Consumers' Behavior and the Bertrand Paradox: An ACE Approach*

Xavier Vilà†
Universitat Autònoma de Barcelona

28th October 2005

Abstract

We analyze the classical Bertrand model when consumers exhibit some strategic behavior in deciding from which seller they will buy. We use two related but different tools. Both consider a probabilistic learning (or evolutionary) mechanism, and in the two of them consumers' behavior influences the competition between the sellers. The results obtained show that, in general, developing some sort of *loyalty* is a good strategy for the buyers as it works in their best interest.

First, we consider a learning procedure described by a deterministic dynamic system and, using strong simplifying assumptions, we can produce a description of the process behavior. Second, we use finite automata to represent the strategies played by the agents and an adaptive process based on genetic algorithms to simulate the stochastic process of learning. By doing so we can relax some of the strong assumptions used in the first approach and still obtain the same basic results.

It is suggested that the limitations of the first approach (analytical) provide a good motivation for the second approach (Agent-Based). Indeed, although both approaches address the same problem, the use of Agent-Based computational techniques allows us to relax hypothesis and overcome the limitations of the analytical approach.

Key words: Agent-Based Computational Economics – Evolutionary Game Theory – Imperfect Competition.

1 Introduction

As Hehenkamp (2002) points out, the Bertrand model of duopoly relies upon two important assumptions regarding consumer's behavior: (i) consumers search for the cheapest firm at price zero and, (ii) switching firms has to be free. Nevertheless, not many papers explicitly take into account consumer behavior when analyzing the competitions among firms. Harrington and Chang (2005), and Ferdinandova (2003) analyze to what extend consumer's behavior can influence market dominance and, thus,

*Financial support from grants BEC2002-02130 of the *Dirección General de Investigación* (Spanish Ministry of Science and Technology), SGR2001-0164 of the *Direcció General de Recerca* (Generalitat of Catalonia), and CREA (Barcelona Economics) are gratefully acknowledged.

†Departament d'Economia i d'Història Econòmica. Universitat Autònoma de Barcelona. 08913-Bellaterra (Barcelona). Spain. `Xavier.Vila@uab.es`

can contribute to "shape" the competition among firms. More related to our work, Hehenkamp (2002) shows that the "sluggishness" (or not) of consumers when searching for the best price can critically influence the outcome of the classical Bertrand model, that is, depending on how consumers behave firms could, in equilibrium, set prices equal to marginal cost or higher. These three papers use different evolutionary techniques for their analysis. Alchian (1950) already suggested that evolutionary models could suit well the dynamics underlying competition among firms.

Here we also present an evolutionary analysis of firm competition when consumers are allowed to exhibit some strategic behavior when deciding from which one of the two sellers they will buy. Moreover, as discussed below, we will use to different approaches to the problem: an analytical one and another using Agent-Based Computational techniques. We will take into account not only the behavior of consumers when searching for the best price but also the role that a positive cost of switching firms might have. Finally, as opposed to the model by Hehenkamp, we will assume that the competition among firms is modelled as a repeated game which, to out understanding, fits better to the social evolutionary dynamics.

The model will consist of two sellers, offering the same homogeneous product, that compete against each other in a market with a given number m of potential buyers. At each period t $(t = 1, 2, \dots, T)$ the market opens (and closes) a fixed number of times R that we will refer to as rounds. It is important to note the difference between periods and rounds. Each period (for instance, a week) contains R rounds (for instance, 7 days). The two sellers decide at the beginning of each period a strategy that specifies the course of action to take at each of its rounds contingent on what happened in previous rounds of the same period. Thus, at the beginning of each round, each of the two sellers sets a price for his product according to their strategies for that period. Then, the market opens and the two prices become known to everybody. The m buyers walk in and each of them decides from which of the two sellers will buy based on the observation of this two prices. Trade takes place and a new round begins. After R rounds (end of the period) both sellers update their *per period* strategies according to some procedure that will be specified.

The idea behind this sequence of events is that each seller uses the R rounds to learn about the behavior of its competitor. Thus, each period represents a step in a learning procedure. The model can also be thought as a evolutionary model where each period represents the end of an old generation of sellers and buyers and the birth of a new generation that inherits the characteristics of its parent generation with, hopefully, some improvements. We will refer to the two intuitions indistinctly.

As said above, we will approach the resolution of the situation sketched above using two related but different tools. The two of them consider a probabilistic learning (or evolutionary) mechanism and in the two of them we will discuss how the consumers' behavior can affect the competition between the sellers.

First, we consider the case in which the learning procedure can be described by a deterministic dynamic system that uses expected values based on the work of Nowak et al. (2004). Using strong simplifying assumptions we are able to solve this case and to produce a complete description of how the learning process behaves. We also discuss the problems that one might face when trying to relax some of the assumptions made.

The second approach is an instance of Agent-Based Computational Economics techniques based on the analysis of the repeated Prisoner's Dilemma using Genetic Algorithms by Miller (1996). We use finite automata (encoded as binary strings) to represent the strategies played by the sellers (and also by the buyers) and a decentralized adaptive process based on the models of genetic algorithms to simulate the stochastic process of learning or evolution. With this technique we can relax some of the strong assumptions used in the first approach and still obtain the same basic results. Additionally, as a methodological agent representation issue, we modify the standard operators used in genetic algorithm techniques to make them more suitable to social (or economic) simulations.

Finally, we like to think that, to some extent, the limitations of the first approach (analytical) provide a good motivation for the second approach (Agent-Based simulations). Indeed, although both approaches address the same problem, we show that the use of Agent-Based computational techniques allows us to relax hypothesis and overcome the limitations of the analytical approach while produces results that can be easily put in contrast with those of the first approach.

