A MODEL OF SOCIAL LEARNING IN AN ECONOMY WITH MANY DUOPOLIES

Paolo Lupi

Msc Dissertation

Department of Economics and Related Studies

University of York

Introduction

It has long been argued that firms use rules of thumb for solving problems and making decisions (see for example Hall and Hitch (1939), Cyert and March (1963), Simon (1947) inter alia); in other words firms are boundedly rational agents without sufficient computational power to take into account all the information they have. It has also been suggested that the firm's rules of thumb might evolve over time by a trial and error learning process that leads to the survival of the more profitable rules and the perishing of the less profitable ones in a way similar to Darwinian evolution (Alchian (1950), Nelson and Winter (1982), Winter (1971)).

The main notion here is that agents (firms in our case) follow "strategies" or rules that tell them what to do, they do not optimise relatively to fixed preferences, they are not assumed to be so "rational" as to correctly guess the other agent's choices, they simply have decision rules for playing games. But agents can recognise whether the strategy they are using is performing "well" or not (at least compared to the strategy used by the other agents). Over time the strategies that perform better will become more common either through a form of propagation or imitation and will tend to predominate. We can therefore explain the strategies of firms as being the result of such a process of social evolution.

This evolutionary metaphor has received increasing attention from economists (and in general from game theorists and studious of behavioural sciences) in recent years. In the process of social evolution (which dates back to "Social Darwinism") such a mechanism of propagation might also be present, that is, successful firms grow and diversify, efficient firms take-over nonefficient firms, unsuccessful firms go bankrupt.

In the following we shall develop a simulation of an evolutionary process in the general framework of a very simple model of Cournot duopoly which is a typical example of strategic environment. The method of analysis will be as follows. We will consider 50 duopolies each of each is formed by two firms: 1 and 2. Each firm "type" is a particular decision rule, which provides the firm's current output as a function of the other firm's output in the previous

period. We start by randomly generating 100 firm types (one for each of the two firms of each duopoly) and match the firms into 50 duopolies. Then we compute the (deterministic) level of output and the profits for each firm. If the firm's profit falls short of a certain level, determined on the basis of the firms' aspirations and a so-called "bankruptcy level", the firm switches to another existing rule[1] by means of an imitation process. Otherwise the firm will continue to use its rule. We allow this process to continue until all the firms have got a level of profits at least equal to their aspiration level.

The reason for using imitation, to be intended as the act of mimicking actions of others, as an improving rule is that imitation is a behaviour that is often observable in the context of human decision making.

Imitation leads to what is called replicator dynamics. A replicator is an entity that has two properties: 1) it replicates itself and 2) it determines strategic behaviour in a game. Originally the replicator dynamics were only interesting for the biological literature where this concept of dynamics finds its justification in the necessity of reproduction: successful genes become more common because they give rise to more progeny. Lately various models of individual behaviour leading to replicator dynamics have appeared in the economic literature. In fact, rules of thumb, codes of conduct, life styles and scientific ideas are all replicators of a kind. But their mode of replication is not biological, they spread from one man mind to another through imitation or replication (and it is for this reason that Dawking (1976) calls them "memes").

In particular, in our case, firms imitate in order to increase profits in a model of limited information. We should stress that in our case a firm can imitate even a rule that gives it a lower profit, but since the evolutionary process leads to the survival of the best rules only, the probability of choosing a rule that implies higher profits tends to increase with time.

---

[1] By existing rule we mean a rule that is actually followed by one of the other 99 firms.

Furthermore, we allow for the presence of noise in the evolutionary process: it might be (with a small probability) that when a firm switches to another rule instead of imitating another firm, it selects a new, randomly chosen, rule. The economic interpretation of this hypothesis, that we can call "random mutation" hypothesis is that a player might exit with some probability and is replaced with a new player who either knows nothing about the game or fails to perceive the rules of the game and so chooses a strategy (rule) at random. Another economic interpretation of this "random mutation" hypothesis is that players can make mistakes in executing their rules[2].

The main differences between the evolutionary approach that we are using in this work and the evolutionary approach that is usually used in biological game theory are that in our case agents do not switch to the rule that had the best performance in the previous period. It does not seem plausible that every individual knows that there is a best rule and what this rule is. Maybe individuals differ in their perception of their environment. Maybe individuals can not perform such complicated calculations.

