make their decisions is largely influenced by firms with their marketing activities, so that competing firms contribute to shaping their own selection mechanism, hence causing the (partially) endogenous character of preferences.

The consumer representation discussed, founded on cognitive theory and available evidence, is implemented in a simulation model which is tested with exercises meant to investigate the properties of a demand formed by such consumers. The results show that the such a demand respects the basic requirements, for example recognizing dominating products when they exist. Moreover, the model proposed is able to investigate the properties of markets for heterogeneous products, suggesting different categories and properties of endogenous market segmentation.

Keywords: Evolutionary Economics, Consumer Theory, Bounded Rationality, Marketing and Preferences, Simulation Models, Market Structure

JEL-classification: C63, D11, D81, L11, L15, M30

*This work is based on my PhD thesis, discussed in the Universiy of Aalborg, Denmark, Department of Business Studies, as a member of DRUID, Danish Research Unit for Industrial Dynamics. I wish to thank Esben Sloth Andersen for his careful supervision and many useful suggestions. Preliminary versions of this paper have been presented to: International J.Schumpeter Society Conference, Manchester, June 2000; Department of Economics, University of Trento, February 2001; Max-Planck-Institute for Research Into Economic Systems, Evolutionary Economics Unit, Jena, January 2002. Usual caveats apply.

Contents

Introduction

Evolutionary economics moves the focus from the properties shown by agents' behaviour (e.g. optimising) to what agents actually do (e.g. apply routines), so that room is made for heterogeneous agents. However, most of evolutionary scholars focused their attention on the supply side of markets, relying on an extremely sketchy representation of how markets actually function. This is reflected in the almost universal assumption of "homogeneous" products, implying that buyers cannot appreciate differences other than price. This allows to use a standard, aggregate demand curve, matching higher price with lower sales. However, demand curves prevents the analysis of the core of consumption behaviour, that is, the solution of the trade-off formed by products with different qualities, limiting the role of consumers only to the choice buy/don't-buy. Instead, an even casual observation of real-world markets cannot but acknowledge the crucial role of heterogeneity of products in explaining, for example, market dynamics and the economic impact of (product embodied) technological innovation. This is particularly important of evolutionary theory that gives large attention to the role of innovation, but then, in most of cases, treat almost exclusively the case of costs reducing innovations.

The role of demand is recently gaining an increasing attention among evolutionary scholars (Metcalfe, 2001). For example, it is discussed the origins of new diversified wants (Witt, 2001) or the role of variety in demand for growth (Saviotti, 2001). However, these works focus on aspects of demand independent from the markets where consumers operate, and generally tend to ignore the way consumers actually reach their purchasing decisions (but see Aversi *et al.*, 1999).

The consumer's representation proposed in the following is based on robust results obtained in experimental and cognitive psychology. Besides its realism, the model proposed qualifies as a working implementation of Simon's bounded rational behaviour (Simon, 1982). The consumer is assumed to make two logically distinct actions, although potentially they can take place at the same time: preferences formation (the criteria to follow) and the actual decision procedure (which option is chosen). The choice procedure Gigerenzer and Goldstein (1996) is based on the requirement for a "compelling reason" to act, and is compatible with the literature on experimental evidence (Shafir *et al.*, 1993). However, the literature does not suggest how the preferences used in the choice procedure are actually developed. However, in experimental studies the process of creating preferences is widely recognized to be influenced by the context in which they are applied: "*[...] 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). In the case of consumers the context is the set of "reasons" for purchase offered by producers via marketing, here generically defined as any action taken by firms for presenting their products on the market, i.e. to the potential buyer. The inclusion of marketing in a theory of consumer seems all the more justified by the immense importance that non-productive and non-organizational expenses in modern companies, and that are largely ignored in economic theory.

The preferences sets pressed for by each producer is likely to be incompatible with that of competitors'. A likely weighting mechanism at work to resolve conflicts emerging from different proposed preferences is identified in the social communication among consumers, since, has been observed, they rely heavily on the observation of their peers for their purchasing activities (Smallwood and Conlisk, 1979; Cowan *et al.*, 1997). Here it is suggested that market shares are used as weights to ponder conflicting marketing pressures, opening the way for endogenous preferences.

The demand resulting from the consumer model proposed highlights the to-way interactions between producers and consumers that contribute to shape market structure and dynamics. It is suggested that besides supply-driven dynamics (e.g. technological evolution) or demand-driven ones (e.g. emergence of new wants) the very mechanisms of trade contribute to shape consumer preferences, market shares and therefore, ultimately, actual industrial patterns.

A relatively simple simulation model implementing sets of consumers offered a heterogenous set of products provides different "virtual histories" depending on various assumptions. The model implements consumers undergoing the development of preferences influenced by producers' marketing efforts and by the current state of the market. The preferences are then used to make purchasing decisions according to the bounded rational paradigm. The series market shares distributions produced by the model are used to discuss how normally observed market configurations can be rooted in particular aspects of demand.

1 An Evolutionary Model of Consumer

The appeal of optimization to represent agents' behaviour is largely motivated by the possibility that, assuming the result (i.e. the optimum), we can neglect the means by which it is reached. However, besides the "standard" criticism to the use of optimality in Economics (Nelson and Winter, 1982), its use for consumer raises further scepticism. Firstly, while it is possible to conceive some optimization target for firms, although disputable[1], in most of cases this is not possible for consumers. Nor, of course, any "*as if*" argument can be invoked, since inefficient consumers are not subject to selection. Secondly, consumers, in respect of producers, should be assumed to be less committed to, and less expert of, the products and services they trade. In fact, in most of cases (and the most economically relevant) the role of a purchase in the buyer's overall life and income is negligible, and therefore consumers can hardly be expected to devote huge amounts of time and attention on a purchase. Of course, people do not like to waste money or buy lemons, as long as they can prevent it, and therefore some form of rationality is called for.

In the following it is presented a model of bounded rational agents representing the behaviour of consumers such that it can be applied in the evolutionary economic theory. A primary requirement for the model is that it must be able to cope with heterogeneous products. In fact, a central argument of evolutionary theory concerns the innovation, and therefore it is necessary that a theory of consumer can cope with differentiated products. Secondly, the model should realistically replicate the available evidence. The model uses the results obtained in the cognitive sciences, implementing an algorithm representing bounded rationality that is compatible with the huge evidence on actual people's behaviour from experimental studies. Thirdly, the model explicitly exposes the dependence of preferences on the suppliers' marketing activities. It is, in fact, evident that in practically every modern market the sellers' own information is one of the primary sources used by potential customers to make their choices. And the importance of marketing is testified by the huge amount of resources devoted for this purposes.

In the rest of this section we present the model by starting with the representation for the objects of trade in a market. Then it is presented the choosing procedure proposed, its implied definition of preferences and how they are represented in the model. Finally, it

[1]For example, optimising short-term profits may undermine longer term measures of success, or vice-versa.

is presented the "demand" built with such consumers in order to test the proposed model, whose results will then be discussed in the second section.

1.1 Product Representation

We can generally assume that the set of products offered in a market can be represented as vectors over a set of dimensions, or characteristics (see, e.g., Lancaster, 1966; Saviotti and Metcalfe, 1984; Gallouj and Weinstein, 1997).

	Char. 1	**Char. 2**	**...**	**Char. m**
Prod. A	v_A^1	v_A^2	...	v_A^m
Prod. B	v_B^1	v_B^2	...	v_B^m
...	...	...	...	...
Prod. N	v_N^1	v_N^2	...	v_N^m

Table 1: Products' quality values

In Table 1 the generic value v_X^i is the measure of product X in respect of characteristic i. This value must be interpreted as a measure of the quality for the "service" that the product provides in respect of a specific use[2]. Of course, it makes sense to compare only values along the same characteristic which, without loss of generality, are assumed to be positive ones defined over the set of real numbers. This means, for example, that one of the characteristics cannot be "price", but rather "cheapness", possibly defined as the inverse of price[3]. Note that a product can be defined absolutely superior to another only if all its values are superior to all the values in the vectors of the other competing products. In general, this representation permits to represent different trade-offs, with some product scoring better than competitors on certain characteristics and worse on others.

Such representation is very general, and can be used for many purposes. For example, it is possible to segment users according to a "minimal requirement" on each dimension. This, for example, can be used for selecting buyers who, although potentially interested, cannot afford the price for a given product, or the ones who necessitate minimal properties. In the following, however, we will consider the dimensions only as long as they are used for preferring one specific product over the competitors, assuming that the whole set of products available satisfies the minimal conditions for being considered for an actual purchase by all buyers.