2 The deterministic system

In order to simplify the analysis and avoid the interference that different demand functions could produce in our model, we assume that, at each period t, the sellers can only choose between two possible prices, say a high price $\overline{p}$ and a low one $\underline{p}$. This is, indeed, a very strong simplification. The Agent-Based approach in section 3 could easily overcome this limitation, but then we would loose the reference of the analytical approach. We further assume that at the beginning of each period, the set of strategies for the period available to the sellers consists only of the following:

- *Always cooperate* (C), which means that the seller using this strategy will always set a high price $\overline{p}$ —and thus will "cooperate" with its competitor— regardless what happened in previous rounds of the current period.
- *Always defect* (D), which means that the seller using this strategy will always set a low price $\underline{p}$ —and thus will " defect"— regardless what happened in previous rounds of the current period.
- *Tit for Tat* (T) , which means that the seller using this strategy will start each period with a high price $\overline{p}$ and then, at each successive round in that period, will mimic what its competitor did in the previous round.

Again, the assumption that only three possible strategies are available to the the sellers is a very restrictive hypothesis. Nevertheless, by adopting it we can build on the analysis by Nowak et al. (2004) and apply it to our case. We will later discuss on the essential difficulties of dropping this assumption.

Without loss of generality, we can normalize the number of buyers so that $m = 2$. To begin with we assign to the high and low prices the values $\overline{p} = 3$ and $\underline{p} = 2$. We will first assume that the consumers' behavior is fixed so that they will always buy from the cheapest seller and, in the case that the two of them set the same price, they will split equally. Matrix A summarizes the per period payoff for each seller in this situation. The first row of the matrix specifies the payoff that a C-strategist receives if its opponent is a C-strategist, D-strategist, or T-strategist respectively. Similarly, the second and third row correspond to the payoffs for a D-strategist and a T-strategist respectively.

So, if a D-strategist meets another D-strategist each one receives a payoff of $2R$ after R rounds (one period) whereas if a D-strategist meets a T-strategist, the D-strategist gets 4 in the first round (all the buyers go to the D-strategist) and 2 for the rest of the period.

$$A = \begin{pmatrix} 3R & 0 & 3R \\ 4R & 2R & 4+2(R-1) \\ 3R & 2(R-1) & 3R \end{pmatrix}$$

Each seller, at the beginning of period t, chooses the strategy s ($s \in \{C, D, T\}$) with probability $p_t(s)$. From a evolutionary point of view, $p_t(s)$ represents the proportion of s-strategists in the total population.

Let p_t be the vector of probabilities of choosing each strategy and let p_t^T denote its transposed.

$$p_t = \begin{pmatrix} p_t(C) \\ p_t(D) \\ p_t(T) \end{pmatrix}$$

The following equation drives the dynamics of the learning/evolutionary process:

$$p_{t+1}(s) = p_t(s)\frac{A^s \cdot p_t}{p_t^T \cdot A \cdot p_t}$$

where A^s denotes the row of A corresponding to the strategy s.

Using the learning intuition, the probability of choosing a particular strategy s is updated according to its relative performance. Indeed, $A^s \cdot p_t$ represents the expected payoff of the strategy s whereas $p_t^T \cdot A \cdot p_t$ represents the expected average payoff.

From an evolutionary point of view, the dynamic equation corresponds to the dynamic system known as *replicator dynamics* according to which each strategy is replicated as a function of its relative performance with respect to the rest.

In order to explore the behavior of this system, we study the corresponding vector field, which is shown in Figure 1. Each point represents a probability vector p_t that specifies with what probability each of the three strategies will be chosen. At each point the arrow indicates the instantaneous direction in which these probabilities evolve.

The strategy D is *locally stable* in the sense that if p_t is pushed slightly away from D by a small perturbation, it will eventually go back to D. In other words, D is *evolutionary stable* (Maynard Smith (1982)). That is not the case for T. If p_t is pushed away from T it will then converge to some point near T in the line TC but not necessarily to T again. Any point in the line TC is a rest point (altough not stable) of the system since in the absence of D-strategists, C and T do equally well.

Next we assume a different behavior for the consumers and consider the case in which they are "loyal", that is, they always buy from the cheapest seller as in the previous case. The difference now is that in the case the two prices are set equal, they will go back to the seller they bought from in the previous round. By doing so, the payoff matrix A changes since, obviously, now the first round becomes much more important as it will determine the market share of each seller. Figure 2 shows the vector field corresponding to this case. Now, any path starting in the interior of the simplex converges to D, which

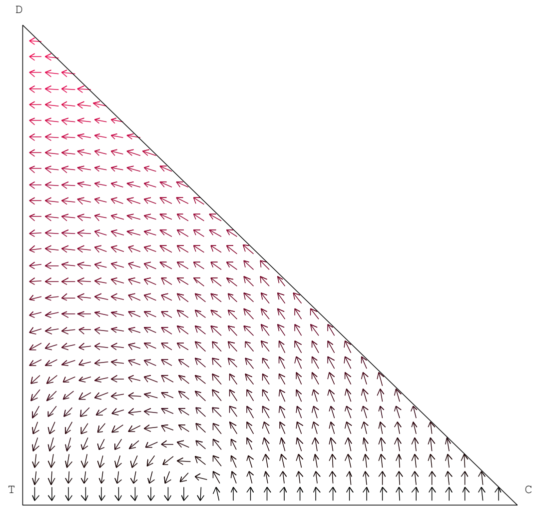

Figure 1: Vector Field with *non-loyal* consumers

turns out to be the only stable equilibrium. The points in the line TC are stationary but not stable.

The change with respect to the first case is not surprising as now the competition for market share in the first round becomes critical. One might argue that in this case we should not consider the *Tit for Tat* strategy (begin with a high price and then imitate your competitor) since it is not a good strategy in the competition for market share. If a seller starts out by setting a high price, chances are that he will not have any one to sell in the next round. If we consider the opposite strategy *Tat for Tit* (Ta) instead, that is, start with a low price and keep it low if the competitor does so, we obtain the results depicted in Figure 3. Basically the result is the same as in the previous case in that the firms end up setting a low price. The difference is that in this case all the points in the line DTa are stationary but none of them (not even D) are *stable*.