Secondly, ours is not a random matching model: firms are randomly matched only in the first period of the simulations (when they are "generated"). After that period firms continue to play against the same competitor unless one of them (or both) goes bankrupt, or is not satisfied with his level of profits or randomly switches to a new rule. Random matching is plausible in biological applications where organisms randomly (and constantly) move in search of other organisms for reproductive necessities. It can also be plausible in a model of an exchange economy, where buyers search out sellers etc. But random matching does not seem plausible in a duopoly context, where firms usually continue to compete in the same market and continue to be located in the same place.

---

[2] And from this point of view this interpretation resembles Selten's (1975) trembling hand story: mistakes are uncorrelated random errors that intervene between the decision to take a certain action and the action itself.

The fundamental question we seek to answer is: what types of behaviour would we expect to find duopolists following in this simple Cournot environment? The traditional reasoning of Cournot himself or of game theorists using the model of "rationality with common knowledge" predicts the textbook Cournot-Nash outcome.

However, when we allow the evolutionary process to take its course, we find a strong tendency towards more collusive outcomes: indeed joint profit maximisation is a very common outcome (although not universal). There are strong evolutionary forces driving behaviour towards co-operation.

The plan of this work is as follows. In section 1 we outline the basic Cournot Duopoly model, in section 2 we outline the method for generating firm types, in section 3 we describe the algorithm used to compute the equilibria in the duopoly markets; in section 4 we describe the evolutionary mechanism; in section 5 we present the results of the simulations; in section 6 we draw some conclusions from the result of the simulations.

1: The Basic Duopoly Market Model

We will be examining a very simple model of Cournot duopoly: there are no costs, and the market price P is a linear function of the two outputs x_i and x_j with the slope and intercept normalised to unity:

(1) $$P = \max\left[0, 1 - x_1 - x_2\right]$$

For obvious reasons, we will truncate each firm's strategy space to the unit interval. The firms' profits are given by the payoff function $U_i : [0,1]^2 \Rightarrow [0,1/2]$ where:

(2) $$U_i(\mathbf{x}) = x_i\left(1 - x_i - x_j\right)$$

For practical purposes, this enables us to concentrate on the set:

$$A: A \equiv \left\{\mathbf{x} \in [0,1]^2 : 1 - x_i - x_j \geq 0\right\}.$$

In this set the standard Cournot-Nash outcome occurs when both firms produce 1/3, and the price is 1/3 leading to a profit, for both firms, equal to 1/9 (=0.1111...) The monopoly point for firm 1 is the point where the firm produces the monopoly output 1/2, and the other firm produces nothing. The line set $1-x_1-x_2=1/2$ is the set of Joint profit maximising outputs. The point where this set intersects with the 45° line represents the symmetric joint profit maximising outcome, at this point each firm earns a profit of 1/8 (=0.125).

Firms play the duopoly game using a decision rule. A decision rule maps the information that the firm has onto its choice of output. We view the decision rule as a primitive, a "black box" which links action to information, and hence runs together the "strategy" of the firm and the process of "reasoning" leading to that strategy[3]. We make three restrictions on the decision rule: first, it is assumed to be linear in information; secondly, the memory of the firm is assumed to be only one period, so that actions in t can only depend on what happens in t-1; thirdly, we restrict firms to decision rules which depend only on the other firm's output. Under these restrictions the model is not only simple (and therefore plausible under complexity constraints), but the decision rules resemble the standard Cournot reaction function, in that current output of firm i depends on the lagged output of firm j.

A Firm type is a decision rule, and under linearity it can be represented as a pair of parameters, the intercept h_0 and the slope h_1: firm type i is thus characterised by a pair $\{h_0,h_1\}$. The firm type (decision rule) is a primitive, which contrasts with much of oligopoly theory which derives the decision rule from some underlying hypotheses (for example profit maximisation, conjectures about the other firms etc.). The space of decision rules includes any decision rule that passes through the set A (hereafter the "unit triangle").

