QUAERERE VERUM

Universidad de San Andres

DEPARTAMENTO DE ECONOMIA

Price Setting in a Schematic Model of Inductive Learning

D. Heymann (CEPAL, UBA, ITDT), R.P.J. Perazzo (Center for advanced studies), A. Schuschny (Center for advanced studies)

CICLO DE SEMINARIOS 1997

Día: Martes 21 de Octubre

9:00 hs.

Price Setting in a Schematic Model of Inductive Learning

D.Heymann †, R.P.J.Perazzo‡, and A.Schuschny‡

† CEPAL (Comisión Económica para América Latina, Un. Nations).
University of Buenos Aires, Faculty of Economic Sciences and ITDT.
Paraguay 1178, (1057) Buenos Aires, Argentina
‡ Center for Advanced Studies, University of Buenos Aires
J.Uriburu 950, (1114) Buenos Aires, Argentina

Abstract

We consider a schematic model of a decentralized market in which many agents with learning abilities, shop in one or two stores gessing in advance the price that the store will post. The stores in turn also adapt their pricing strategies in order to maximize their income. The inductive learning of the stores and the agents is represented by a genetic algorithm. We study situations of markets that are far from equilibrium and relax towards it through learning and adaptation. We find that the systems that we analyze may live long periods in metastable states that are better described as dynamic (as opposed to static) equilibria in which the evolution constrains the system to visit a reduced portion of the available configuration space.

1 Introduction

Economic models typically make strong assumptions about the ability of economic agents to process information. The underlying hypothesis is that individuals can somehow draw optimal plans, with an accurate knowledge of the systematic properties of the environment. This mode of analysis can produce interesting characterizations of steady states - where it can be assumed that the behavior of agents has evolved so as to become adapted to the environment -, but it hardly looks like a general approach to the study of economic processes. Since economic outcomes derive from a complex interaction of the decentralized decisions of large number of agents there is no a priori reason to postulate that individual behavior is optimal or, for that matter, that the coordination of activities takes place automatically [1]. There are reasons, therefore, to investigate analytical alternatives. One of them is to represent agents as capable of inductive learning [2], i.e. of revising their own "picture of the environment" based on the flows of information they receive. This type of analysis, in which agents are "boundedly rational", seems apt to generate useful results, in order to understand behavior patterns of economic systems which seem problematic for the conventional approach [3], [4].

Although an increasing volume of work is being carried out within that general framework, the search for commonly acceptable modelling procedures and for robust results is still in an incipient stage. The exploration of the properties of very stylized environments then appears to serve a meaningful purpose, as a preliminary step towards the representation of actual economic interactions. In this regard, the study of elementary processes of exchange seems a natural starting point.

When considered in some detail, even highly simplified systems of exchanges appear to require quite intricate representations. The competitive process through which, say, a large number of buyers and sellers of a certain good coordinate their decisions (as in the elementary supply-demand model) derives from the actions of agents based upon certain conjectures, which they revise according to the observed outcomes. The "law of one price" and a market-clearing configuration of prices and quantities would not be established without the activity of the agents. One may choose to focus on the end-state of the system, but the fact remains that this comes about endogenously, through the decisions of the agents and the effecets they have on those of others. While the system is in transition, ther decision problems of individuals become complicated, due to the information processing that must be undertaken to forecast the movements of the evolving environment.

In this paper we consider extremely simple systems of transactions. Although the emergence of specialized traders as "market makers" is a particularly relevant analytical issue, we focus here on the study of schematic pre-specified market settings, where "stores" are assumed to exist from the start. In the models we analyze how buyers with limited knowledge, endowed with certain learning procedures, interact with price-setting stores who must also find out the properties of the behavior of their prospective clients. The dynamics of the system are generated by the decisions of those agents on the basis of their learning.

The inductive learning of the agents is represented here by a genetic algorithm [5]. The "genome" of a buyer encodes the internal representation that the agent has of the pricing policy of the store(s) he deals with. Price forecasts are determined accordingly, and the purchasing decisioms follow from them. The stores are similarly endowed with a learning mechanism that is used to anticipate the behavior of buyers and to set prices. THe pricing policies are also encoded in a genome that is used to predict future sales for a given choice of prices. The learning process works through revisions of the internal representations as a response to the observed performance. We then analyze how (and if) the learning and co-adaptation of the set of adaptive agents makes the system relax towards an equilibrium state.

2 The three models

We work with three simple models. In the first one that we label **Model A**, we analyze the working of the competitive pressures that establish a single market price in a system where buyers search for low-price offers and suppliers choose a selling price. We consider a system composed by a (large) number N of buyers of a good and a seller, who obtains the good at constant costs (assumed for simplicity to be zero). There is an "outside source", where the buyers can avail themselves of the good at a given price p^*. Each trading "day", the store faces a binary choice of the price it will post, which can be set either at $p_>$ or $p_<$. Buyers decide (before observing the actual price) whether to visit the store or not. In case a customer does visit the store, he purchases one unit of the good at the posted price. Since we assume that $p_< < p^* < p_>$, a buyers will go to the store if and only if he conjectures that $p = p_<$.

We assume that the traders must learn about certain properties of the environment. In the model, customers are aware of the existence of the outside source selling at p^* (a value that remains constant throughout the exercise) but do not necessarily know the pricing strategy of the store. Here, the (strong) simplifying assumption that prices are chosen from a restricted set of values becomes relevant [1].

If there is a gap between p^* and both $p_<$ and $p_>$, and customers only care about the price averaged over repeated purchases, the equilibrium where all agents optimize would obtain at a state in which the store "randomizes" the price, so that there is no predictable pattern [2] and the mean value is p^*. This behavior seems rather artificial. We therefore concentrate on the case where p^* is slightly above $p_<$, implying that the equilibrium is established at the price $p_<$.

In any case, consumers must use a predictor to anticipate the price of the store and decide an action. We postulate that the predictor is a Boolean function of three inputs (the last three

[1] This assumption was made for the sake of computional tractability, given the use of Boolean funcitions in the learning devide

[2] It may be noted that, for the purposes of the model, the predictability of the price sequence is relative to the learning ability of the buyers (or the resources they invest in learning, as the case may be), not to the judgment of an outside observer.

observed prices) and one output (the expected price). Each predictor is encoded in a string 8 bits long. Agents choose one predictor from a set, whose elements are initally selected at random from the universe of possible predictors. Associated with a given predictor is a *strength*, which increases (decreases) whenever the action taken by the economic agent according to that predictor is a success (a mistake). The predictors of each agent are tested against the current price strategy of the store and are ranked by their strength. The next action taken by each agent follows the "better" predictor. i.e. that with the largest strength. After T_c time steps the predictors of each agent are updated with a genetic algorithm, so that they undergo crossover and mutation. The agents update their strategies asynchronically: in each time step, a given fraction of the population is allowed to invoke the genetic algorithm to update predictors. The next updating for this subset of agents takes place T_c time steps later.