Note that such representation is readily available to represent product embodied innovation, which can cause the increase of quality along one or more dimensions. But it is also possible to represent an innovation as the increment of the very set of dimensions, in case the innovation introduces new services in addition to an existing product.

1.2 Bounded Rational Decisional Strategy

The most challenging question concerning a consumer model is how to represent the choosing behaviour. In Gigerenzer and Goldstein (1996) it is convincingly sustained that the Take-The-Best strategy (TTB) is, both, experimentally supported by observation of actual people's behaviour and very efficient under the normally occurring circumstances of

[2]Although we will mainly refer to "products", all the analysis may well be applied to markets for services.

[3]As we will see below, the absolute values for products' qualities are not important, since the procedure used needs only to make comparisons determining whether one product is superior, inferior or equivalent to another product in respect of one dimension.

uncertainty and poor information. Concerning the general case of choosing one alternative options among many, each defined over a set of characteristics, this strategy is composed by the following steps:

1. Choose one characteristic, among the m available.
2. If one single option scores highest in respect of that characteristic, that is your choice.
3. Otherwise, if more than one option scores similarly, remove the dominated options, and restart from 1. with another characteristic.

The authors proposing this strategy as a representation for human behaviour concerning decision making argue convincingly that it is an algorithm respecting the principles of bounded rationality (Simon, 1982), which seems quite adapt to represent the generality of consumers' behaviour. In fact, most of the purchasing decisions, and by far the ones of larger economic impact in modern markets, are made by people buying items whose costs and importance is very limited in respect of their overall life and income. Therefore, they have relatively little interest in investing time and intellectual resources just for being sure of picking the optimal choice, and would rather risk to choose a dominated alternative, possibly a little more expensive in price or of slightly lower quality, but by far easier to be decided upon. Concerning the realism of TTB, there is a huge amount of literature suggesting that, when people face the choice between different alternatives "*[...] they resolve the conflict by selecting the alternative that is superior on the more important dimension, which seems to provide a compelling reason for choice*" (Shafir *et al.*, 1993, p.15). TTB is nothing more than a generalization of this "reason-based" procedure, extending to consider more than two dimensions and the possibility of equivalence of some alternatives in respect of some dimension[4].

In the TTB it is crucial the order in which characteristics are used to make a decision. It is obvious, in fact, that in general one obtains different choices depending on which characteristics are initially used to select upon the available alternatives. This provides a useful insight and a new problem. In fact, we obtain a pretty precise notion of "preferences" defined as the order in which characteristics are used when applying the TTB strategy. That is, the relative importance of products' characteristics make up the set of preferences. For example, suppose only two characteristics exist on a given product: quality and cheapness. Then there are two possible preferences: highest quality or lowest price. The "quality oriented" buyers will prefer the cheaper products only among those scoring highest in quality, while "price oriented" buyers will buy the best product among the cheapest. In the general case of many dimensions the number of preferences is given by the number of all possible permutations of the characteristics.

The new problem concerns the origins of preferences. In the following paragraph it is suggested that firms offering their products have a large say on this.

1.3 Marketing Induced Preferences

Given this notion of preferences, we need to ask ourself where they come from. A frequent opinion among economists is that preferences are exogenous, referring to psychological, social and other aspects independent from the actual action of purchase. Actually, most of what is vaguely referred to as exogenous preferences can be roughly grouped in two sets:

[4]Another interesting property of TTB, not used here, is that it can deal with missing values, typical when people ignore some of the characteristics on the available choices. However, this aspect is not considered in the model presented here and therefore it is not discussed.

the constraints (e.g. few people can afford a Ferrari, and therefore consider it when deciding which car to buy); and the purchase's aim (even though many people may afford an opera ticket, few actually desire to attend an evening). So, preferences, properly defined, should be invoked only when discussing the decisions made by people with the same constraints and the same aims, simply preferring one option over another. In these cases people still reach different decisions for reasons only due to preferences.

Some authors have discussed how new wants stem from the general economic growth (Witt, 2001), or how purchasing behaviour is influenced by social considerations (Cowan *et al.*, 1997; Aversi *et al.*, 1999). However, an important determinant of buyers preferences is rooted in the very action of consumption. In fact, it is well known that the way a choice set is presented heavily influences the final response (see, e.g., Kahneman *et al.*, 1982). That is, "*[...] preferences are actually constructed - not merely revealed - during their elicitation.*" (Shafir *et al.*, 1993, p.34).

The activities of firms generally considered in econonic theory focus on the production process and their results (i.e. the products), and, in some cases, on the internal organization of the firm. However, an important (and increasingly so) part of firms activities concern how their products are presented to the public of potential buyers. If we look at the balance sheets of modern companies we see that huge amounts of resources are devoted to marketing. Although differing widely between industries and individual companies, marketing is always responsible for several tens of percentage points of a company overall expenses. What is the function of this item? For sure, part of this is required in order to have a company's brand known to the public[5]. However, established firms do not need, if not marginally, such brand awareness efforts. Here it is suggested that the best way to describe the effects of marketing is as an effort to create and bias buyers' preferences in such a way to better highlight your own product against competitors.

The marketing of a firm can be embedded in any form: particular packaging, commercials, sponsorships, etc. In any case it presses buyers for adopting a particular perspective of the product that exalt its product against competitors. In our setting this translates in an effort to convince buyers that some characteristics are more important then others, so urging potential consumers to develop a specific preferences set.

Given the definition of preferences as a ranking of products' characteristics, the logical way is to assume producers to have their own "ideal" ranking, that they would like to push through consumers, exalting their own product against competitors'.

	Char. 1	**Char. 2**	**...**	**Char. m**
Prod. A	k_A^1	k_A^2	...	k_A^m
Prod. B	k_B^1	k_B^2	...	k_B^m
...	...	...	...	...
Prod. N	k_N^1	k_N^2	...	k_N^m

Table 2: Producers' marketing strategies

The generic element in Table 2 must be interpreted as the relative importance that the producer gives to the characteristic for its marketing strategy. The more important a characteristic, the higher its value compared to the other values for the same producer on the other characteristic. In the following, when describing how buyers develop their preferences, the exact meaning of these values will be clearer. For the time being, consider that, in principle, a coherent producer should set its k_X^i higher the more its product is better than competitors' in respect of the characteristic i.

[5]Recent surveys on internet startups showed that some 50-70% of seed money was spent on marketing.

1.4 Social Influences

We have seen how suppliers can influence buyers, but we still need to specify how a buyer determines its preferences out of the set of marketing pressures from all the suppliers. In other terms, what makes more effective the marketing pressure from one firm in respect of competitors' (assuming they apply incompatible marketing strategies)?

In Smallwood and Conlisk (1979) the authors propose a model where buyers choose their purchases randomly with probabilities proportional to the competitors' market shares. The justification for this is obvious: there is no better advertising that just having many users showing your product around. This method seems even more adapt to represent the relative diffusion not of products, but of preferences. Consumers pass each other "perspectives" of the product, and the probability of choosing one given perspective is likely to depend on the number of "members" that already adopted it. Cowan *et al.* (1997) support this view, suggesting that consumers act according two principles: distance yourself from lower class elements, and imitate upper class ones. Since we are studying the properties of buyers induced by market trading, it seems reasonable to use only the second aim, that is, buyers tend to imitate each other's behaviour.

As a first attempt to implement this kind of behaviour we can therefore assume that buyers' preferences are determined by choosing randomly among suppliers' marketing strategies weighted with their market shares. Formally, the algorithm used to determine a buyer preferences is the following. We define preferences as the set of integers representing the m characteristics:

$$< c_1, c_2, \ldots, c_m >, c_i \in \{1, 2, \ldots, m\}$$

For example, if m=3, the possible preferences are:

$$< 1, 2, 3 >; < 2, 1, 3 >; < 1, 3, 2 >; < 3, 1, 2 >; < 3, 2, 1 >; < 2, 3, 1 >$$

The model determines the preferences for a consumer using the following indicators:

$$p_i = \sum_{j=1}^{n} (k_j^i s_j)^{\delta} \tag{1}$$

where s_j represent the market shares of supplier j, k_j^i is the marketing level of firm j in respect of characteristic i, n is the number of firms, and δ is a coefficient flattening or steepening the indictors[6].

In short, the indicators (1) show the importance that the supply side of a market as a whole gives to the i^{th} characteristic. The more firms (weighted with their market shares) will press for one characteristic, the more likely users will consider it important for their preferences.