In our analysis, our intention is to talk about decision rules in economic terms with respect to their intercept and slope coefficients. Recall that the market game summarised by (2)

---

[3] When a player switches to a new rule he does not do so because he thinks that the switch is optimal he is just reacting to a stimulus-response mechanism.

is one where firms' outputs are strategic substitutes: the marginal profit of one firm's output is decreasing in the other firm's output. We will describe a firm as being more aggressive if it has a larger intercept coefficient. A large intercept term in effect precommits a firm type to a larger output (for most decision rules, given the slope coefficient, the firm's market share is increasing in its intercept h). The firm's slope coefficient is also important. If a firm type has a negative slope, we will call its behaviour accommodating. If a firm type has a slope of -1, then it is competitive in the sense that it chooses its output to maintain a constant total industry output (a profit maximising firm has a reaction function with slope -1 if it has Bertrand conjectures, as does the Walrasian rule in Table 1). If a firm has a slope coefficient between 0 and +1, then we will call its behaviour co-operative: its rule leads it to move in the same direction as its competitor. In the language of conjectural variations, positive conjectures about your rivals responses elicits co-operative behaviour. Firm types with 0 slopes will be noted as Nash firms. Firms with slopes outside the interval [-1,1] are unstable in the sense that if they play themselves, they give rise to an unstable system (as explained below). Clearly, these terminologies derive from the analysis of the Cournot duopoly model when reaction functions are derived under the hypothesis of profit maximisation, possibly with conjectural variations or some strategic behaviour. In the sense that our decision rules are primitives, these terms should be seen as having only limited significance: they serve merely as suggestive analogies which help us to orientate ourselves in a complicated new environment.

2: The firm generating Algorithm

The way we have chosen to generate firm types is based on the result found in Hey and Martina (1988), Klemperer and Meyer (1988), Bone and Dixon (1992). If we take any point in the interior of the unit triangle, we can generate a reaction function by taking the tangent to firm 1's isoprofit curve at that point. Let us consider some point $\mathbf{x}' \in$ A, at $\mathbf{x}'$ the payoff of firm 1 is $x_1(1-x_1-x_2)$. It is easily verified that the tangent to the isoprofit curve at $\mathbf{x}'$ is the ratio of

the marginal profits $\partial U_1/\partial x_i$ at **x**'. Hence the tangent at **x**' is characterised by the slope and intercept terms:

(3a) $h_0 = 2x_1 + 2x_2 - 1$ intercept

(3b) $h_1 = (1 - 2x_1 - x_2)/x_1$ slope

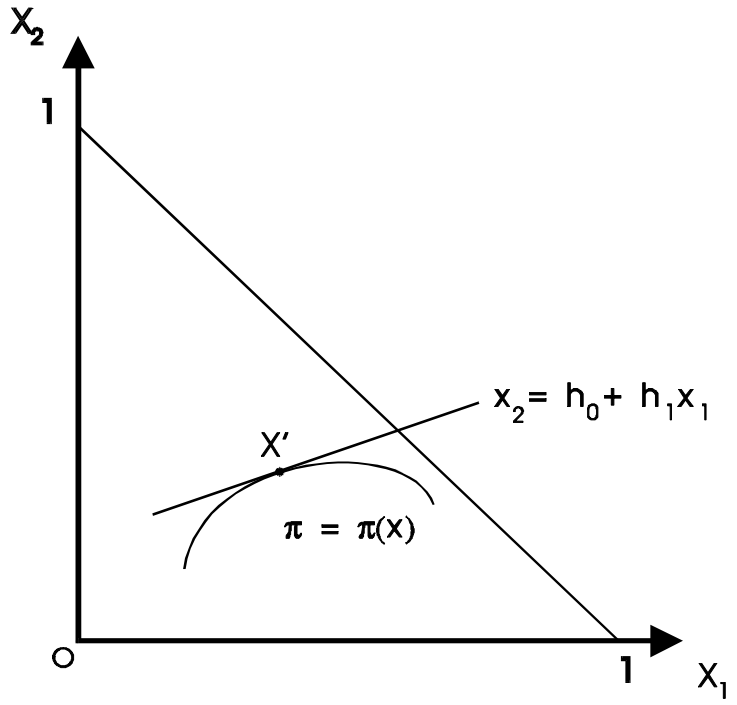

Figure 1: The algorithm generating firm types.