The store also has a pricing strategy that is given by a Boolean function of four inputs (the past four prices) and one output (the next price). The *pricing strategies* play a similar role to the predictors of each client and are encoded in 16 bit strings. The store has a set of strategies and also updates them by using a genetic algorithm. After T_s time steps, the population of 16-bit strings is subject to crossover and mutation. The available pricing strategies are tested against the current predictors of a sample of clients. This is done during a session of T_v virtual trading time steps. The strategies are then ranked by the total profit (in this case: the value of sales) they generated during such rounds of virtual trading; the store then chooses as its strategy the one which produced the largest benefit. This process allows a co-adaptation between the strategies of the agents on both sides of the market.

In **Model B**, closely related to the previous one, we assume that the price p^* is intermediate between $p_>$ and $p_<$), and that customers must incur a *mobility cost S* when they shift from one supplier to the other. Clearly, if this cost is uniform among consumers and sufficiently high (i.e. larger than $(T_c - 1)(p_> - p^*)$), the store has "captive" clients, and would set a price $p_>$ in a profit-maximizing equilibrium. This is as if in **Model A**, p^* was higher than $p_>$. However, the asymptotic distribution of clients that shop at each supplier may now depend on the initial (random) conditions and on the particular path followed during the relaxation process. A richer variant of this simple model appears when it is assumed that the mobility costs can vary accross consumers, and p^* is near $p_<$. If there is a fraction of (actually) captive clients (due to their high switching costs), the store has to learn about the price elasticity of demand: depending on that fraction, the "optimal price" may be either $p_>$ or $p_<$, and this optimal price may itself be path-dependent.

In the third model, labelled C, customers do not face a mobility cost. Now we have two stores, each one with the choice between prices $p_>$ and $p_<$. These stores are confronted with N buyers searching for a convenient purcchasing strategy. We have thus a duopoly (if the outside source is irrelevant for the exercise), and the model can be used to analyze whether the stores will "learn to cooperate" (against the consumers), ot they will end up acting as competitors.

The predictors used by the buyers are now more complicated, since they must take into account the price patterns of the two stores: in this case, the predictors are Boolean function of six inputs (the three past prices of both stores) and two outputs (the expected prices for both stores). The genetic algorithm operates as in **Model A**, and the pricing strategies for the sellers have the same structure as before. The co-adaptation process - which also takes place during a session of T_v virtual trading time steps - takes place as both stores test their possible price strategies against the strategies of the other store and those of a sample of customers.

3 Results

Model A shows a strong tendency to converge to a "competitive equilibrium", where the seller posts a price $p_<$ and all the buyers shop at the store. Still, once the equilibrium is reached, occasional bursts of high prices can be observed. These are due in the model to the occurrence of random mutations in the genome of the store. During such bursts the seller takes advantage of the "surprise" caused to customers by the unexpected change in the strategy of the store. However, in this fundamentally stationary environment, the deviations from equilibrium are however corrected in increasingly shorter periods, since the clients improve their predicting accuracy as the number of trading periods increases and buyers revise their forecasts.

Although the system is definitely attracted towards the equilibrium, it does not follow a simple convergence path. During the transition, there are cases where the market spends a considerable number of periods in intermediate states where the seller alternates between low and high prices, and the buyers learn to adapt to such alternation. This is seen in the bursts shown in figure 1(a), where the number of "wrong" clients declines over time while the store maintains a pattern of alternating prices. Another example is shown in figure 1(b) in which the metastable situation is magnified on one side of the picture, to indicate in more detail the type of transitional states which can result from the co-adaptation of the strategies of the seller and the clients. The metastable states can last for a sizeable number of trading periods before the equilibrium is approached. The lifetime of such coadapted states clearly depends upon the size of the fraction of the population that has not learned its features.

The appearance of those metastabilities makes it difficult to establish in advance the duration of the relaxation process. There are two extreme situations: $T_s \ll T_c$ and $T_c \ll T_s$. In both cases, the length of the relaxation period is governed by the the largest of the "strategy-revision" intervals. If $T_s \ll T_c$, the store adapts its pricing stratregy to a population of clients most of which are acting on the basis of wrong predictors. The store can then obtain high profits, with a significant fraction of the population of buyers making mistaken choices. In the opposite case, the coadaptation is governed by the clients. The system then shows successive metastable states, as the clients adapt to new price-fluctuating strategies of the store. In the intermediate situations in which $T_c \simeq T_s$ the relaxation process is more complicated. This can be seen in the examples of figure 1. In figure 1 (a) we show a simulation in which $T_c \ll T_s$ while in figure 1(b) we have taken $T_c = T_s$.

The population of clients does not behave as an average representative agent. Different agents adapt at different speeds and they also revise their strategies asynchronically. The diversity among the current strategies within this population helps to understand the form of the relaxation process. If each client is labelled by the strength of its current predictor, it is found that the initial random (gaussian) distribution evolves into a bimodal one, with one peak of well adapted agents, and another of poorly adapted ones. These two sharply differenciated populations finally merge into a single group of "informed" traders as the equilibrium is reached.

The characteristics of the relaxation process are robust with respect to changes in the parameters. Clearly, the static equilibrium is stable; this fact is verified numerically by the short transients that occur after a random change in the price of the store once the equilibrium is reached. In addition, the movement towards equilibrium takes place even when the agents or the store are endowed with longer "memories" of past prices and incomes in order to make predictions. However, the relaxation process is altered because the collective search has to cope with an exponential increase of the alternatives which are considered when determining the strategies. For the same reason, metastable configurations grow both in number and complexity. More "sophisticated" learning procedures, therefore, need not result in a quicker or more "direct" convergence to equilibrium.

In **Model B**, the search undertaken by customers is made more expensive through the introduction of the "mobility cost". The asymptotic distribution of clients buying in each store now depends on the (random) initial conditions and on the particular path followed during the relaxation process. As could be expected, when the switching cost S becomes very large, the selected strategy of the store is to fix the price at $p_>$ (figure 2). However, there is a lower critical value S_0 of the mobility cost, given by $(T_c - 1)(p_> - p^*)$ -or, equivalently $(T_c - 1)(p^* - p_<)$ -, below which there are not captive clients, because buyers then have incentives to make the "investment" of shifting away from a high-price store.