The results we want to obtain is to assign the preferences to a consumer. Let's begin with the first, highest ranking, characteristic in the preferences, c_1. This is obtained drawing randomly one of the m characteristic each with probability:

$$\Pr(i = c_1) = \frac{p_i}{\sum_{h=1}^{m} p_h}$$

[6] A value of δ approaching 0 means that the indicators for all characteristics will tend to be equal. Conversely, higher and higher values of δ will make one p_i close to 1 (the one with higher marketing support) and all the others near 0.

The same indicators are used for obtaining the probabilities to draw the second characteristic in the preferences, just setting to 0 the indicator p_{c_1}:

$$\Pr(i = c_2) = \begin{cases} 0 & i = c_1 \\ \dfrac{p_i}{\sum\limits_{h=1, h \neq c_1}^{m} p_h} & i \neq c_1 \end{cases}$$

The same procedure is used to assign the probabilities to all the remaining characteristics for the ranking positions, each time setting to 0 the indicators for the characteristics already used, and renormalizing accordingly.

1.5 Products' Values Perception

The TTB strategy defined to represent buyers' purchasing decisions relies on comparisons among products' values on individual characteristics. Considering the largest majority of cases in real world consumption markets it seems logical to distinguish between what the real values of products' qualities are, and what consumers *think* they are, and use in their decisional processes. In fact, in most cases buyers are not experts in the technology embedded in the products, or do not consider worth to make extensive research for finding the exact quality values. Moreover, sellers may be reluctant to make public detailed information concerning their products, preferring to remain vague on their true products' qualities.

We consider two types of "distortions" from real values to perceived ones. The first consists in a random error, so that the observed value of a product's quality is a random draw from a normal distribution centered on the real value, and whose variance depends on the buyer's "expertise". Buyers are represented as using, in place of the actual values v^i_X for product X along characteristic i, a random draw from a normal distribution:

$$\hat{v}^i_X = Norm(v^i_X, \Delta_t)$$

The normal random function has the mean value identical to the true value of the product, and a buyer's specific deviation Δ_t changing through time, which will be discussed below. Such draws are repeated any time a buyer need to observe the quality value of a product.

To appreciate the effects of the use of stochastic errors, consider that the values observed are used only to make a comparison between two products, and that the normal function is centered on the true values. If $\Delta_t = 0$, then the buyer observes the true values of products, making no mistakes, and therefore will always choose the superior product in respect of the characteristic she is currently using. Conversely, for very large values of Δ_t, the result of the comparison will be totally random, with (almost) 50% of probability for each product to be perceived as superior, irrespective of their true values. For intermediate values of Δ_t, the probability for an inferior product to appear as superior to the true optimum is less than 50% and proportional to the actual difference between the products: the higher the difference, the lower the percentage of "mistakes".

The value of Δ_t is a buyers' specific variable, changing through time with an exogenous dynamics. This represents the fact that novice buyers are more likely to make mistakes (higher Δ_t's). The longer they use the product, the sharper they become in assessing the true product quality levels (lower and lower values of Δ_t). The initial and limit level of Δ_t, as well as its decreasing speed, is exogenous and identical for all buyers. The simulation examples presented in the next session differ, among other things, for the limit levels which Δ_t for each consumer is allowed to reach. A limit level equal to 0 means that, in the long term, buyers develop the capacity to perfectly judge products values. A positive limit level, instead, represents the fact that, though improving through time, buyers' expertise does never become good enough to completely remove the possibility of errors.

Besides perception errors, buyers use a margin of tolerance when comparing a product with the (apparently) optimal one, in respect of a characteristic. That is, if the difference between the (observed) values is less than a given percentage, the two products are considered equivalent:

$$\hat{v}^i_{Max} \simeq \hat{v}^i_X \iff \hat{v}^i_{Max} \cdot \tau < \hat{v}^i_X$$

where τ is a coefficient in the [0,1] range. The use of a tolerance margin represents an obvious protection against the discarding of very slightly inferior products in the first stages of the TTB algorithm (when using the early characteristics), which are possibly much better in respect of further aspects.

1.6 Model Dynamics

We have defined the supply side and demand side of the model. Supply is made of a set of products that, over a set of product characteristics, are defined by quality values and marketing strategies. Demand is made of a set of buyers that set their preferences and make purchasing decisions. We now need to define how a simulation run puts together the above elements to represent the dynamics of a general market.

We consider a simulation run as representing the development of a market for a novel product: at the beginning, only one buyer is an early user of the product, and, later, more buyers are brought in the market up to a fixed number of consumers, reaching the saturation of the market. One of the very few universal facts observed in many different markets is the s-shaped dynamics of the number of buyers in new markets. This phenomenon can be simply represented assuming that the rate of entry of new buyers in a market is determined by a contagion diffusion model. That is, early consumers show their novel purchase to their friends, who mostly ignore it. Each of the new consumers brings in new consumers, giving rise to the initial increasingly steep growth. In later stages, most new consumers have friends who already use the (no more so) novel product, so that we observe a diminishing increment rate of consumers, until the saturation level is reached and no new consumers enter the market (apart the demographic turn-over, that we ignore)[7].

To simplify the interpretations of the simulation results we assume that a new consumer decides the preferences at time of the very first purchase, and then apply always the same preferences for each subsequent purchase. This is obviously a quite unrealistic assumption, although we can invoke some sort of learning lock-in. For example, a new entrant in the early stage of the personal computer markets may "absorb" the preferences sponsored by Apple, and therefore choose one of their computers. Later, although the market shares of Apple have sensibly fallen, this consumer will apply again the same preferences, e.g. requiring fancy graphical interfaces (and likely choosing an Apple again, at least until IBM-compatible will offer a comparable windowing system). Another kind of justification of this assumption is that preferences, even if changing, are modified quite slowly, and anyway slower than the time scale we consider in the model. In any case, this assumption

[7] This buyers' entry mechanism opens the possibility, not exploited in the current version of the model, to construct cohorts and "family trees" of consumers, depending on when they entered the market and introduced by whom. Such information can be exploited to represent differentiated reactions for, say, consumers entered in the early stages of the market evolution in respect of later entrants.

does not change radically the results, producing a clearer representation of the phenomena discussed below.

A consumer in the market is supposed to make a purchase every few time steps, assuming the most general case of the product being a semi-durable[8]. This causes a differentiation between market shares measured on sales (determined by the buyers who make a purchase at a given time), and those measured on "installed bases", that is buyers who bought a product in the past and are still using it. For the determination of preferences seems obvious to make use of the installed bases distributions, more likely to be observed by new consumers than daily sales shares.

2 Simulation Results

In this section we will present some of the possible applications of the model presented above. The goal of these simulation exercises is not to replicate realistically a given phenomenon. Rather, the goal is to show "what happens if" the elements discussed in the previous sections are set in a specific way and allowed to run through time. The results will not be considered for their realism, however measured, but in respect of the possibility to causally link the aggregate phenomena generated during a simulation run with the micro interactions among the elements of the model. That is, we are interested in investigating the causal relationship of the events taking place in the simulated environment, and it is these "explanations" that are compared with real world evidence.

For these purposes, the specific numerical values used to initialize the simulation runs, as well as the data produced, do not matter beyond their capacity to represent a phenomenon and to suggest an explanation for it. For each experiment presented below we will declare only the initialization properties relevant for that configuration, without specifying in detail all the values. Similarly, the results will be presented by means of graphs where, in some cases, the individual series do not need to be identified, since it is their aggregate properties that matter.

2.1 General Setting

Our interest is focused on the market configurations emerging from specific demand characteristics only, excluding modifications on the supply side of the market. Of course, this is a highly unrealistic case, but allows to identify the contributions of demand to the overall market dynamics. In particular, we will consider the effects of three elements. Firstly, we will set up different exogenous supply sides, that is sets of sellers with products' quality levels and marketing strategies. Secondly we will consider the effects of different "learning" limit levels. That is, although in every case buyers improve their capacity to judge correctly products' qualities, in some cases they end up making no errors at all, while in others they continue to make small errors for ever. The last effect we consider is the effect of different levels of tolerance, so that, in some cases, minimal differences in quality levels cause a product to be discarded, while in others all products whose qualities close the best product are considered equivalent.

The model contains thousands of parameters, which severely affect the simulation results, besides the ones we are interested in studying. If we were to make simulation analysis as a complete representation of a real world counterpart, and we would like to claim

[8] Remember that our model does not allow for supply dynamics, and therefore it is not possible to model a technological competition causing replacement of obsolete products (Richardson, 1996). However, this phenomenon may be easily implemented with few modifications to the current model.

universal results, then we would need to test all those potentially relevant parameters. However, as already repeated, this is not our goal. We consider the model as an artificial construction, where we can play artificial histories. Our aim is to understand what happens during this histories, and why. It will be these explanations that will be compared with real world examples, where they surely will apply together with many other elements not included in the model.