Next we consider different values for the high and low prices. Figures 4, 5, and 6 are the equivalent to Figures 1, 2, and 3 respectively with the difference that now the high and low prices are $\bar{p} = 5$ and $\underline{p} = 2$. The size of the gap between the two prices is important since now it might not be optimal for the sellers to retaliate in order to gain market share if it requires to lower the price by more than a half. That would not make sense, for instance, if all that you expect to gain is to double your market share. The direction of the arrows in Figures 4, 5, and 6 seem to corroborate this intuition. In the first one, corresponding to the case in which consumers are not "loyal", it seems that unless there is a high probability of ruining into a D-strategist as a competitor, the mechanism is going to push the probabilities to some point in the line TC.

If consumers are loyal (Figure 5) the result is similar. It seems that we need the probability of meeting a D-strategist to be more than one half in order to converge to the low price equilibrium. In both cases, though, D continues being the only *stable* strategy.

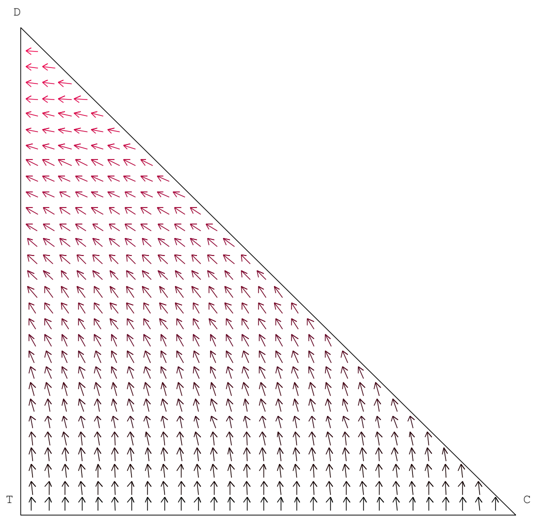

Figure 2: Vector Field with *loyal* consumers

If we consider Tat for Tit instead of Tit for Tat (Figure 6) we obtain very similar results to the ones in Figure 3. This is important since now it seems that being a T-strategist is not so bad after all (according to Figure 5) due to the bigger price differential and, therefore, maybe we should not substitute T by Ta in this case.

In the light of these results one might say that it makes sense for the consumers to be *loyal* (at least in the very weak sense used in this model) for that makes the sellers behave in such a way that competition works in the best interest of the consumers, that is, they get low prices.

Even though this intuition can be thought as a nice complement to the static model of Bertrand competition, we think that the model proposed has two important limitations that makes us consider a different approach. These two limitations are:

1. In the model we only consider three strategies at a time. One might argue in favor of this that C, D, and T are the strategies that seem more likely to emerge from some learning or evolutionary process and, therefore, there is no need to consider more strategies for they will not be played as we approach to a limit solution. This reasoning, though, does not always work. We have an example of that in Figures 2 and 5. According to Figure 2, *Tit for Tat* is not a good strategy, but it turns out to be good when we consider a different set of prices. Thus, a strategy can turn out to perform relatively well or not depending on the particular set up that we have. Since we do not know before hand (if we knew we would not need to solve the problem), we do not have any reason *a priori* to include a particular strategy and discard another.

2. Another limitation is that the interpretation that consumers will be "loyal" because if they do then competition will "evolve" in their best interest involves a strong rationality assumption. Indeed, we are not only assuming that consumers

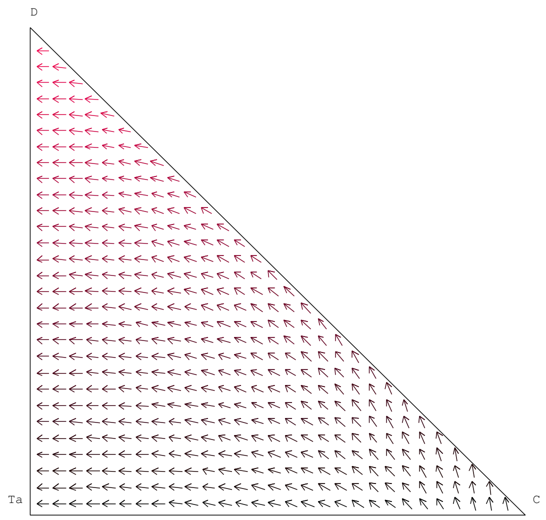

Figure 3: Vector Field with *loyal* consumers (and *Tat for Tit*)

are rational maximizers perfectly informed but also that they are able to understand how evolution works so that they can predict and even control the path of evolution. This is hard to believe. We should, somehow, incorporate consumers' behavior in the system so that all the agents (both sellers and buyers) learn or evolve at the same time.

On the other hand, it is easy to recognize the difficulties involved in relaxing any of these assumptions. Imagine that we want to expand our set of three strategies. The question now is what additional strategies do we want to consider and why, since choosing an additional strategy can be as arbitrary as not choosing it. One might use some criteria to guide this decision. For instance, a good criterion would be: "consider only strategies according to which the action taken by one of the players depends only on his past move and on the past move of his rival". The problem now is that there are 26 different such strategies. It is clear that now we cannot use the same technique since it would imply having to work with vector fields in a 25-dimensional simplex. In fact, working with more than 3 strategies is already a challenge.

Next we propose a model of adaptive learning or evolution based on genetic algorithm techniques to try to overcome these limitations.