This simple model can be extended by introducing some heterogeneity among agents, in the form of different mobility costs. In this extension, the population of potential customers is divided into two groups: one with a low (zero) cost of search, the other with a high cost. Given the "planning horizons" of clients, if these were fully informed, the latter group would represent a set of captive customers, while the first one would only buy at the store at the price $p_<$. The store thus faces a downward-sloping demand curve (a curve which varies over time, since the mobility cost applies both when a buyer switches from the store to the outside source and viceversa); it must learn about its elasticity, which depends on the fraction of buyers with high and low mobility costs.

Here too, the behavior of the system is path-dependent. Figure 3 shows the distribution of some variables over many experiments, for different values of the proportion of the population of buyers that has high mobility costs. There can be seen that, when this parameter is small, in most of the experiments the co-evolution of the strategies of agents makes the price set by the store go to $p_<$: the distribution of the price has a well defined peak at that value. For very low values of the parameter, the store often ends up selling (at the price $p_<$) to a large majority of the potential customers. As the fraction of high-mobility-cost clients become higher, the distribution of posted prices tend to become bi-modal, with a sizeable number of experiments where the average price is "high" (only slightly lower than $p_>$): here, the store has "learnt to exploit" the set of captive customers, even if this means losing the low-cost customers to the outside store (as shown in the distribution of the average number of actual buyers at the store).

The simple model with *mobility costs* can be used - in a naive fashion - to study the possible emergence of a *fix-price strategy*. This may be done by allowing a population of clients (with the same switching costs) to evolve in the presence of a store whose prices fluctuate randomly between $p_>$ and $p_<$, and a second store with a fixed price $p^* = \frac{p_> + p_<}{2}$. Although clients are seen to "prefer" the fixed price instead of the fluctuating one, the preference shows not to be robust in this case. This is so because the agents of the model are only sensitive to time-averaged prices, and therefore do not properly penalize the uncertainty about future prices (or the costs associated with processing data).

In **Model C**, there are two stores, competing with one another and (possibly) also with an outside source. When this outside source has a price between $p_>$ and $p_<$ (say, the average of these two values), the system admits a static equilibrium where both stores set the price $p_<$ and share the population of clients. However, given the complexity of the co-adaptation of the strategies of buyers and sellers, the system does not simply relax to that state. In addition, the steady state, once reached, seems to be easily upset by sudden changes in the price of one store due to random mutations in the genome. In figure 4 we show two simulations that differ in the size of the price gap $p_> - p_<$. When the gap is big (figure 4(b)) the system quickly reaches a configuration showing a pattern of prices that vary over time: either both stores "take turns" in setting the high price, or one store sets the price $p_>$ while the other produces high-frequency oscillations. In such situations, the income of both stores is relatively high, because many clients atre kept operating with inappropriate predictors.

The behavior of the model can be studied also in a time- independent framework, indicating the frequency with which the state visits different regions of the relevant spaces (figure 5). It can be seen that the system spends most of the time in a quite restricted area of the C plane, performing large oscillations along the straight line $I_1 + I_2 = Np_<$. This indicates that, while the total income of the stores does not vary much, the search for price strategies induces wide movements in the relative incomes of both suppliers. Also, the symmetric distribution of customers is frequently visited, but, because some customers go to the outside source. Some highly asymmetric distributions of clients are observed in some instances: this occurs because random fluctuations in the number of clients at one store are not rapidly reversed. In any case, although the time average of the number of clients generally does not satisfy, for the average incomes of the stores, holds with a considerable degree of approximation. In this regard, the system would seem to go to a dynamic equilibrium, where the state variables are not constant over time, but their evolution is such that some properties (symmetries and time averages) are preserved, and the configuration remains statistically confined within a small portion of the feasible space.

If the price of the outside source is higher than $p_>$, we have in fact a duopoly, given that informed customers would always choose to visit one of the stores. One could choose to view this setting as one where the stores are engaged in a strategic interaction with one another. It is known that repeated duopoly games have no unique, well defined equilibrium when the players are assumed to know the strategies that all of them are playing. Here, the "payoff landscape" is evolving as a consequence of the co-adaptation of the agents, so that the way in which the system performs can be expected to depend on the way in which the agents go about learning about its features.

When there is a large gap between the high and low prices ($p_>$ and $p_<$), it is found that most of the time the stores post $p_>$, with a distribution of clients that oscillates with a sizeable variability, but remains on average around an even share of sales. It is as if the stores were "cooperating" instead of competing with one another. Occasionally, one store switches to the low price (see figure 6), but it eventually revises its strategy. That is, one finds in this example some instances of "price wars", but they do not last, and they are not sufficient to generate the competitive equilibrium as an outcome. This result is probably generated by several features of the model. The binary choice of prices implies that, in order for $p_<$ to be "perceived" as a profitable alternative by the price-setting algorithm, the response of the customers must be quite fast: since the clients show a delay in their reaction, and the store is "forced" (by the exogenous determination of $p_<$) to lower the price significantly, if it does reduce the price at all, then it may lose in terms of the value of sales. In other words, the implicit coordination among the sellers appears to be induced by the combination of a discrete set of widely different possible prices and the lag in the adjustment of buyer's strategies, which results in a relatively low "demand elasticity" as seen by the firms during the time interval that they consider for the purpose of establishing the pricing strategy.

The dependence of the results on the parameters can be checked when the gap between $p_>$ and $p_<$ is lowered (see figure 6). Here, too, the distribution of customers varies around a 50% proportion in each store. However, although the price $p_>$ is often observed to be posted by the stores, the system shows considerable price oscillations, and no simple pattern appears to emerge after a number of trading "days" as long as 3000. In this case, the fact that $p_>$ is higher (relative to $p_<$) strengthens the "temptation" for the stores to engage in price competition in order to attract customers starting from a configuration where both sellers set $p_>$. But also, conversely, the learning lag of customers makes the competitive equilibrium easy to disturb through price increases on the part of sellers. The system thus shows a dynamics which does not settle down to a well defined steady state.