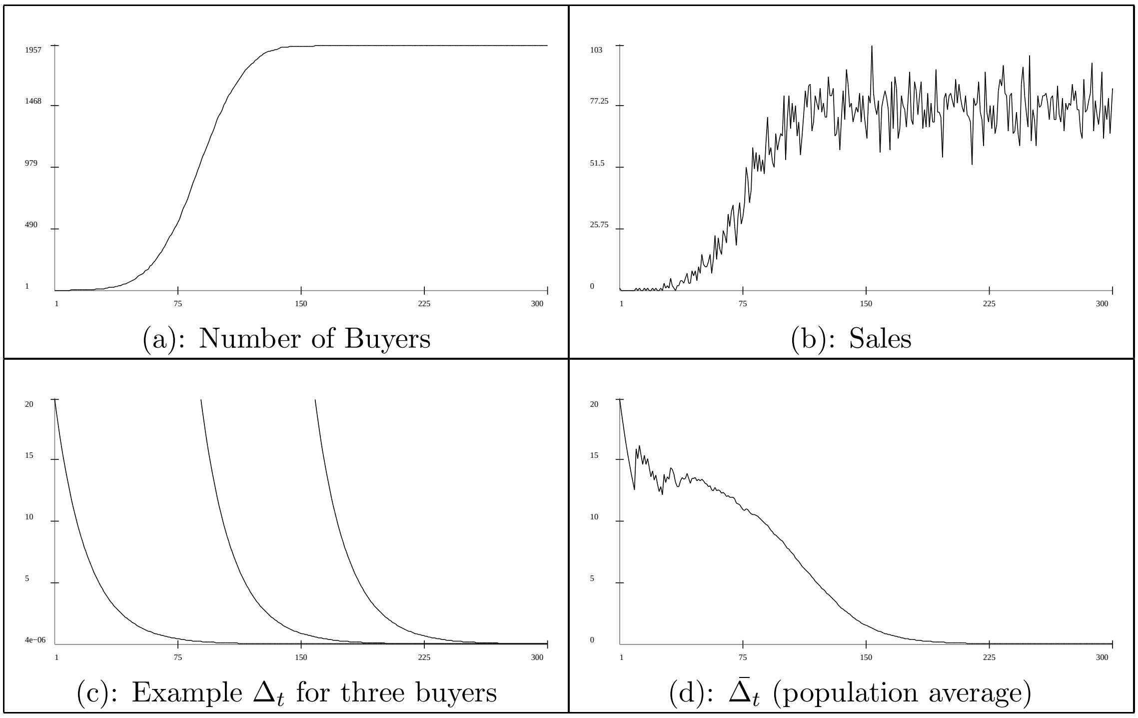

Figure 1: General settings. Each simulation run uses the same initialization producing the following dynamics: (a) the dynamics of buyers' entry; (b) the level of sales; (c) the level of Δ_t for single buyers (starting and limit values change for different exercises); (d) average of Δ_t for the whole population.

In all the simulations we will use some common settings. We consider a market for a novel product, where the buyers enter (i.e. start to purchase, use and from time-to-time re-purchase a new product) following a contagion diffusion dynamics. At the beginning of a simulation run only one buyer is present in the market; this will introduce 6 new buyers in its lifetime, each of which will bring 5 new buyers. Eventually, the market dimension is made of almost 2000 buyers (see Table 1:(a)).

Each buyer makes a purchase at the first time step it is introduced in the market, and then replaces the product every few time steps. The time between two purchases is randomly determined, within a given range (in between 20 and 40 time steps). Therefore, the number of actual sales at each period is always a fraction of the total number of buyers present in the market at that time step (see Table 1:(b)).

Each buyer develops independently its own capacity to read products quality, represented by Δ_t (Table 1:(c)). While we will vary the Δ_t limit values, they will always follow the same dynamics. Buyers enter at different times, and therefore the average level of Δ_t's reflects the population composition (Table 1:(d))

Most of the evidence presented below will concern the values of sales per producers. It will be presented showing the levels of installed bases per producer, which is the number of buyers who bought from that producer at their latest purchase. Therefore, the sum of installed bases over all producers sums up to the total number of buyers in the market.

2.2 Simple Monopoly

In order to make clear the functioning of the model, let's consider first an extremely simple case. The supply side is made of 10 products defined over one single characteristic. All products have the same value for the characteristic quality, but one product which is set to a superior value. Buyers enter in the market (with the contagion diffusion dynamics) having a very poor capacity to assess products' quality (making large errors). Through time, however, each buyer develops increasing expertise of the product and therefore reduces the scope for errors, being able to identify with higher and higher probability the best product. In the event, buyers reach perfect capacity to judge products' qualities ($\Delta_t = 0$)

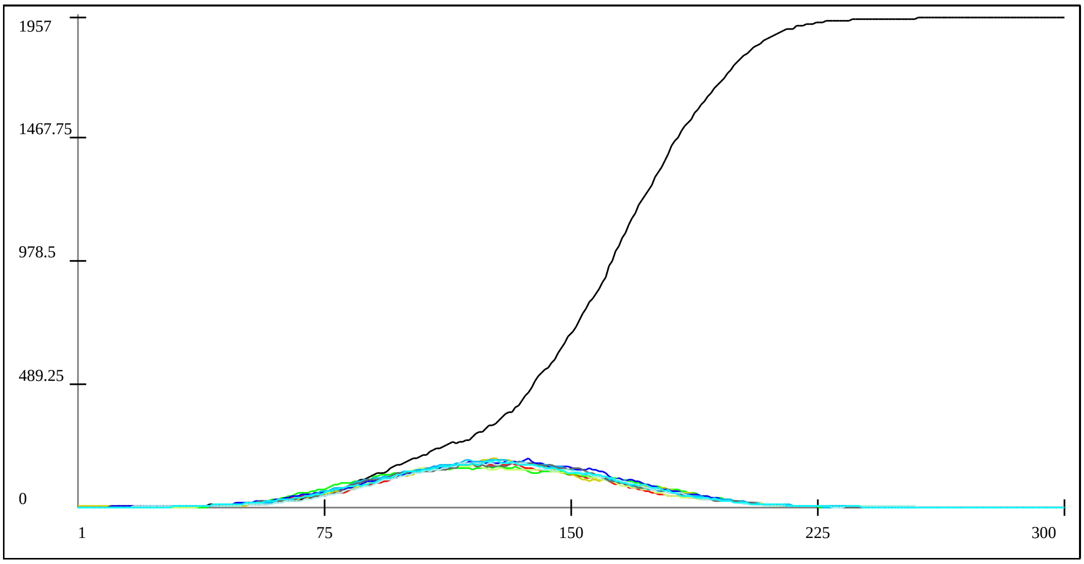

Figure 2: Installed bases dynamics for Standard Monopoly. All producers have identical quality and marketing values, but for one producer being superior than competitors. The superior product eventually dominates the market, after that buyers have reached perfect capacity to observe the true quality values ($\Delta_t \simeq 0$).

Figure 2 shows the series of buyers using each of the products through time, that is the level of installed bases for each product. Up to about time step 100 there are new buyers entering the market. The new buyers prefer mostly the "monopolist" product, superior to competitors. However, the dominated products can still make positive sales because of the poor capacities of buyers to assess products' values. That is, when comparing the products' values, the distortion in the observed values make appear dominated products as superior to the actual best product (high value of Δ_t). Later, in the final stages of the simulation, the value of Δ_t falls to zero for all the buyers, and therefore all buyers consistently choose the best product, producing the expected monopoly.

Even though almost trivial, this exercise shows the typical properties of new markets (see, e.g., Abernathy and Utterback, 1978): a confused initial stage, where many competing products enjoy positive and increasing sales, followed by the emergence of a dominant product. There are many explanations for this robust finding in the literature. In this exercise, given the assumption of constant products' qualities, the only explanation is that the buyers' limited knowledge concerning the new product causes them to make comparison errors, which later can be avoided because of the increased capacity to correctly assess the true products' quality.

2.3 Monopoly with errors

As a second exercise, let's consider an identical set up as in the previous example, but we limit the maximum level of knowledge that buyers can develop (that is, $\lim_{t\to\infty} \Delta_t > 0$).