3 The adaptive system

In our adaptive system, *strategies* for the repeated game will be represented by *finite automata* as in Miller (1996). A finite automaton is a system that reacts to discrete inputs producing discrete outputs. Formally, a finite automaton –or Moore machine– is described by a five-tuple $< \Omega, Q, q^*, \lambda, \delta >$ where:

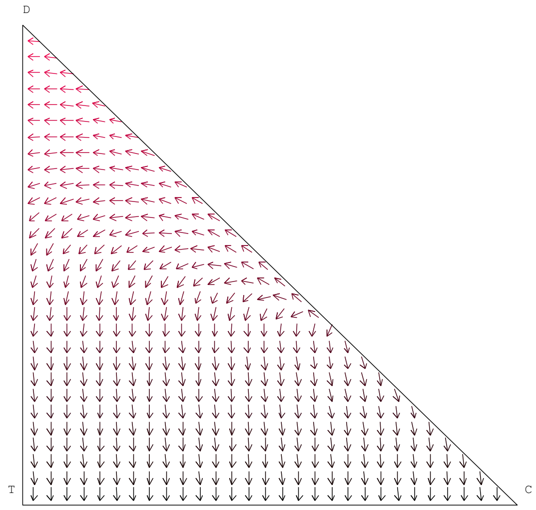

Figure 4: Vector Field with *non-loyal* consumers

1. Ω is the *alphabet* (actions) the automata understands ($\Omega = \{\bar{p}, \underline{p}\}$ in our case)
2. $Q = \{q_1, \ldots, q_n\}$ is a finite set of internal *states*.
3. $q^* \in Q$ is the *initial state*
4. $\lambda : Q \to \Omega$ is the *output function* that assigns an *action* to each internal state
5. $\delta : Q \times \Omega \to Q$ is the *transition function* which maps the current internal state and the input ω that the automaton gets from the environment to a new internal state.

Thus, this automaton consists of a set of internal states. Each internal state has associated an action and a transition function depending on the input that the automaton gets. That transition function points to another internal state that will be the active state for next round. For instance, the automata in Figure 7 represents the Tit for Tat strategy

In our model, the strategies for the sellers will be represented by automata of size 2 (two internal states, labeled 0 and 1). With this we restrict ourselves to consider only 26 different strategies. These automata can be represented by a binary string of 7 bits according to the following scheme:

- *First bit* : Initial state (0 or 1)
- *Second bit* : Action when at state 0. 0 means that the low price $\underline{p}$ is chosen and 1 corresponds to the high price $\overline{p}$.
- *Third bit* : Transition state if the rival sets a low price $\underline{p}$. (0 or 1)
- *Fourth bit* : Transition state if the rival sets a high price $\overline{p}$. (0 or 1)

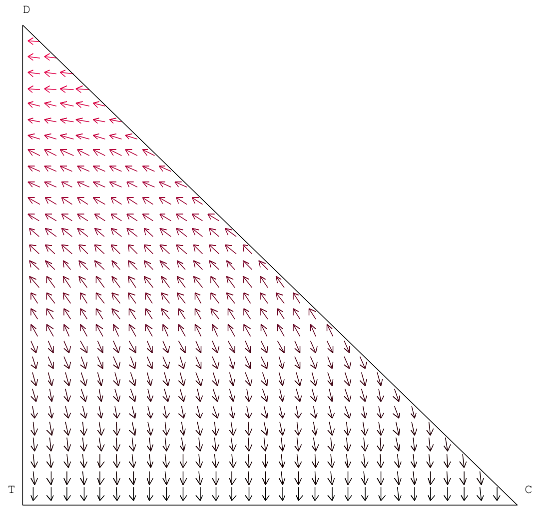

Figure 5: Vector Field with *loyal* consumers

- *Fifth, sixth and seventh bits* : Equivalent to second, third and fourth bits but corresponding to state 1.

Thus, the automaton 0001101 represents a Tit for Tat strategy as in Figure 7.

One of the novelties of this adaptive model with respect to the deterministic model discussed in section 1 is that now the buyers will also be represented by finite automata (of size four) that will evolve jointly with the global evolution of the system. This is a much realistic assumption rather than taking their behavior as given. In this sense, a buyer's strategy is represented by a binary string of 30 bits corresponding to the following structure:

1. The first two bits form a number in base 2 that specifies the initial state (0, 1, 2, or 3)
2. The rest of the bits are structured in 4 groups of 7 bits each, corresponding each of this groups to an internal state. In each of these internal states:
   (a) the first bit corresponds to the action to take at that state. 0 means that the buyer using this strategy will buy from the same seller than in the previous round (loyalty) and 1 corresponds to the buyer switching to the other seller.
   (b) The six remaining bits correspond to the transitions states (expressed in base 2) depending on the price paid by the buyer in the market being higher, equal or lower (respectively) than the price set by the other seller.

The evolutionary nature of the system will be driven by an adaptive mechanism based on a modified genetic algorithm that can be outlined as follows:

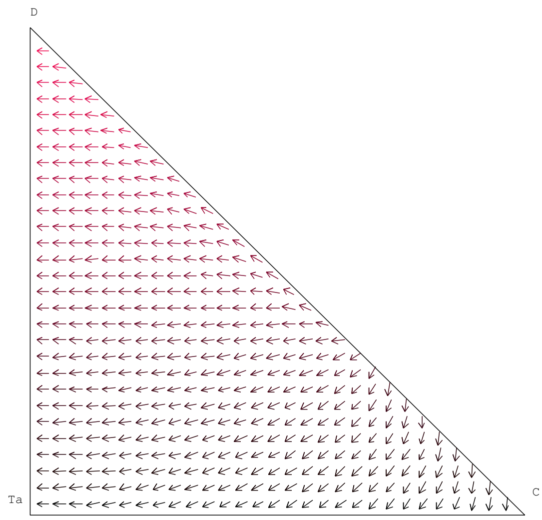

Figure 6: Vector Field with *loyal* consumers (and *Tat for Tit*)

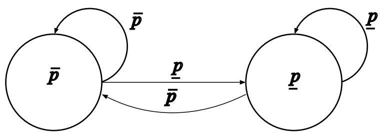

Figure 7: The *Tit for Tat* automaton

1. Two initial populations, one of 32 strategies for the sellers, and another of 1000 strategies for the buyers are randomly formed.
2. Randomly match each strategy of the first population (sellers) to another strategy in the same population.
3. Bring each of these pairs to a simulated market in which each strategy will receive some payoff $s_i(t)$ according to a market mechanism that we will describe later on.
4. Form a new population of strategies for the sellers as follows:
   (a) Include the top 16 strategies from the previous population.
   (b) Create 16 new strategies:
      i. Randomly select two "parents" from the previous population. The probability of strategy j being selected as a "parent" is:

$$prob(j) = \frac{s_j(t)}{\sum_{i=1}^{32} s_i(t)}$$

ii. Form two new strategies applying the *crossover operator*[1] to the two parents.
iii. Apply the *mutation* operator to the newly created strategies.
iv. Repeat a, b, and c 7 times so that 14 more strategies are created.