If the fitness of the strategies of both stores is measured only by the number of clients rather than by the income, the static equilibrium is rapidly reached. This is shown in figure 7 in which we plot the frequency with which the system visits each point (C_1, C_2), when $p_< < p^* < p_>$. In a brief first stage of the relaxation process (that can hardly be read in the figure) the outside source (selling the good at p^*) looses clients. In a second stage both stores share evenly the clients among them. As we mentioned above this equilibrium is brittle. If a small perturbation is produced in the number of clients, the clients lost by one store shop at the other and the one that is loosing clients has no strategy at its disposal to correct this fact. In addition the correct strategy of keeping its price at $p_<$ looses fittness and is replaced by another that makes the situation even worse causing that the optimal strategy is found again only after large fluctuations in the clientelle have taken place. This can be seen in the figure as an extended set of peaks that are lined up along the straight line $C_1 + C_2 = N$ that are nevertheless very narrow in the direction perpendicular to that line.

If the fittness is measured by the income instead than by the number of clients a further degeneracy occurs because the same income can be obtained with less clients and with a higher price. The number of alternative strategies is therefore larger making more difficult the process of reaching the static equilibrium. When the gap $\Delta p = p_> - p_<$ goes to zero (as shown in fig 5(a)) the dynamical situation approaches the limit in which the utility function depends only upon the number of clients and not of the income. Correspondingly the static equilibrium is more easily found. When $\Delta p > 0$ the equilibrium $C_1(t) \simeq C_2(t) \simeq N/2$ no longer holds because there is a greater possibility of compensating a loss of clients in favour of the external store by setting a higher price (see fig 5).

Inspite of all those internal degeneracies the system spends most of its time in a rather restricted region of the C_2 vs C_1 plane performing large amplitude oscillations along the straight line $I_1 + I_2 = Np_<$ (see fig. 5). The values that correspond to the static "equilibrium" ($C_1 \simeq C_2$ is frequently visted but $C_i < N/2$ because the external store is always visted with some frequency. There are also other highly asymmetric distributions of clients ($C_1 \leq N$ and $C_2 \simeq 0$ or $C_2 \leq N$ and $C_1 \simeq 0$). This is so because an accidental fluctuation in the number of clients has very low negative feedback causing one of the two stores to loose rapidly all its clients in favor of the other. The region of the plane that is mostly visited by the system depends upon the value of the "external" reference price p^*. For $p^* \simeq p_<$ the outside store is visited frequently while for $p^* \simeq p_>$ all the clients are shared by the two stores.

The seemingly coordinated ("oligopolic") actions of both stores is a consequence of the symmetries of the model that force both stores to optimize the same utility function facing similar environments. The dynamic exchange clients has very stable time average values. Indeed in all numerical simulations we have found that inspite of the fact that the time averages of the number of clients *do not* in general fulfill $< C_1 > \simeq < C_2 > \simeq N/2$, the average incomes instead fulfills to a high degree of appproximation $< I_1 > = < I_2 > = p_< N/2$ that is the value that corresponds to the static equilibrium. As long as time average values are considered, the sort of time evolution that is obtained displays a large degree of indifference with respect to the static equilibrium. We are thus inclined to consider it as a *dynamic equilibrium* i.e. one in which the relevant parameters are not constat in time but its evolution is such that some properties (symmetries and time averages) are preserved while the system satistically visits a reduced portion of the available configuration space.

The stability of the equilibrium can also be studied with the addition of a very small "mobility cost". As expected, its robustness is increased impairing the large fluctuations of clients between one store and the other. This however happens at the price of a strong "path dependence". The population of clients tends to freeze into an uneven division between both stores that depends upon the particular evolutionary path followed by the system in each case.

It is instructive to study the model as a non cooperative game played by both stores in which each "move" corresponds to choose a set of prices for the next few days (i.e. to fix a price strategy). The pay-off matrix is then determined by the income obtained through the combined effect of all the current predictors of the whole set of clients. There are two extreme situations. The simplest one is one in which it is assumed that all clients have unlimited learning capacity, i.e. that fully learn the strategy of the stores instantaneously. In this case the pay-off matrix is static and is obviously independent of the past history of the system. In addition, since the model is fully symmetric with respect to both stores, the pay-off matrix seen by each store is the traspposed of the other. This game has an obvious static equilibrium situation in wich both stores set their prices at $p_<$ and share 50 % of the population of clients.

The second situation is one in which one takes into consideration that the current predictors of each client change as a consequence of their adaptation. In this case the pay-off matrix can be evaluated numerically and should be considered as a dynamical variable. The symmetry between the pay-off matrices seen by both stores is now lost. In these conditions the "game" seems actually to be based upon a moving "payoff landscape" instead of a common payoff matrix. In practice this means that both stores actually play different games and is therefore less likely that a common strategy based upon a common perception will emerge. Note the difference with the case of a mixed strategy when no dominance takes place. In our case the pay-off matrix changes all the time and the relaxation to equilibrium corresponds to a process in which both pay-off matrices approach a common value given by the static value mentioned above.

In figure 8 we show the minimax strategy of both stores obtained numerically from the "real market" of all clients every a fixed number of time steps T_s i.e. every time that the strategies are updated through the genetic algorithm. We show a case in which the utility function is given by the number of clients and one in which the utility is given by the income. As discussed above, in the former case a static equilibrium is rapidly found wile in the latter the minimax strategies for each store are very different yet complementing each other to a high degree. In order to simplify the graph the strategies correspond to a future of only three days and are labelled by the average price for that period.

4 Conclusions

We have considered several schematic models to study de stylized features that may be found in markets that are far from equilibrium. We have shown how even in very simple market environments. learning and co-adaptation induce non-trivial dynamics. In the simplest settings. the interaction between the strategies of "intelligent" agents leads to a stable equilibrium but can also induce the appearance of meta-stable states, making it hard to guess in advance the duration of the relaxation process.

In the present models equilibrium is reached when the agents are able to coordinate their actions using only global information (the price set by the store) that is in turn the cooperative result of many individual actions. The fact that all agents end up by using the same strategies may seem at first sight not surprising because all agents are optimizing the same utility function and have also the same information. This argument is however not acceptable because the same type of information and processing is not enough when an additional store is added to the system. In addition, the random and asynchronous nature of the adaptive process causes that during relaxation towards equilibrium. the learning process proceeds differently for different agents. We observe that this produces clearly differentiated subpopulations of highly and poorly adapted agents. A byproduct of the present analysis is therefore to cast doubts on the validity of economic models involving a single representative agent. that attempt to describe the transient process towards equilibrium.