Note that if firm two chose the decision rule thus generated, then the best profit that firm 1 could achieve is by choosing a decision rule passing through **x**'. That is, **x**' solves the problem:

$$\max U(\mathbf{x}) \text{ subject to } x_2 = h_0 + h_1 x_1$$

where the slope and intercept are as in (3). A Nash equilibrium in decision rules occurs at **x**' when firm 1 chooses as its reaction function the tangent to firm 2's isoprofit curve at **x**', and vice-versa. Such an equilibrium can be constructed to support any point in the interior of the unit triangle (see Bone and Dixon (1992) for a further analysis).

This method of generating firm types has two advantages. First it is defined over a compact convex set, second it provides a simple visual and graphical way to represent a firm: we can reverse the algorithm and represent the two-dimensional parameterization of firm 1 $\{h_{0i}, h_{1i}\}$ by the point in the unit triangle which generated it (note that the mapping represented by (3a,b) is 1-1).

3: The Computation of the equilibria in the markets

Since we are dealing with 50 duopolies. We generated 100 randomly chosen decision rules than are matched in 50 duopolies. To be more precise we generate 100 random points in the interior of the unit triangle (points whose co-ordinates x_1 and x_2 are such that $0 < x_1 + x_2 < 1$); from these random points we, using the HMBD result, derived the actual firm rules. Then we matched them in 50 pairs to play a duopoly game. We need to be more explicit about how a constituent duopoly game is played. Each firm has a decision rule, which gives its output in period t as a function of the other firm's output in the previous period. This is a dynamic system of the form (for firms type 1 and 2):

$$\mathbf{x}_t = \mathbf{h} + \mathbf{H}\mathbf{x}_{t-1} \tag{4}$$

where $\mathbf{h} = [h_{01}, h_{02}]$, and $\mathbf{H} = \begin{bmatrix} 0 & h_{11} \\ h_{12} & 0 \end{bmatrix}$

In order to save computational time and to take into account the non linearities of the system (only equilibria inside the unit triangle have an economic meaning) we have used a very simple even not very efficient algorithm to compute the firms' profits in each of the 50 constituent games.

Given the solution x_i and x_j of the linear system we have that:

$$(5)\qquad \begin{cases} \text{if } x_i \geq 0 \text{ and } x_j \geq 0 \text{ then } x_i^* = x_i \text{ and } x_j^* = x_j \\ \text{if } x_i < 0 \text{ and } x_j < 0 \text{ then } x_i^* = 0 \text{ and } x_j^* = 0 \\ \text{if } x_i \geq 0 \text{ and } x_j < 0 \text{ then } x_i^* = h_{0i} \text{ and } x_j^* = 0 \end{cases}$$

On the basis of the levels of output for each firm in each of the n=1..50 markets we computed the duopoly equilibrium prices p_n n=1..50.

4: The evolutionary algorithm

After having determined the level of output and the price in each of the 50 duopolies we computed the level of average profits over all the 100 firms at time 1:

$$(4)\qquad \overline{\pi}_1 = \frac{\sum_{n=1}^{50}(x_1 + x_2)p_n}{50}$$

At this point the evolutionary process starts to take its course, in fact we hypothesised that firms have an "aspiration level of profits" Λ determined from the average level of profits on the basis of an "aspiration parameter[4]" λ (that for the sake of simplicity we consider constant over all firms):

$$(5)\qquad \Lambda_t = \lambda \cdot \overline{\pi}_t$$

Firms also have a "bankruptcy" level of profits γ, to be intended as a minimal level of profitability, that we arbitrarily fixed at $\gamma = 0.05$[5]. Firms whose profits at time 1 fall short of the maximum value between the aspiration level and the bankruptcy level switch to another rule, in other words:

if $\pi_{it} < \max(\Lambda, \gamma)$ then firm i (i=1..100) switches to another rule.

---

[4] We have ran simulations for different values of the aspiration parameter.

[5] From a pure computational point of view, the presence of a "Bankruptcy" level speeds up the process of learning by eliminating more fastly the less effective rules.

The rule that a firm who decides to switch will adopt at period t+1 is randomly chosen with equal probability between all the rules used by the other 99 firms.

It is worth noting, at this point, that during the evolutive process a firm only knows its own payoff but not the rules of its competitors, the only additional informations that he has is the average level of profits.