As we have mentioned above, a market mechanism is simulated to test the performance of both the strategies of the sellers and of the buyers as well. This market works as follows:

1. Two sellers constitute a market that to opens 5 times (5 rounds).
2. At the beginning of each round, 1000 consumers walk in the market and observe the prices announced by each seller.
3. Based on these prices, each consumer decides from which one will buy
4. Trade takes place and each seller and buyer receive their payoffs for period t according to the following scheme:
   (a) For the sellers: Let $d_{i,r}(t)$ denote the number of consumers that buy from seller i at round r and let $p_{i,r}(t)$ the price set but seller i at round r. Then, seller's i payoff is given by
   $$s_i(t) = \sum_{r=1}^{5} p_{i,r}(t) d_{i,r}(t)$$
   (b) For the buyers: Let $b_j(t)$ denote the payoff obtained by strategy j belonging to generation t. Let $k_j(t)$ be the number of times that strategy j pays a low price $\underline{p}$. Then,
   $$b_j(t) = 160(\overline{p} - \underline{p}) k_j(t)$$

In the same way as we do with the strategies for the sellers, the payoffs obtained for the strategies for the buyers are used to create a new population of strategies following a procedure parallel to the depicted for the sellers.

We perform several simulations of the procedure outlined above varying some of the parameters of the model. More specifically, we consider 9 scenarios corresponding to the following (c represents a switching cost for the consumers):

- Case 3A: $\overline{p} = 3 \quad \underline{p} = 2 \quad c = 0$
- Case 3B: $\overline{p} = 3 \quad \underline{p} = 2 \quad c = 1$
- Case 3C: $\overline{p} = 3 \quad \underline{p} = 2 \quad c = 1.5$
- Case 4A: $\overline{p} = 4 \quad \underline{p} = 2 \quad c = 0$

[1]This and the mutation operator will be explained later

- Case 4B: $\overline{p} = 4 \quad \underline{p} = 2 \quad c = 2$
- Case 4C: $\overline{p} = 4 \quad \underline{p} = 2 \quad c = 3$
- Case 5A: $\overline{p} = 5 \quad \underline{p} = 2 \quad c = 0$
- Case 5B: $\overline{p} = 5 \quad \underline{p} = 2 \quad c = 3$
- Case 5C: $\overline{p} = 5 \quad \underline{p} = 2 \quad c = 4.5$

The values for the rest of the parameters were:

- Size of populations = 32 strategies for sellers and 1000 strategies for buyers
- Number of generations (iterations) = 500
- Rounds played by each generation = 5
- Level of relative performance α = 0
- Probability of mutation = 0.001
- Crossover operator = Partial imitation

In Vilà (1997) we discuss the problems generated by the classical *mutation* and *single-cut crossover* operators used in canonical genetic algorithms and propose alternatives to them. First, we use a modified *mutation* operator in which mutations do not take place on a locus-wise basis of the binarily encoded automata but on a state-wise basis of the underlying automata. The reason for this is that locus-wise mutations induce a non uniform distribution over the set of automata that can be the result of the mutation of a given automata. As a consequence of this, when an automaton mutates, some automata are more likely to be the resulting "mutant" than others, which could produce an unwanted erratic behavior of the genetic algorithm dynamics. Second, we modify the standard *single-point crossover* operator to furnish it with a social interaction meaning. In this sense, we introduce the *partial imitation crossover* operator according to which two parent strings (i.e. two parent automata) can only exchange information that makes sense to exchange. This is important due to the special characteristics of the binary representation of finite automata. Indeed, it can be easily showed that the *single-point crossover* operator could produce, from two identical "parents", two automata whose behavior is exactly the opposite from the behavior of their "parents". This is not only awkward but could also induced an strange behavior in the underlying dynamics. More specifically, the *single-point crossover* operator works as follows:

1. Randomly generate a fictitious history of moves by a virtual opponent of a given length. For instance: 0010.
2. Determine the move that each of the two parents would make given the sequence of inputs described by the history generated.

3. Form two new automata that reproduce the two parents but "switching" the move reported in step 2. Hence, the first new automaton will use the move reported by the second automaton if the sequence of inputs it gets is the one described by the history considered and vice versa.

Finally, in order to avoid the problems generated by negative payoffs and, more important, to make the payoff structure immune to affine transformations we work with the payoffs normalized according to the formula:

$$\hat{s}_i(t) = \frac{s_i(t) - \mu}{\sigma} + \alpha$$

where:

$\hat{s}_i(t)$ = Normalized score
$s_i(t)$ = Payoff
μ = sample mean
σ = sample standard deviation
α=parameter that determines the importance of relative performance.

Automata with negative normalized payoff are assigned a zero payoff and, hence, a zero probability of contributing to the formation of the new generation. Therefore, a choice of α=2 implies that the automata whose normalized score is less than two standard deviations from the mean are "killed". The new generation of each population is created in the following way: The first half are the automata of the old population with the highest scores. The other half is created by "reproduction". Two parents are selected randomly being $\frac{\hat{s}_j}{\sum_i \hat{s}_i}$ the probability that automaton j is selected. Then, the *partial imitation* crossover operator is applied.

This reproduction process is repeated as needed to fill in the second half of the new population. Then, the mutation operator is applied to this new population as follows: at each one of the internal states, each transition state is "mutated" to another transition state with small probability ϵ and also each move is mutated to its complementary with this small probability. After the mutation operator has acted we have finally completed our two new populations. Then the whole process starts all over again until the terminal generation is reached.