The limitation that is placed upon the price to have only two possible alternatives is mainly made to place no limiting assumptions upon the kind or type of predicting strategies of both the store and the clients and keep nevertheless the computation tractable. The model is therefore *not biased by any a-priori assumption upon the algorithm used to determine future prices.* During the relaxation process the system can explore in principle *all* the functional space of predicting alternatives. In spite of this exponential explosion, the evolutive process of updating and improving the strategies seems to provide a rather robust model to describe the learning process leading to a global optimum. One price that has to be payed by the freedom assumed in the present model is to place no cost to (unrealistic) strategies in which the price is repeatedly changed from $p_>$ to $p_<$. On the other hand. from a purely qualitative point of view, because of this same fact the model is able to "predict" the (realistic) situation of a coadaptation between the store and the clients that gives rise to increasingly complex metastable (dynamical) equilibria that impair the relaxation to a final, static configuration.

In a more realistic situation the "complexity explosion" that is already present in this simple model makes matters worse. The evolutive process can only involve minor changes in each step in order to improve the existing population of predictors. Major changes in the pricing policy or in the purchasing strategy leading to genomes that may be optimal in the long run are doomed to be discarded by the environment. Obtaining an optimal strategy therefore implies an (unrealistic) exponentially long evolutive process. After a finite training period we should therefore expect to obtain a robust albait *suboptimal* strategy.

The gap between this very schematic market model and more real situations of learning and adaptation has therefore to involve two different aspects. On the one hand the process must involve additional learning strategies. One possibility is to assume that the agents share an internal model of the behaviour of the store (a "theory of the mind" of the store that may be implicit in the cultural background of the population) that allows in each moment to discard offhand a large fraction of unrealistic strategies and predictions. This internal representation may also perhaps in turn be thaught of to undergo a slow adaptation. A second alternative for the learning process - that is absent in the present model - is to assume an additional coordination mechanism such as the transmission of "local" information (as opposed to the "global" information provided by the price posted by the store). This could for instance be about how well neighbouring agents have performed. The exchange of this information opens the possibility of imitating each other strategies.

A second important aspect that has to be considered to bridge the gap with more realistic situations is to think that a realistic system does *not* in fact reach a stable equiliubrium but rather spends most of its time in slowly fluctuating, coadapted situations explained above in which the store and the clients actually complement their actions keeping the system in a metatstable state that has the advantage of preventing random and therefore unpredictable fluctuations.

All the elements that are sufficient to give rise to an equilibrium in the simplest models prove inadequate to produce a coordinated stable state when there are two competing stores. This latter case shows no obvious convergence to a fixed state. The system engages instead in regularly changing configurations with long term fluctuations that have a robust average properties. This situation parallels that found in the metatstable equilibria of model A in which a large fraction of the agents learn a particular price strategy of the store and traps the system in a pattern of regularly changing configuration that may last for a long time.This suggests that the concept of a static equilibrium appears to be exceedingly stringent and a wider concept of a *dynamical equilibrium* should perhaps be more conveniently used to encompass both situations.

The toy market model that we have examined allow to extract some other side conclusions. The first lies in the fact that the results obtained adding a mobility cost to Model C are of no use to justify the appearance of a "fix price" strategy. This may lead one to conjecture that the advantages of such strategy do not derive mainly from the existence of such costs, but rather from "learning costs", i.e. from the fact that buyers may accept to pay a higher average price in order to save on the trouble of engaging in an expensive learning process when the price fluctuates.These costs that have been largely disregarded in the literature, may also provide an explanation of the persistence of suboptimal strategies. The second is the analysis that has been done casting the model into a non cooperative game. The main conclusion in this aspect is that the relaxation processes that one can expect to arise from learning and adaptation, may be difficult to cast into the language of the theory of games mainly due to the dynamical change of the pay-off matrices that cause the players to actually see and play different games.

The simple models shown in the paper are not intended to represent actual market conditions: they are exploratory constructions to obtain indicative results. However, most of the stylized facts that occur in this very schematic framework are realistic enough to draw relevant guidelines to analize more real situations. The outcome of the exercises suggest that in real market situations cooperative actions may derive from rather simple and schematic individual actions. By the same token, complex behaviours of the system should be considered as emergent properties of the ensemble of interacting agents considered as a whole.

References

[1] Arthur, W.Brian (1992). *On learning and Adaptation in the Economy*. Working Paper SFI, 92-07-038.

[2] Arthur, W.Brian (1994). *Complexity in Economic Theory: Inductive Reasoning and Bounde Rationality*. AEA Papers and Proceedings, **84**,2 , May.1994 , 406-411.

[3] Sargent, T.J. (1993). *Bounded Rationality in Macroeconomics*, The Arne Ryde Memorial Lectures. Clarendon Press - Oxford.

[4] Kirman, A. and Salmon, M. (1995). *Learning and Rationality in Economics*. Blackwell Publishers Ltd.

[5] Goldberg, D.E. (1989). *Genetic Algorthms in Search. Optimization. and Machine Learning*, Addison Wesley Pub.Co.

[6] Kauffman, S. A. (1993). *The Origins of Order*, Oxford University Press.

Figures

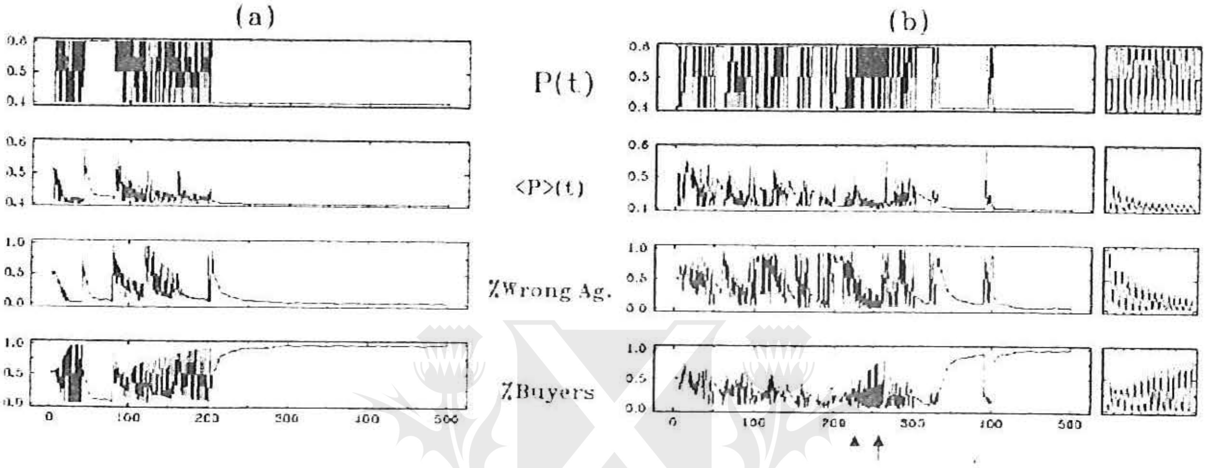

Figure 1: Two simulations of **Model A** with a universe of 400 buyers, each with a set of 16 predictors; the store has a set of 32 predictors. The two prices are $p_> = 0.60$, $p_< = 0.40$ and $p^* = 0.41$. Upper plot: price set by the store, second plot: price averaged over the 400 buyers, third plot : % of buyers that predicted a wrong price, fourth plot: % of buyers that decided to buy at the store. The evolution takes place during 500 time steps.The updating times are for figure **(a)**: $T_s = 60, T_p = 30$, $T_c = 10$; for figure **(b)**: $T_s = 10$, $T_p = 30$, $T_c = 10$. A magnification of a metastable regime indicated by the two arrows at the bottom is shown in the right part of the figure.