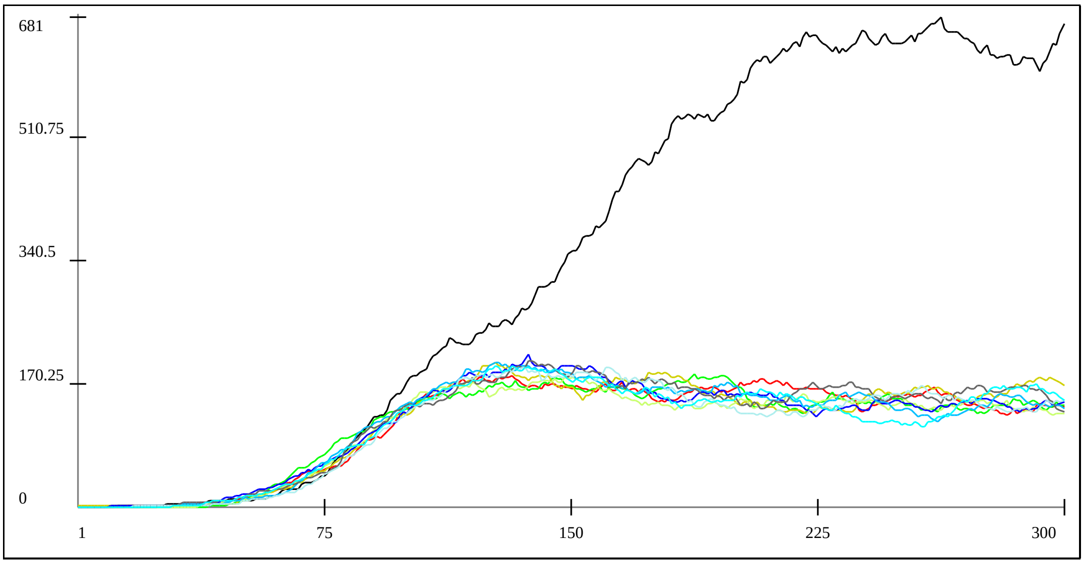

Figure 3: Installed bases dynamics for monopoly with limited learning. Buyers cannot reach perfect capacity to observe products' qualities.

The results in Figure 3 show that the "monopolist" is not able to gain the whole market, although enjoying a superior market share than competitors. The obvious reason is that buyers continue to make errors in judging products' quality, although with lower probability than in the beginning. In real world cases this may be due to either an objective difficulty in comparing products, or to scarce interest by buyers in spending time to assess products' relative quality. In marketing literature has been noted that one possible strategy for producers is to exploit this aspect (Karnani, 1983). That is, in markets where the buyers do not have very strong incentives in taking the best products, a producer can target the share of buyers who pose little attention to their purchases, and therefore are likely to misjudge products' qualities. As a result, even with grossly inferior products any producer can ensure a minimum level of sales.

2.4 Monopoly with quality tolerance

In this exercise we consider the effect of parameter τ. We use the same set up as in the first exercise: all products identical, but for one superior to competitors', and buyers allowed to reach perfect capacity to judge products in the long term. However, the value for τ is set to less than 100%, and the quality difference between the superior product and the others is within the tolerance margin.

As expected, the level of installed bases series in the long term is identical for all products in the final steps of the simulation (see Figure 4). In fact, at this point, all buyers are able to acknowledge that the higher level of quality of the best product is not relevant, and therefore choose randomly among all available products. However, this

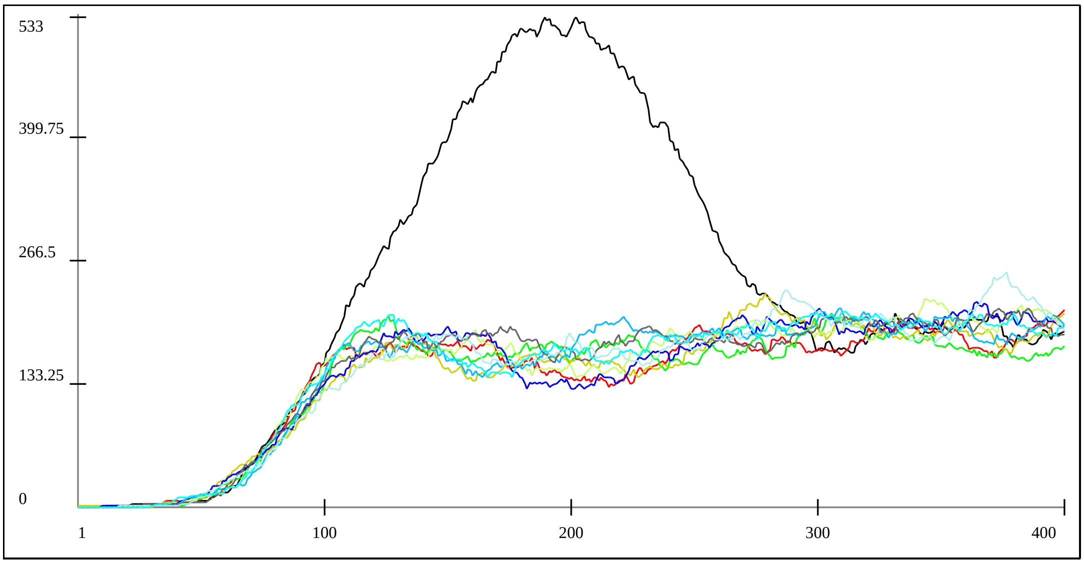

Figure 4: Installed bases for monopoly with tolerance. In the long term, buyers consider the quality difference of the best product as not relevant, but in the medium term the positive level of Δ_t makes appear that difference higher, therefore favoring the best product.

final configuration is reached after a period in which the share of the superior product is systematically higher than competitors. This phenomenon is due to the combined effect of the errors made by buyers and by the tolerance. Although the actual difference between the quality of the best product and the others is within the tolerance margin, when buyers are still making errors (relatively high Δ_t) it is possible that the observed values of quality have a larger difference than the one admitted in the tolerance:

$$v^i_{Max} \cdot \tau < v^i_X$$

and

$$\hat{v}^i_{Max} \cdot \tau > \hat{v}^i_X$$

where $\hat{x} = \text{Norm}(x, \Delta_t)$.

This effect has an obvious interpretation, and a quite counter-intuitive consequence. The interpretation is that when buyers are not perfectly able to judge products' qualities, they tend to favour higher quality ones, even when this superiority is, in fact, so small that shouldn't make any difference. Think, for example, the competition between AMD and Intel to produce "the fastest processor". In some stages of this competition, the two companies made grand statements about minimal clock superiorities of their processors, which no user would ever be able to detect. The same applies to many commercials concerning hi-fi systems, where producers advertise distortion levels very far below the audible threshold. The counter-intuitive consequence is that superior products may take advantage of not informing perfectly their customers. This is, in effect, what frequently happens: in their technologiacal race, Intel and AMD claimed alternatively to sell the "world fastest processor", not that theirs chips offer 1000 MHz while the competitor's reaches only 999 MHz. . .

2.5 Self-reinforcing oligopoly

Let's consider now a market where products are defined over many (10) characteristics. In this first example of a "natural oligopoly" we consider a group of producer each of which specialises in one characteristic.

Each product has a high level of quality in respect of one characteristic, and a low level for all other characteristic (values for "high" and "low" qualities are identical for all products). Correspondingly, the value in the marketing strategies vectors for the producers are high for the characteristic on which they are strong, and low for all other characteristics. Concerning buyers, we use the default configuration where the initial level of Δ_t is positive and decreases to 0 in the long term, with no tolerance for quality differences. That is, in the long term all consumers "learn" perfectly how to judge products' qualities, and any small difference make the inferior product to be discarded.