But this process of imitation does not happen with certainty; in fact, as we have outlined in the introduction there is also a small probability δ that a "random mutation" happens, that is to say that a firm instead of imitating one of the other 99 firms (and therefore adopting its rule) switches to a new, randomly chosen[6], rule. The effect on the evolutionary process of this random mutation noise is to widen the set of rules existing at each time. It is worth noting that the probability of a random mutation decreases with time and with the diffusion of the best rules since only firms who are not satisfied with their level of profits can mutate.

After that all the firms whose profits at time t were lower than $\min(\Lambda, \gamma)$ have switched, either by imitation or by random mutation, to a new rule, the equilibrium levels of outputs, price and profits are newly computed in each of the duopolies in which at least one firm has changed its rule. The evolutionary process outlined in this paragraph continues until all the firms have reached a level of profits at least equal to their aspiration level.

5: The results of the simulations.

We have developed a program in Maple V[7] to simulate the evolutionary process explained in the previous sections. At the end of each simulation, that is to say when all the firms have acquired a profit at least equal to their aspiration level, the simulation program computes (among the other things) the level of average profits (over the 100 surviving firms), the profits of the firm who had the worst performance and the profits of the firm who had the

---

[6] Using the HMBD result.

[7] The actual version of the program is reported in appendix I.

best performance. With this program we have tested the behaviour of the profits with 3 different aspiration levels: 0.8, 0.9 and 0.99. For each of the aspiration levels we have run 100 simulations and we have computed some summary statistics. The following table summarises our results[8]:

Table 1: Summary statistics of the simulations

Aspiration level	average profits	st. dev. of profits	worse profits	best profits	avg num. of periods
0.8	0.108736	0.012436	0.065409	0.124636	36.34
0.9	0.110541	0.011695	0.077621	0.124917	40.87
0.99	0.119203	0.006054	0.095105	0.124992	68.91

The second columns of table 1 shows the average level (over all the 100 simulations) of the average profits of the 100 firms. As we can see there is not much difference in the average level of profits deriving from an aspiration level of 0.8 and 0.9; both levels are very near to the Cournot level. On the other hand, the average profits deriving from an aspiration level of 0.99 show a much higher degree of co-operation between firms.

As expected the standard deviation of the average profits (reported in the third column of table 1) tends to reduce with the increase in the aspiration level (and in this case, too, there is no much difference between the 0.8 and the 0.9 aspiration level case); with an aspiration level of 0.99 the set of rules that survives the evolutionary process is usually populated only by two or three elements (and quite often only by one element). These results can be observed more clearly in following histograms showing the distribution of profits over the 100 simulations that were run.

[8] The detailed results of the simulations are reported in appendix II.