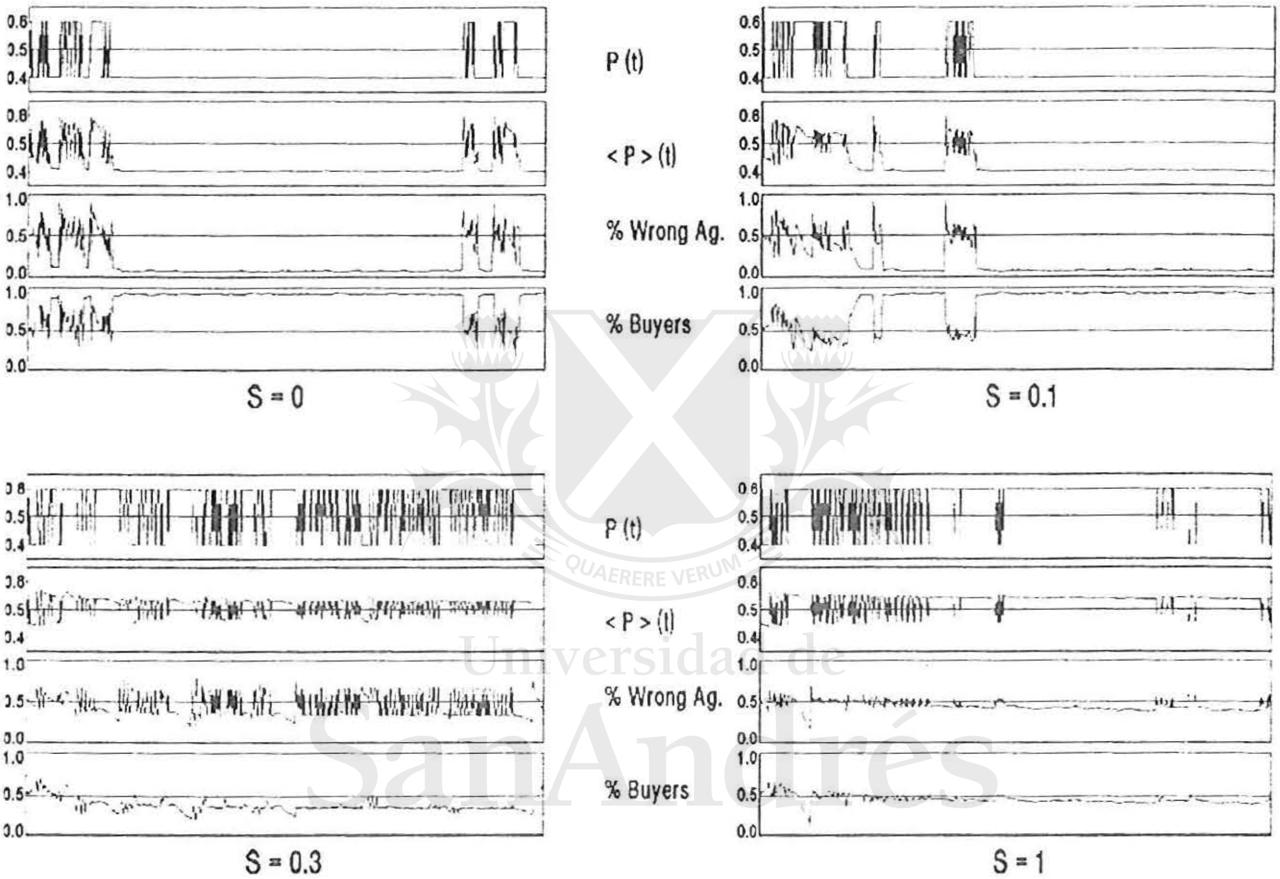

Figure 2: Simulations of **Model B** with increasing values of the *mobility cost* (S). The simulation universe is the same as in figure 1. The titles have the same meaning as in fig1.The value $S = 1$ gives rise to a population of captive buyers in accordance with the estimate given in the text. The reference price is $p^* = 0.41$, and the updating times are $T_s = 10$,$T_v = 30$, $T_c = 10$.

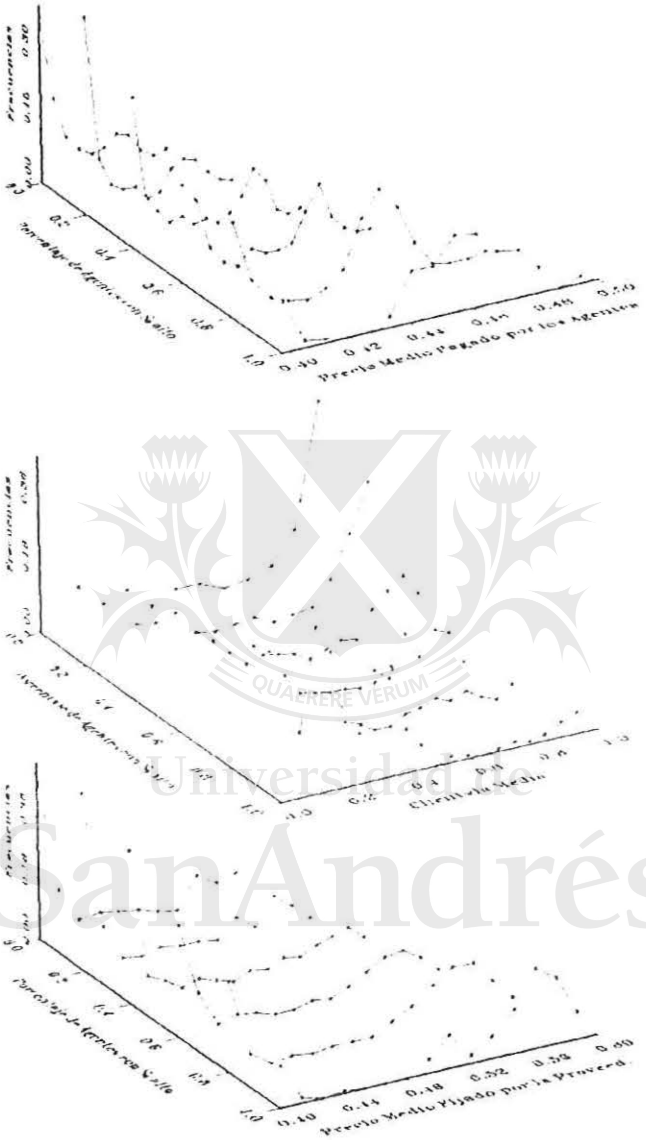

Figure 3: Distributions of the average prices paid by the agents, of the buyers and distributions of the prices that the stores fix as a function of the % of buyers with high "mobility cost" ($S_o \to \infty$, $p_> = 0.60$, $p_< = 0.40$ and $p^* = 0.41$.