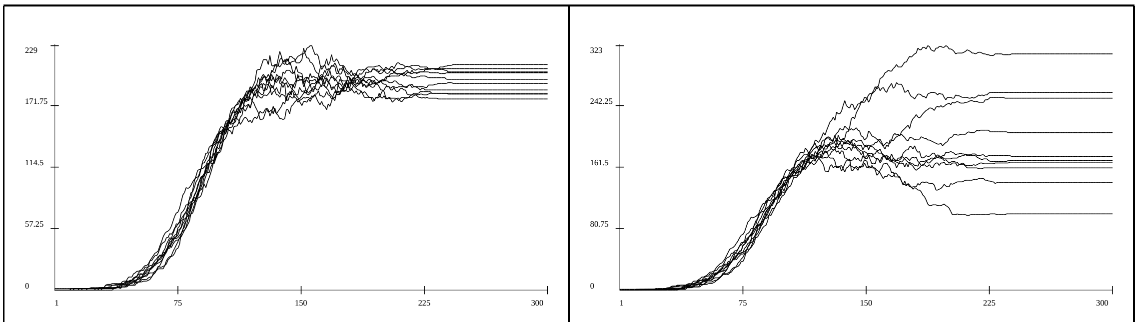

(a): Low δ (b): High δ

Figure 5: Installed bases for specialised producers. Each product has one characteristic higher than competitors', and focuses its marketing on that characteristic. Low δ simulation (a) has lower concentration, while high δ has higher concentration

As can be seen in Figure 5 the distribution of market shares among producers is random: each buyer sets its preferences depending on the market shares distribution at the moment it enters in the market. Since then, and until the level of Δ_t is high, the user makes mistakes in reading products' quality levels, so that it selects almost randomly the products. But in the long term (when Δ_t falls to 0) the buyer will choose consistently the product scoring better in respect of the highest ranking characteristic in its preferences.

The number of buyers who develop preferences for a given product is purely random, depending on the market shares distribution at the time of its preferences formation. However, the market *concentration* is not random. It depends on the strength of the self-reinforcing cycle between preferences and sales. The higher the sales for a given product, the more effective is its marketing pushing up the number of buyers adopting preferences that advantage that product, which translates in higher sales. The parameter controlling the strength of the self-reinforcing mechanism is δ, which determines the concentration of probabilities to adopt preferences (see equation 1, pag. 8).

In Figure 6 it is shown the average (inverse) concentration for different values of δ. The indicator used is the inverse herfindal, indicating the number of firms with identical size which would determine the same concentration of the measured market.

The Figure 6 shows that for low values of δ practically all firms have identical size, while with higher δ values the market becomes more and more concentrated.

This phenomenon is similar to the effect produced by increasing returns to adoption (Arthur, 1994). That is, if products' quality is a function of the number of users, then the market will concentrate on a single product even if there is no objective reason for

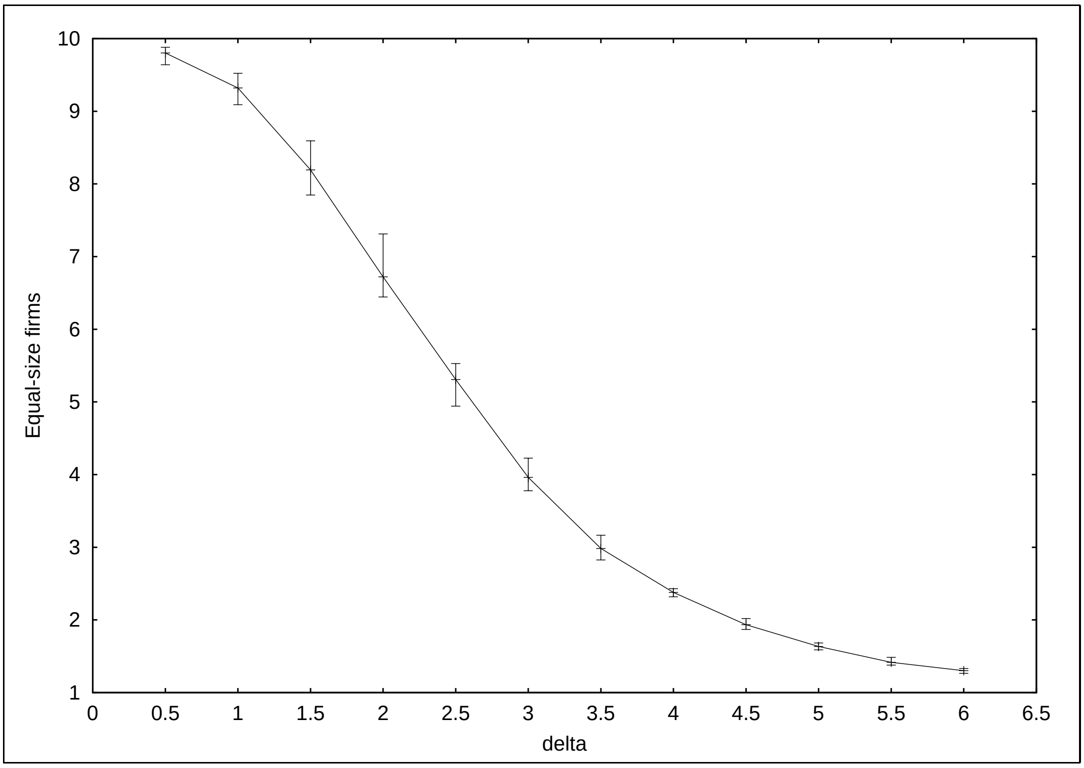

Figure 6: Concentration for different δ. The figure shows the minimum, average and maximum inverse herfindal index value computed at the end of 10 runs for each value of δ from 0.5 up to 6. The inverse herfindal index indicates the number of firms with equal size which would provide the same concentration as the one actually measured.

preferring it against competitors'. In our case, the effect is extended beyond the case of network externalities: although the quality of products does not depend on the number of users, still the self-reinforcing cycle applies to the development of preferences. It is the preferences, influenced by market shares, which are a function of the number of buyers.

The same effect has been noted in Smallwood and Conlisk (1979) concerning the diffusion of information concerning products' quality. More in general, it is possible to suggest that self-reinforcing cycles can be observed any time market shares matter concerning buyers' decisions. This implies that the path-dependency is the norm, and that δ is the crucial indicator to predict market concentration. The interpretation of δ is quite straightforward. It measure the degree of trust that buyers assign to signals from markets. The higher consumers rely on market signals for information concerning a product, the higher the concentration of the market. The more information is obtained by other, market-independent, means (e.g. individual experience, technical literature, personal independent preferences, etc.), the less concentration can be expected since the market will be divided in relatively insulated niches.

2.6 Simple Market Segmentation

In the previous example we did not approach the issue on whether the market is efficient in selecting the "optimal" product. In fact, we considered a model with heterogeneous and non-comparable products: each of them is the best in respect of one dimension.

Let's consider now an exercise where a random initialization for suppliers. The values

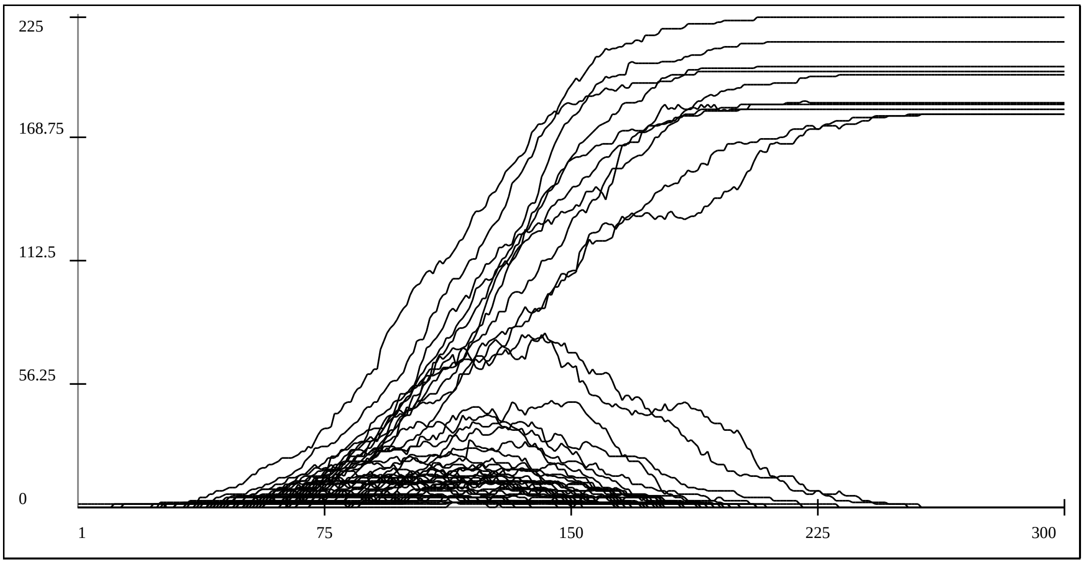

Figure 7: Random oligopoly (50 firms are assigned random values for quality and marketing strategies). Buyers are set with perfect capacity to read product's values and no tolerance ($\lim_{t\to\infty} \Delta_t = 0$ and $\tau = 100\%$)

for products' qualities and for marketing strategies are set randomly. The setup used does not enforce the consistency between products' quality values and marketing strategies. That is, some firm may be very good on some aspect of the product, but its marketing fails to highlight this aspect. This problem is partly overcome by the use of many firms (50) which ensure that at least some of the producer do have some consistency. Even though the aim of these exercises is not to provide totally realistic representations of actual markets, the lack of an explicit consistency requirement should not be understood as totally unrealistic. In fact, the results obtained by marketing are known to be highly uncertain, and there are many examples of marketing failures where marketing campaign actually obtain the result opposite to the desired one.