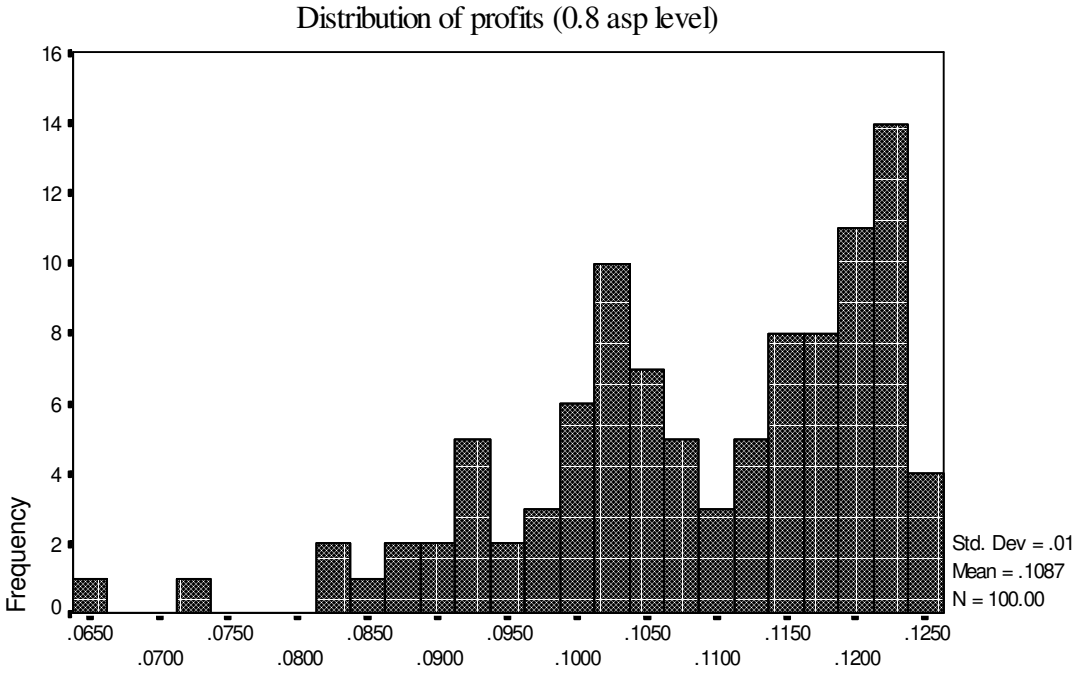

Fig.2: Distribution of the average profits over 100 simulations (aspiration level: 0.8).

The distribution of the average profits over 100 simulations in the case of an aspiration level of 0.8 shows an average value lower than the Cournot's profit level (even if in 51 simulations the average level of profits was higher than the Cournot and in 22 simulations it was higher than 0.12).The standard deviation of the average profits is relatively high; the reason of this variability depends on the high sensitivity of the evolutionary process to the random initial positions caused by the low aspiration level. In fact, when the aspiration level is particularly low the evolutionary process takes a limited number of periods to complete (in fact the average duration of the evolutionary process over the 100 simulation was 36.34 periods); this, as a consequence, leaves no much space for random mutation and, if the initial set of rules is composed by rules that do not perform very well, there is an high probability that these rules will continue to survive until the end of the process.

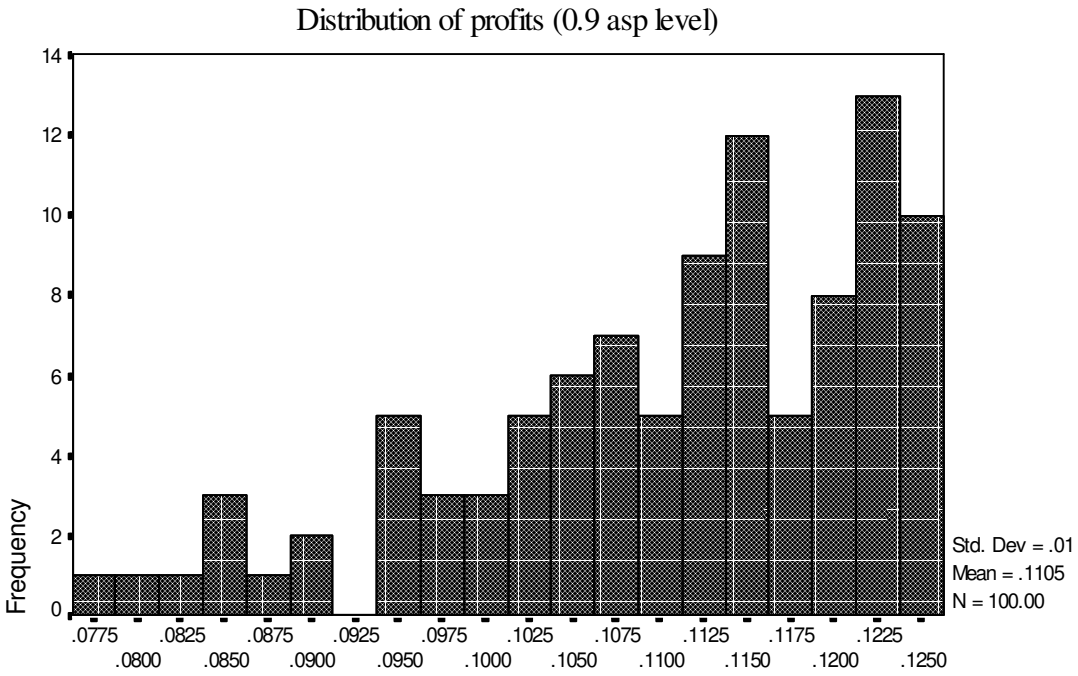

Fig.3: distribution of average profits over 100 simulations (aspiration level: 0.9)

With an aspiration level of 0.9 the results do not differ too much from the 0.8 case. The average over the 100 simulations of the average level of profits is slightly higher than in the previous case and is very similar to the Cournot level. In 59 simulations the profits were higher than the Cournot profits and in 27 simulations they were higher than 0.12. The variability of the average profits was practically identical to the variability of the 0.8 case, so the number of periods needed to complete the evolutionary process, and the same conclusions can be applied in this case.