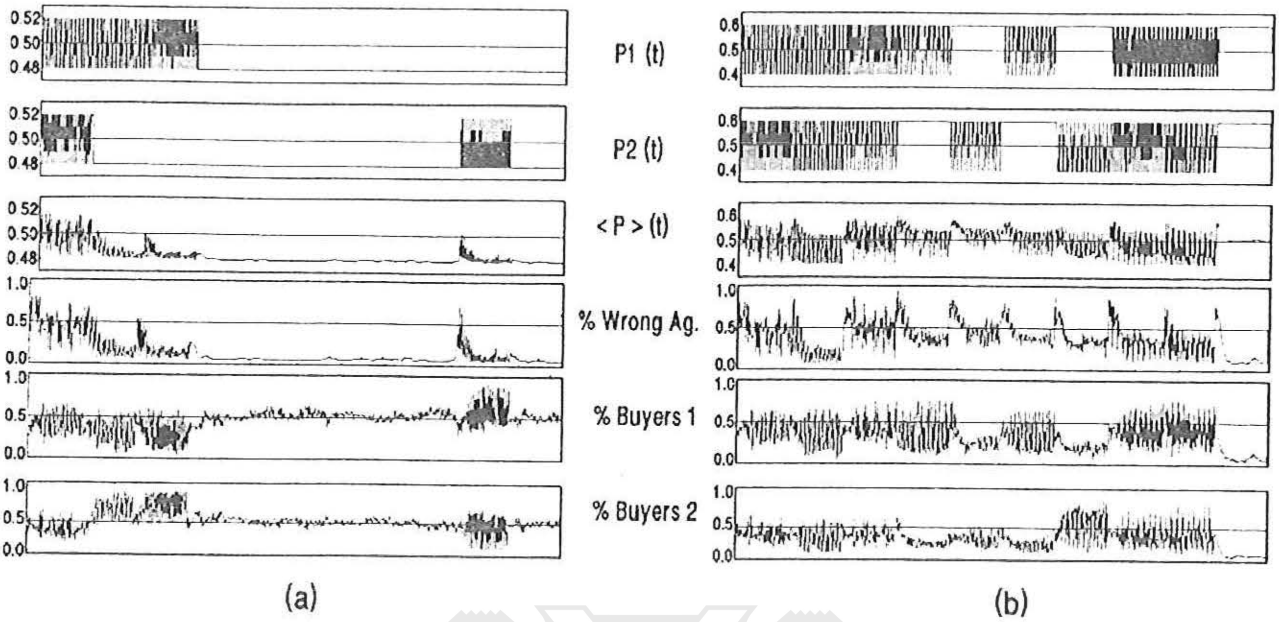

Figure 4: Two simulations of **Model C** with a universe of 48 buyers each with a set of 32 predictors. The stores have 32 predictors each. The titles have the same meaning as in figure 1. The updating times are $T_s = 60, T_v = 30, T_c = 10, p^* = 0.50$. Both prices are for figure (**a**): $p_> = 0.52$ and $p_< = 0.48$, for figure (**b**): $p_> = 0.60$ and $p_< = 0.40$

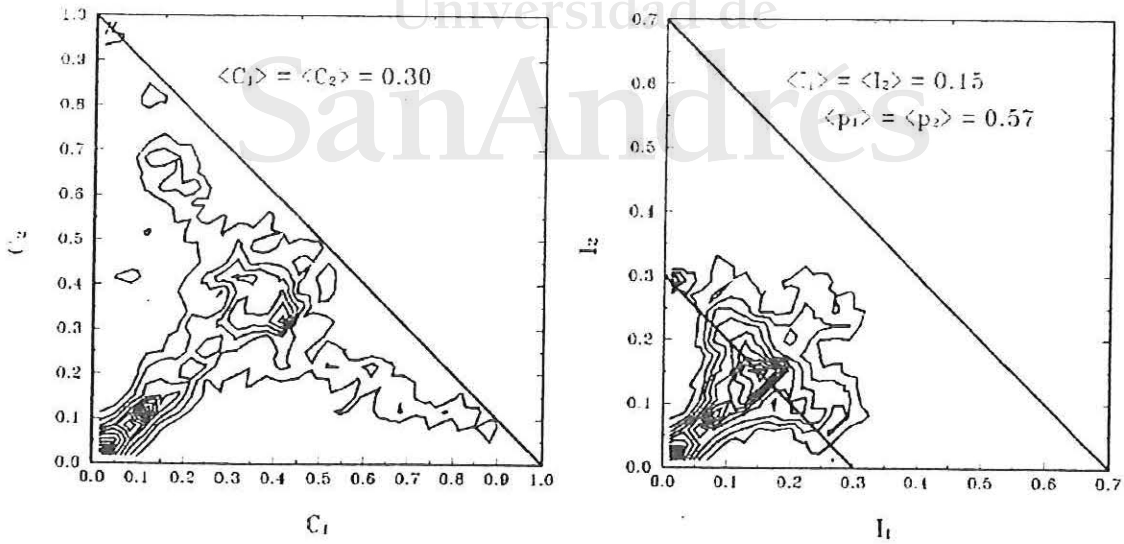

Figure 5: Contour plot of frequencies in the planess (C_2, C_1) and (I_2, I_1) with $p_< = 0.30$, $p_> = 0.70$ y $p^* = 0.50$. During the evolution both stores tend to optimize their profit function that depends only upon their respective income I_i).