As a first test with this set up we use buyers that are allowed to reach perfect capacity to observe the true quality level ($\lim_{t\to\infty} \Delta_t = 0$) and no tolerance ($\tau = 100\%$). The results are shown in Figure 7. The usual dynamics can be observed until buyers' reach their "learning" limit value, even more chaotic than in the previous examples due to the presence of 50 firms instead of 10. Eventually, however, the market is perfectly divided in niches, where each buyer selects the product scoring better in respect of one characteristic, and therefore only 10 products have positive market shares, one for each product characteristic.

The picture changes sensibly when we use the same set up as before, but we prevent buyers from reaching the capacity to perfectly read products' qualities ($\lim_{t\to\infty} \Delta_t > 0$). In this situation buyers keep on making "mistakes", and therefore they happen to choose not only the best product in respect of their most preferred characteristic, but also other ones with inferior quality levels. The probability for a product to be chosen is inversely proportional to the difference between its quality level and the best product on each characteristic. In this case there are two circumstances in which a firm can attract customers. The first, and obvious, one is to be "good" in some characteristic of the product in order to appeal to as many customers as possible who value mostly that aspect. In this case, the firm appeals to many customers having that characteristic as top ranking in their preferences. The second strategy is to be "reasonably good" in many characteristics, so to be

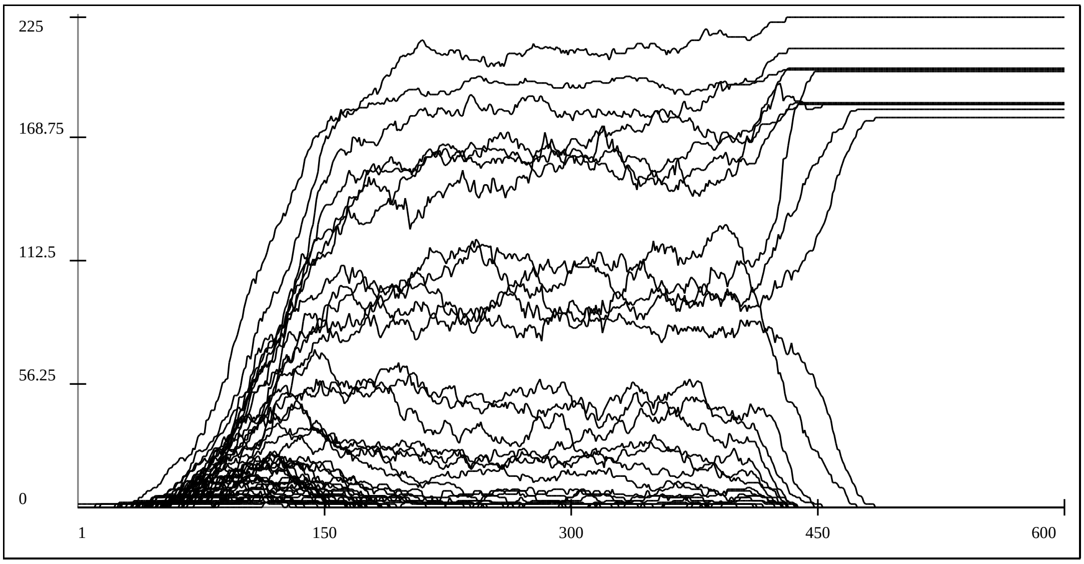

Figure 8: Random oligopoly (50 firms are assigned random values for quality and marketing strategies). Up to time 400 buyers are set with imperfect capacity to read product's values ($\lim_{t\to\infty} \Delta_t > 0$). Many firms have positive shares because customers can make mistakes in reading products' qualities. After time 400 the limit learning is set to 0, so that only the 10 best firms on each characteristic keep on making positive sales.

chosen by a minority of customers from many "niches". In this latter case, the customer base for the firm is variegated, formed by customers with different characteristics on their preferences.

For appreciating this phenomenon consider the Figure 8, where the simulation is modified at run time to make clearly emerge the two types of firms. Up to time step 400 buyers are prevented from reaching perfect capacity to read products' qualities ($\lim_{t\to\infty} \Delta_t > 0$). After this period we changed the simulation parameter for the $\lim_{t\to\infty} \Delta_t$ setting it to 0 and continued the simulation for other 200 steps. Consequently, all buyers quickly stop making errors and choose consistently only the best product in respect of their most preferred characteristic. Therefore, only the 10 firms scoring the maximum in one characteristic continue to make positive sales, while the others, that were previously "exploiting" buyers' errors, drop to zero market shares. The firms that disappear after period 400 are the ones that were relying only on buyers' errors, not having a definitive advantage over competitors over any characteristics.

These two types of firms can be considered as specialists vs. generalist: the former face a homogeneous demand that chooses them for a well specified reason; the latter serve heterogeneous customers, choosing the products for many different reasons. This different types of demand clearly imply that producers require different strategies for reaching their goal, like increase customer base, or defend from competitors' innovations, and therefore being able to affect the direction of technological innovation (Bresnahan and Malerba, 1997). However, this work, and the model, is limited to consider only the demand dynamics, and in the following paragraph we consider more in detail, and more realistically, the formation of market segments.

2.7 Complex Market Segmentation

In the previous experiment customers' preferences were defined only by the top ranking characteristic, because that was the only one used (possibly erroneously, choosing not the best product). As a consequence, we could have divided the market in 10 niches, one for each characteristic, plus a group of firms exploiting customers' errors.

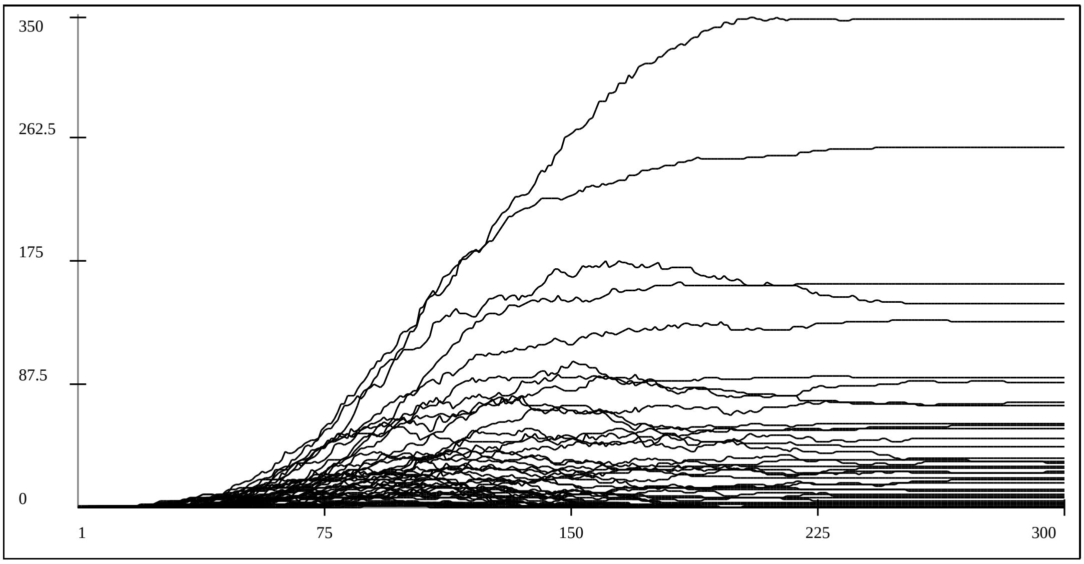

Figure 9: Random oligopoly (50 firms are assigned random values for quality and marketing strategies). Tolerance level τ is set to less than 100% and $\lim_{t\to\infty} \Delta_t = 0$. A complex structure of market niches is formed by means of heterogeneous preferences.

However, even with the simple structure of this model it is possible to represent quite complex market segmentation. In this last experiment (see Figure 9) we use an identical set up as in the previous experiment but we set the τ to less than 100%. We allow buyers to reach perfect capacity to judge products qualities ($\lim_{t\to\infty} \Delta_t = 0$) so to better appreciate the results by removing random fluctuations due to this effect.