4 Results

The results obtained by the final generation seem to corroborate those outlined in the section 2 (see the evolution of the average payoffs at the end of the paper). The important difference is that what was an assumption before is now a result, namely that in the case that the two sellers set the same price, the consumers will remain loyal to their previous seller. We have also introduced the possibility for the consumer of having an ex-ante switching cost.

A brief summary of the results obtained follows:

1. *Cases *A* When there are no switching costs ($c = 0$), the buyers will buy from the cheapest seller and will stay with their current ones if prices are equal (loyalty strategy). The sellers will set a low price regardless the price set by the opponent in the previous round (always defect strategy, D). The only difference is when the low price is too low compared to the high price ($\overline{p} = 5$, $\underline{p} = 2$) that the sellers are better off by charging always the high price. (Case 5A).

2. *Cases *B* When the switching cost equals the price difference ($c = \bar{p} - \underline{p}$) the buyers will buy from the cheapest seller and will stay with their current ones if prices are equal (loyalty strategy). The sellers will set a low price regardless the price set by the opponent in the previous round (always defect strategy, D). In this case, the fact that consumers bear a switching cost equal to the difference between the two prices does not make them "more loyal" in the sense of making them willing to stick to a seller (even when he sets a high price) given that there is no possible gain by switching to another. They do not want to do so because then the sellers would have some sort of monopolistic power as the next case indicates.

3. *Cases *C* When the switching cost exceeds the price difference ($c > \bar{p} - \underline{p}$) the buyers will never switch, not even in the case that someone else offers a better price. Then, sellers take advantage of this monopolistic power and always set a high price (always cooperate strategy, C).

As we look at the charts describing the evolution of the average payoffs, we observe that the convergence reached is very robust and is obtained very rapidly. Indeed, in all cases, once the average payoff reaches a steady state it remains there except for very minor fluctuations. This is consistent when the simulations are repeated several times with the same parameter values.

5 Conclusions

We have analyzed the classical Bertrand model of (repeated) imperfect competition using two different but related approaches. The two of them share a common learning (or evolutive) spirit and consider the strategic behavior of the sellers and also of the buyers. One builds on the analysis of Nowak et al. (2004) and uses standard analytical tools to approach the model. It is shown, by using strong simplifying assumptions, that the consumers are better off by developing a "loyal" strategy for that makes the learning procedure of the buyers to result in lower prices. The second approach consists of an adaptive model based on Miller (1996) that uses genetic algorithms to simulate the joint evolution of the behavior of both sellers and buyers. Although his approach allows for a much higher flexibility with regard to the number of parameters to consider and also to the plausibility of the assumptions made, we find that the results obtained coincide almost completely with the predictions of the analytical model.

We believe that putting this two different approaches side-by-side serves to illustrate how the use of this kind of Agent-Based computational models can help in the analysis of problems that otherwise would require strong unrealistic assumptions, or that would be difficult to solve using other techniques. Additionally, the computational model has the appealing of the evolutionary reasoning in the sense there is no need for making strong assumptions about rational players that know and fully understand the game that takes place in the market.

The adaptive model used, nevertheless, also has its limitations. The fact that only two prices are possible is a very unrealistic assumption. Also, the number of buyers is fixed, so that it does not matter what the prices are, the number of buyers never goes up or down. Ideally, we would like to have different types of buyers with different reservation values that would allow for a wider range of prices. Finally, another important feature of market dynamics, namely entry and exit of firms, is missing in our analysis. All these limitations, though, could easily be avoided in more ambitious computational settings and thus approach the analysis of complex imperfect competition models that would be hard to analyze using analytical techniques.

References

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

Ferdinandova, I. (2003). On the Influence of Consumers' Habits on Market Structure. Mimeo. Universitat Autònoma de Barcelona.

Harrington, J. E. and Chang, M.-H. (2005). Co-Evolution of Firms and Consumers and the Implications for Market Dominance. *Journal of Economic Dynamics and Control*, 29(1-2):245–276.

Hehenkamp, B. (2002). Sluggish Consumers: An Evolutionary Solution to the Bertrand Paradox. *Games and Economic Behavior*, 40:44–76.

Maynard Smith, J. (1982). *Evolution and the Theory of Games*. Cambridge University Press, Cambridge.

Miller, J. H. (1996). The coevolution of automata in the repeated Prisoner's Dilemma. *Journal of Economic Behavior and Organization*, 29(1):87–112.

Nowak, M. A., Sasaki, A., Taylor, C., and Fudenberg, D. (2004). Emergence of cooperation and evolutionary stability in finite populations. *Nature*, 428(6983):646 – 650.

Vilà, X. (1997). Adaptive Artifical Agents Play a Finitely Repeated Discrete Principal.Agent Game. In Conte, R., Hegselmann, R., and Terna, P., editors, *Simulating Social Phenomena*, volume 456 of *Lecture Notes in Economics and Mathematical Systems*, pages 437–456. Springer-Verlag.