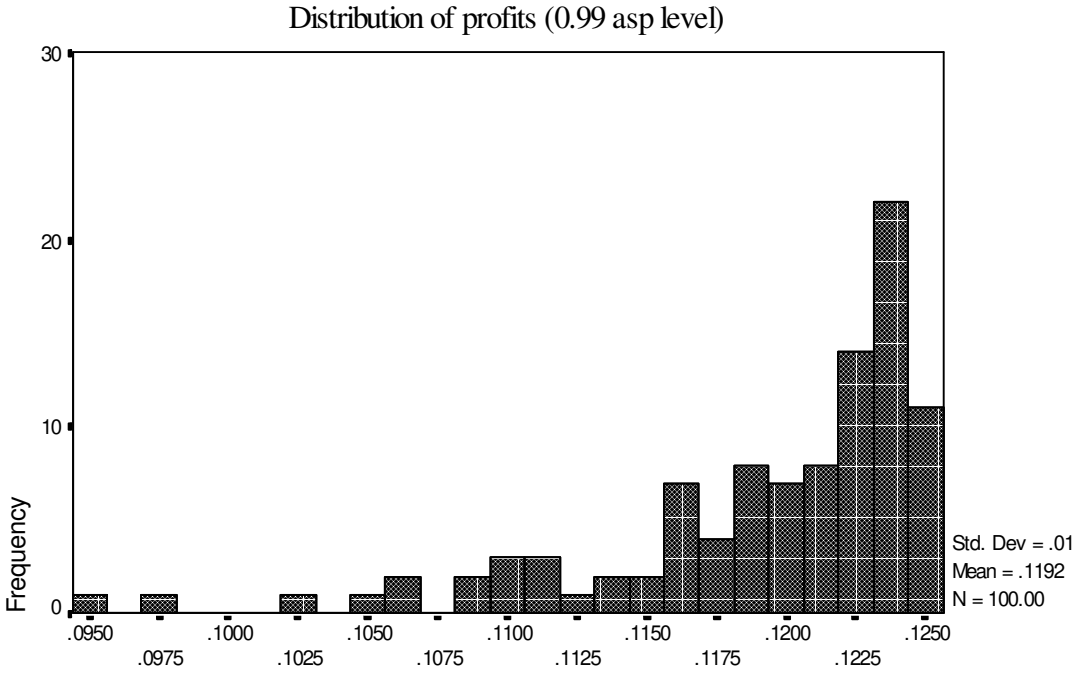

fig 4: distribution of profits over 100 simulations (aspiration level: 0.99)

With an aspiration level of 0.99 we have very clear results. The average value of the average profits is 0.192; the collaborative nature of this results emerges clearly from the distribution of the average profits shown in fig. 4: most of the mass of the distribution is located at profits levels higher than 0.12 (in 58 simulations the average profits were higher than 0.12 and in 17 they were higher than 0.124). in this case the number of periods necessary to complete the evolutive process is sensibly higher than in the previous case (the average value was 68.91)

Further information can be drawn from the observation of the evolution of profits over time in each single simulation.. The following pictures show this process. We have chosen, for each aspiration level, the simulation whose final average profits were more similar to the average value over all the 100 simulations of the average profits[9].

[9] So, for example, in the 0.8 aspiration level case, we have chosen the simulation n° 45 because among all the simulations that we have run with that aspiration level it was the one whose final average level of profits were nearer to the average value of the average profits.