After the usual chaotic phase, during which buyers are still unable to perfectly judge products' qualities, we observe a completely stable market (series stop to fluctuate because each buyer consistently chooses the same product). However, a large number of producers (33 out 50) makes positive sales, and this is not due to "errors" as in the previous example. This is the effect of the tolerance, which makes buyers consider equivalent products scoring slightly less than optimally in respect of a characteristic. In this case, buyers' use more than one characteristic to make a decision (4.3 on average in the latest steps), and the full ranking of characteristic does matter. In fact, consider a buyer's preferences:

$$< c_1, c_2, \ldots, c_m >, c_i \in \{1, 2, \ldots, m\}$$

With characteristic c_1 some of the available products are discarded, so that characteristic c_2 will be used to select only among the set of the remaining products, which possibly does not include the absolute best product in respect of c_2, because it has been previously discarded. If more than one product is still eligible for purchase, c_3 will be used to select among those, but only products equivalent to the maximum in respect to c_1 and c_2 will be considered. The behaviour of buyers, in this case, represents the fact that when the most important aspects of the product are not sufficient to reach a unique choice, buyers make their decisions on the base of less and less important aspects.

This configuration of the model generates many market niches, that is, groups of buyers who consistently choose a specific product. Note that the number of "niches" is much larger than the number of product characteristic used. In fact, the choice of buyers does not represents only the search of a optimal product (in respect of some characteristic), but it is a conditional optimum. These conditions varies largely, so that it can happen, for example, that two buyers make a purchase using the same characteristic as criterion, but choosing different products, because they have followed a different pattern (selecting different products with different characteristics) before resolving to use the same characteristic as decisive criterion.

The structure of market niches in this set up, with buyers having heterogeneous preferences, is very complex and changes sensibly even with slight modifications of the environment. For example, setting slightly different values for the tolerance level causes large changes in the number of buyers choosing a given product. In Figure 10 the simulation results shown above are modified at step 400 setting a slightly different level of tolerance: the structure of sales changes sensibly as a result.

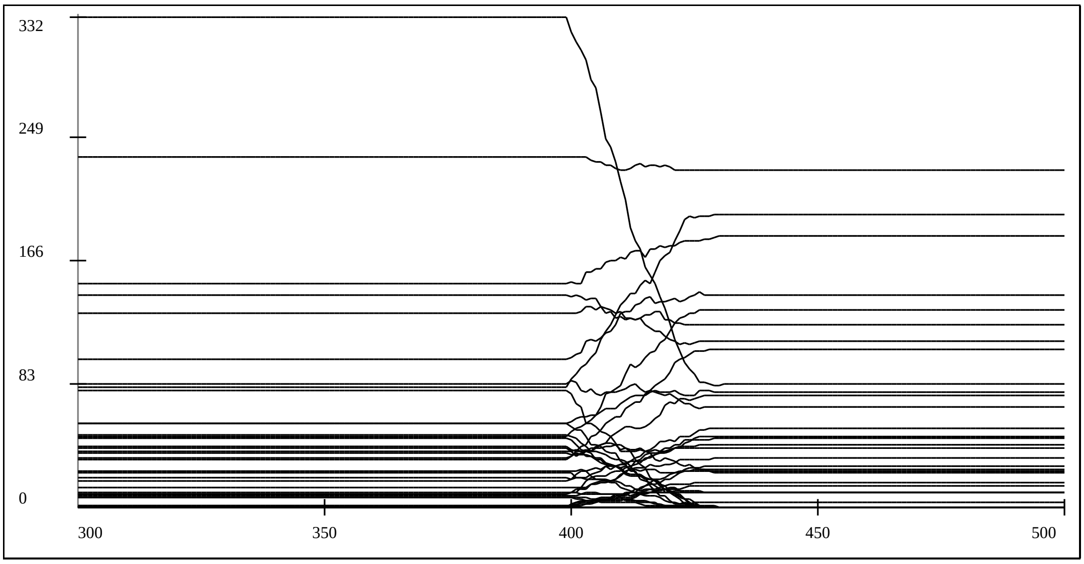

Figure 10: The same simulation run in Figure 9 is modified at step 400 changing slightly the level of τ. The structure of market niches change sensibly because now the decisional processes of consumers are modified.

Comparing the results obtained with the latest result with the ones obtained in the previous paragraph we can see that they share a strong similitude. In both cases we have the market formed by many firms making positive sales, but the reasons for this situation are strongly different. In the previous case we have consumers with similar preferences, chosen in a restricted set, that make mistakes, therefore allowing even inferior products to make positive, albeit small, sales. In this second case we have consumers with wildly different preferences that, with perfect insight and no errors, can choose the same product as a result of totally different reasons. In the first case sellers are part of a specific market niche, and their buyers choose on the basis of one single reason. In the second case, each buyer is a specific market niche, and the actual motivations for each sale is specific to the relative comparison between one product and all the others. This difference between the two cases have obvious consequences on the strategies that a seller willing to improve its position should take. In the first case only one characteristic is relevant, while in the second it is the whole product dimensions that matter.

3 Conclusions

This work proposes a general representation of consumers based on bounded rationality and evidence from cognitive studies. The decisional algorithm used to represent the consumer behaviour offers a natural definition of preferences as a ranking of the importance of products' characteristics. It is sustained that sellers can influence preferences by means of their marketing activities weighted with the levels of market shares, therefore provoking the partial endogeneity of preferences.

The model for the single consumer is used to build a model of demand for a novel product. Such model assumes initially that consumers are all identical before entering in the market, and allows to study the dynamics of market shares as determined only by the properties of demand.

The results obtained with such a model vary depending on the set of offered products and on few parameters determining consumers' behaviour. The simulation exercises show a good degree of qualitative adherence to the reality observed. Moreover, the results suggests that different types of market configurations can be linked to different types of demand. In particular, it is studied how different types of endogenous market segmentation emerge as a function of different assumptions on consumers' behaviours. Apparently similar aggregate patterns, showing many sellers with positive sales, are, in fact, produced by two totally distinct underlining dynamics, with obviously different consequences for possible R&D strategies trying to improve economic results through quality improvements.

This work is meant to show that an evolutionary theory of consumers can be constructed on the basis of the same approach used for the evolutionary theory of the firm. Moreover, the methodology used, computer simulations, ensures both a logically consistent development of the theory and the possibility to link this to other parts of the theory. For example, the model proposed here unrealistically assumes sellers to not modify the qualities of their products, for the obvious purpose to study the results as a function of demand-only motivations. A natural extension of this model is to introduce product-embodied technological innovation to allow for products' improvements.

References

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

Arthur, B. W. (1994), *Increasing Returns and Path Dependence in the Economy*, University of Michigan Press.

Aversi, R., Dosi, G., Fagiolo, G., Meacci, M. and Olivetti, C. (1999), "Demand Dynamics With Socially Evolving Preferences", *Industrial and Corporate Change*.

Bresnahan, T. and Malerba, F. (1997), "Industrial Dynamics and the Evolution of Firms' and Nations' Competitive Capabilities in the World Computer Industry", in R. Nelson and D. Mowery, eds., "Sources of Industrial Leadership", Cambridge University Press.

Cowan, R., Cowan, W. and Swann, P. (1997), "A Model of Demand with Interaction among Consumers", *International Journal of Industrial Organization*, **15**, pp. 711–32.

Gallouj, F. and Weinstein, O. (1997), "Innovation in Services", *Research Policy*, **26**.

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

Kahneman, D., Slovic, P. and Tversky, A., eds. (1982), *Judgment under Uncertainty. Heuristic and Biases*, Cambridge University Press.

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

Lancaster, K. (1966), "A New Approach to Consumer Theory", *Journal of Political Economy*, **14**.

Metcalfe, J. (2001), "Consumption, Preferences and the Evolutionary Agenda", *Journal of Evolutionary Economics*, **11**(1), pp. 37–58.

Nelson, R. R. and Winter, S. G. (1982), *An Evolutionary Theory of Economic Change*, Belknap Press, Cambridge, Mass. and London.

Richardson, G. B. (1996), "Competition, Innovation and Increasing Returns", working paper 96-10, DRUID.

Saviotti, P. and Metcalfe, J. (1984), "A Theoretical Approach to the Construction of Technological Output Indicators", *Research Policy*, **13**.

Saviotti, P. P. (2001), "Variety, growth and demand", *Journal of Evolutionary Economics*, **11**(1), pp. 119–142.

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.

Witt, U. (2001), "Learning to consume - A theory of wants and the growth of demand", *Journal of Evolutionary Economics*, **11**(1), pp. 23–36.