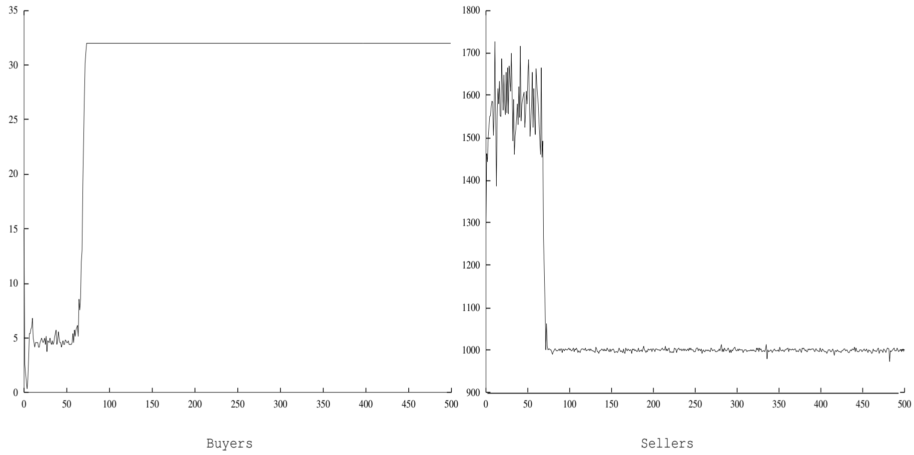

Case 3A

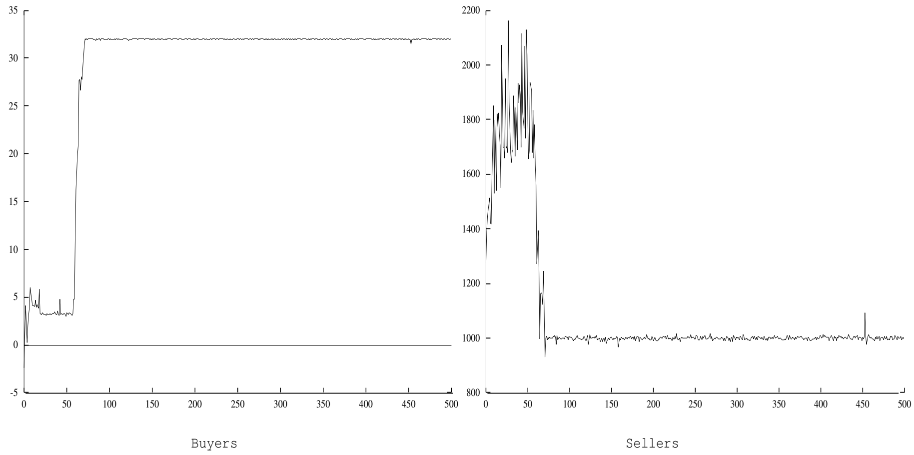

Case 3B

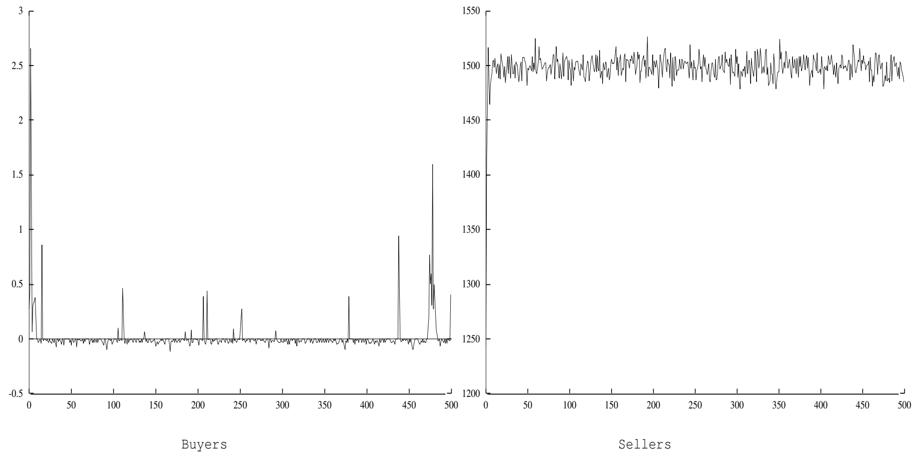

Case 3C

Figure 8: Cases 3*

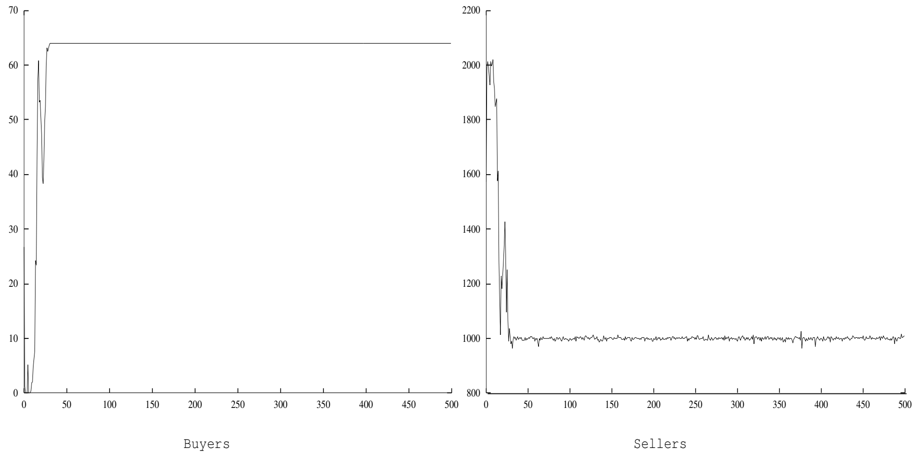

Case 4A

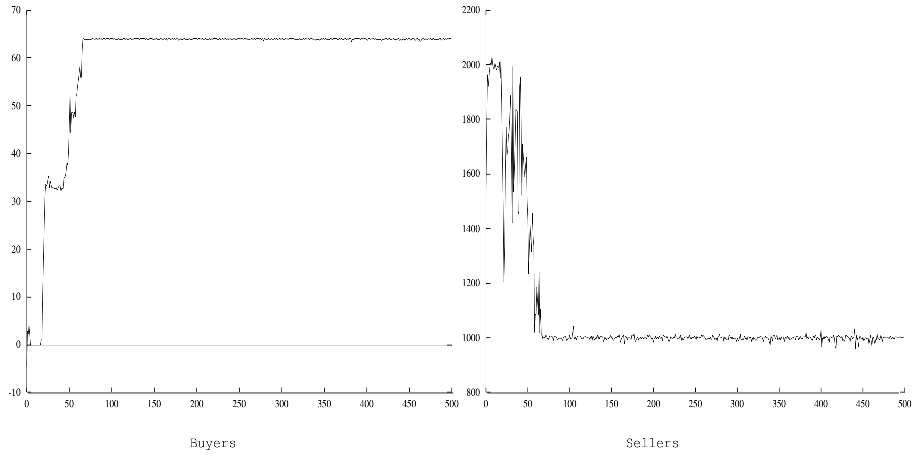