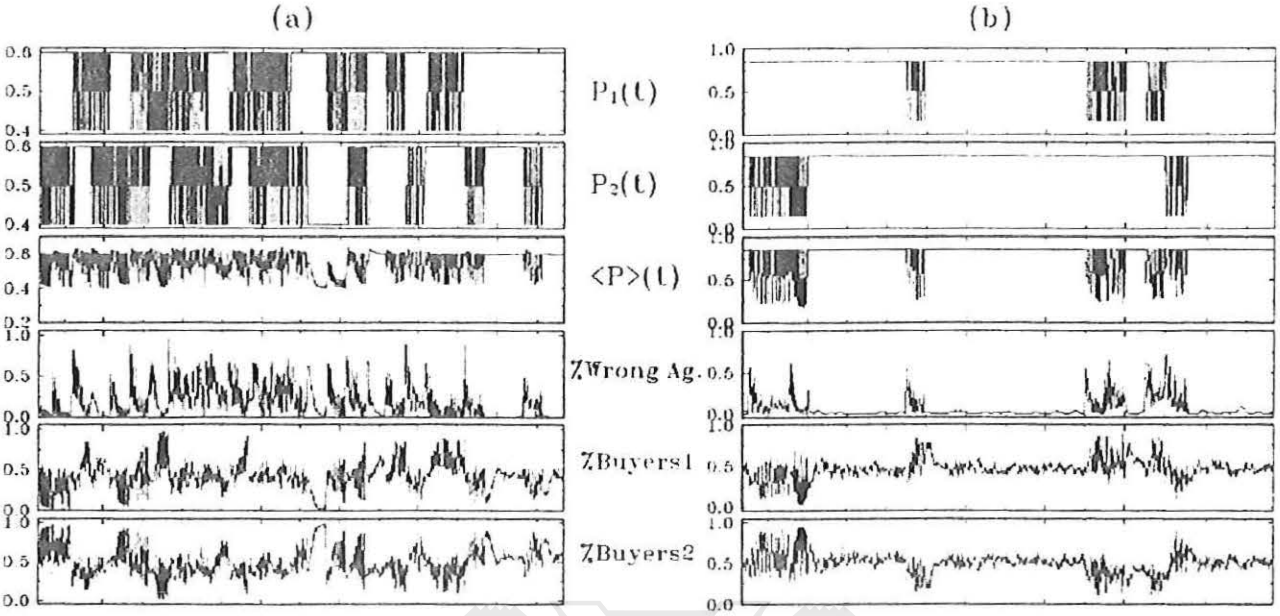

Figure 6: Two simulations of **Model C** with a universe of 48 buyers each with a set of 32 predictors. The stores have 32 predictors each. The titles have the same meaning as in figure 4. The updating times are $T_s = 60, T_b = 30$. $T_c = 10$. $p^* = 0.90$. Both prices are for figure **(a)**: $p_> = 0.60$ and $p_< = 0.40$. for figure **(b)**: $p_> = 0.85$ and $p_< = 0.15$

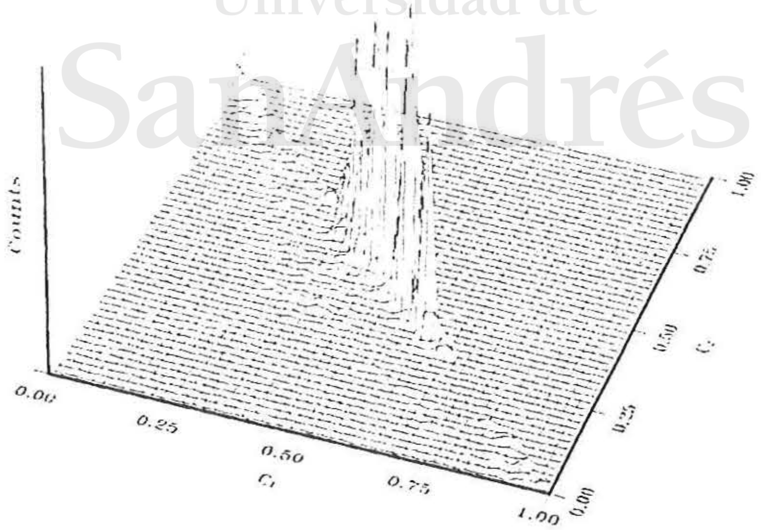

Figure 7: Frequency with which the system composed by two stores and 100 agents visits a point (C_2, C_1). During the evolution both stores tend to optimize a utility function that depends only upon their respective number of clients C_i.

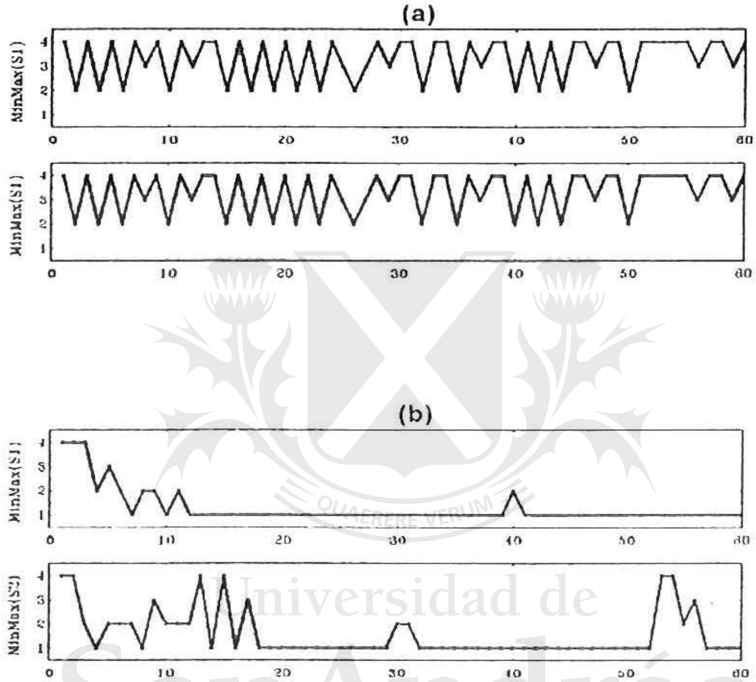

Figure 8: Minimax strategy of both stores as a function of time. Upper part (a): the utility function of both stores depends on the total income. Lower part (b): the utility function of both stores depends only upon the number of clients. The value of the minimax strategy is given by the average price $p_{av.}$ set by each stores for the next three days. The possible values of $p_{av.}$ that are labelled as 1,2,3 and 4 in the ordinate axis of both figures are $p_<$; $\frac{2}{3}p_< + \frac{1}{3}p_>$; $\frac{2}{3}p_> + \frac{1}{3}p_<$; and $p_>$.