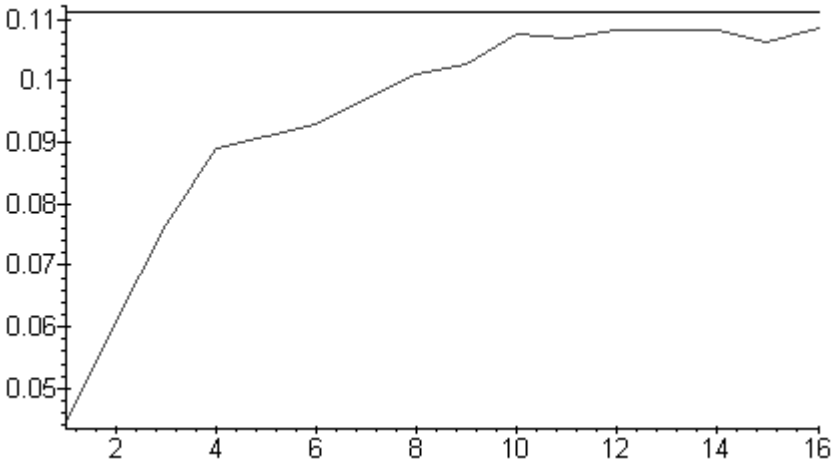

Fig. 5 - Graph of the evolution of average profits (simulation n° 45, aspiration level: 0.8)

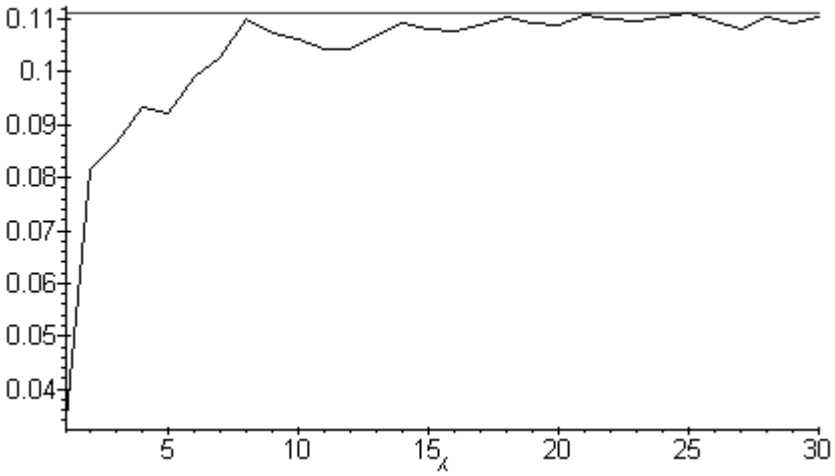

Fig. 6 - Graph of the evolution of average profits (simulation n° 11, aspiration level: 0.9)

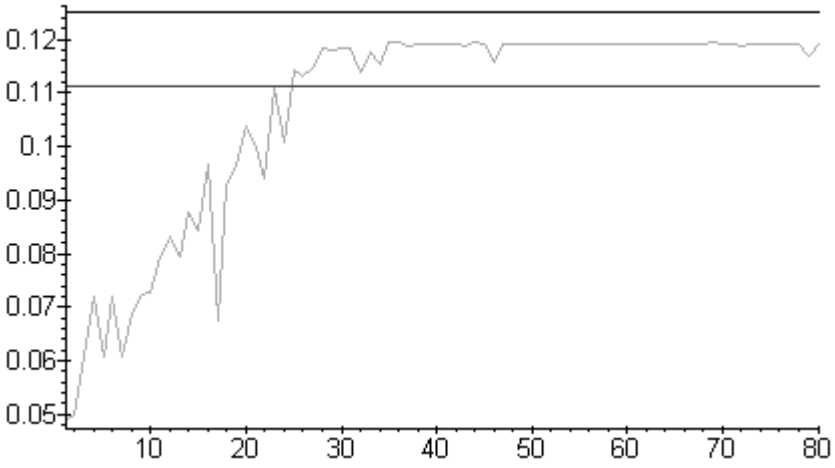

Fig. 7 - Graph of the evolution of average profits (simulation n° 40, aspiration level: 0.99)

As we can see, the learning process is initially very fast (in all cases after less than 5 periods the average profits were higher than the bankruptcy level). Then after a number of periods approximately equal to one third of the total number of periods the increase in the average profits become tendentially more flat. Quite often we can see a steep decline in the average profits curve. This happens when a rise in the average profits causes many firms to switch rule, in this case very often happens that the firms that switches rule either imitate a rule who is not performing very well[10] or switch to a new randomly chosen whose performance is worse than its previous rule.

In conclusion we can say that when an evolutionary approach is applied to the most commonly used oligopoly model, there is a strong tendency toward collusion. This result emerges very clearly when the aspiration level is particularly high and approaches to unity. Clearly this approach is more of a metaphor than a model (altough no less abstract than many economic models).

[10] remember that a firm can imitate even a rule that gives her a lower profit.