Case 4B

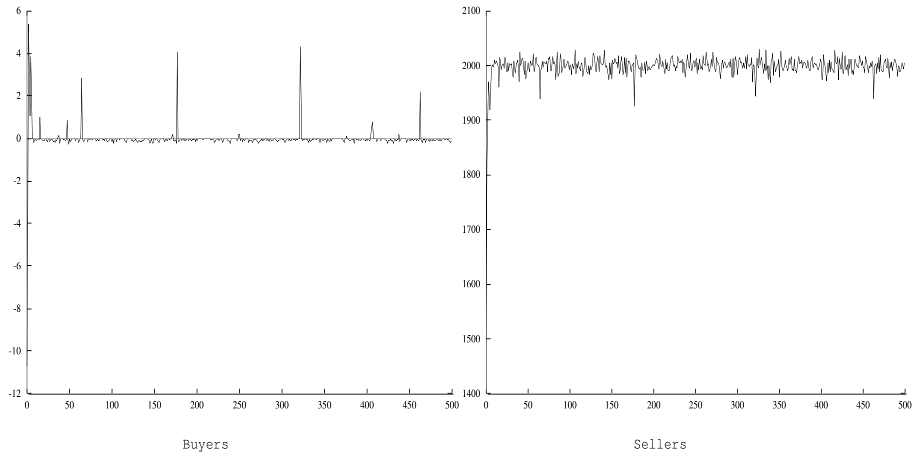

Case 4C

Figure 9: Cases 4*

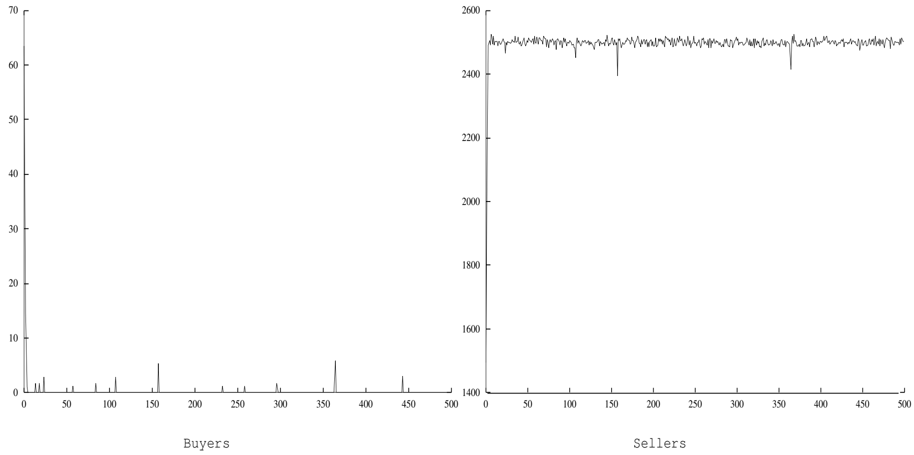

Case 5A

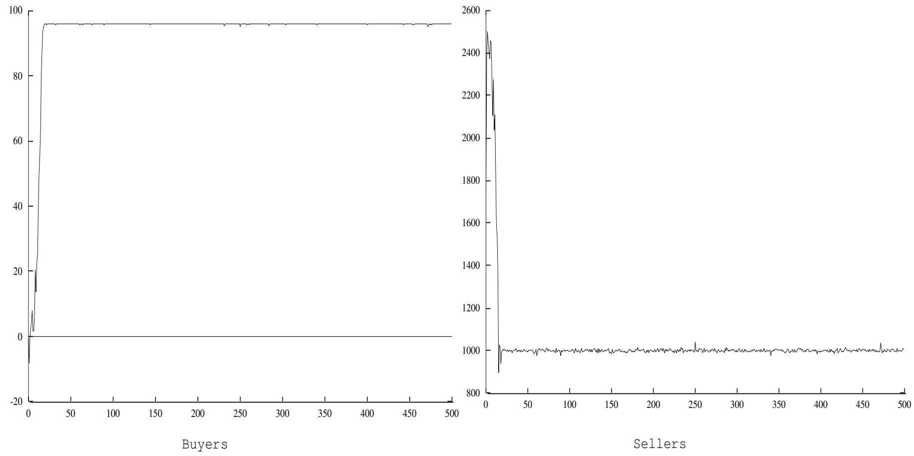

Case 5B

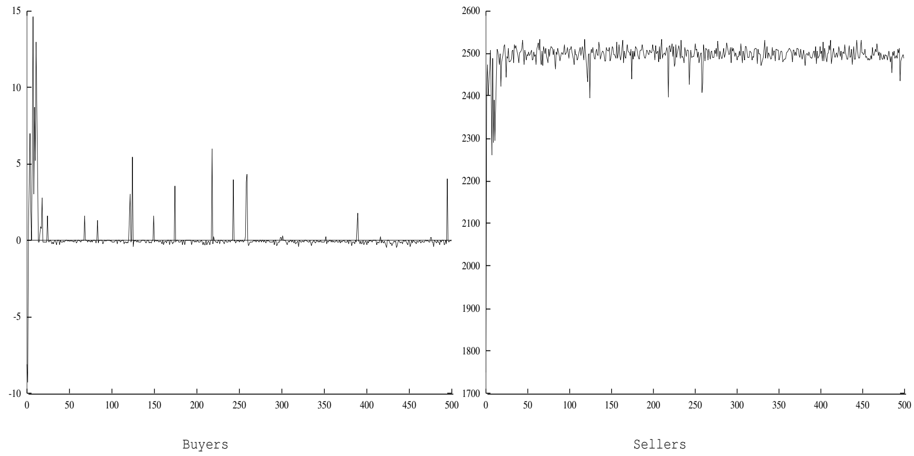

Case 5C

Figure 10: Cases 5*