The propagation of cooperation in a model of learning with endogenous aspirations.

Paolo Lupi*
Department of Economics and Related Studies
The University of York, York YO1 5DD, UK.
E-mail: paolo@shiva.york.ac.uk

Abstract

In this paper we build a spatial, aspiration-based model of learning in the context of Cournot oligopoly from which we want to explore the conditions that lead to the emergence of cooperation among firms. We consider an economy consisting of many identical duopolies; each duopoly is placed on a square of a torus. The duopolists are boundedly rational agents which adopt a very simple behavioural rule: if they are earning at least average profits, they do not change their strategies; if they are earning below-average profits they imitate the strategy adopted by one of their neighbours. We consider many variations to this simple setup and, in almost all cases, as in Dixon (1998) and Dixon and Lupi (1997) we get results that support cooperation among firms.

Keywords: Oligopoly, local interaction, learning, aspirations, evolution.

*The development of this paper started during the III Graduate Workshop in Computational Economics (14–29 July 1997) held at the Santa Fe Institute, whose hospitality and financial support is gratefully acknowledged. The author would like to thank John Miller, Scott Page, Agostino Manduchi, Francisco Rodriguez and all workshop participants for helpful comments and suggestions. This work is part of the ESRC project "Evolution, Oligopoly and Competitiveness".

1 Introduction

The issue of the emergence of cooperation among individuals has been widely discussed by a large number of game theorists, economists, biologists and in the behavioural sciences in general. The reason for all this interest, especially by part of economists and game theorists, probably lies in the fact that cooperative behaviour is something that is frequently observable in many economic and social contexts and that, as many argue, traditional game theory fails to explain.

In the past two decades many theories have been devised to explain how and why cooperation, at least to a certain degree, may emerge as the final result of an evolutionary process. Several of these theories are grounded on the seminal work of Axelrod and Hamilton (1981) and Axelrod (1984) who showed, by means of computer simulations, that cooperative behaviour is the possible outcome of repeated interaction among agents that change their action following some adaptive behavioural rule. They identified the reason underlying the emergence of cooperation in the fear of future retaliation by part of other individuals in case of individualistic behaviour. In Axelrod (1984), in addition, it is shown that the probability of the evolution of the system towards a cooperative equilibrium is greatly increased if, in a model of imitative behaviour, a form of localization of the players and local interaction is assumed and modeled. If, in fact, agents are placed on a spatial structure or on a network of social interactions and the probability of being matched with another agent depends on some criterion of proximity that puts higher weight on the interaction between nearer agents, then clusters of agents that adopt cooperative strategies[1] may form and propagate over the space of interaction.

More recently a few papers have appeared in the literature on evolutionary games and learning that explicitly study the implications of local interaction in games when some form of limitation on players rationality is assumed. Blume (1993), for instance, considers the case of 2×2 symmetric coordination games played by agents who live on a lattice and interact only with a limited number of neighbouring agents; he finds that, under myopic best response dynamics, possibly perturbed by random choices, the risk dominant equilibrium[2] is selected in the limit; the same findings, unfortunately, not always apply to $n \times n$ games. In a similar fashion Ellison (1993) study a local interaction version of the (Kandori et al., 1993) model of stochastic best-reply and finds that with local interaction, convergence toward the final

[1]That are usually strategies that perform well when played against themselves, but not so well when played against others.

[2]Which, as known, is the equilibrium with largest basin of attraction in the set of possible mixed strategies.

pareto-efficient equilibrium is achieved more quickly than with global interaction, and that the speed of convergence is inversely related to the size of the neighbourhood.

Anderlini and Ianni (1996) study a similar setting characterized by a finite population of agents that are at every time t matched with one of their neighbours to play a pure coordination game. Agents update their strategies following a majority rule whose outcome depends on the proportion of agents playing each of the two available strategies. The agents' choice of action is also affected by some form of inertia. They find that the dynamical system converges with probability one to a steady state, but since the system shows path dependence, many limiting states are possible including a state in which only local coordination occurs.

Our model differs from the literature above in a few ways. The main difference lies in the fact that our unit of investigation is not the single agent, but the pair of agents. In fact, we consider a population of duopolies placed on a two-dimensional connected square lattice (a torus). Matching, as in Dixon (1998) and Dixon and Lupi (1997), is therefore "permanent" in the sense that agents, firms in our case, face the same competitor over time, and are not allowed to move and/or change competitor.

Every cell on the lattice can be thought as a location or a market where two competitors face each other in duopolistic competition. The more evident interpretation of a location is of a position in a geographical space, but it may also be considered as a position in the "space of products" where neighbouring cells are occupied by firms producing a slightly differentiated product.

Firms at every duopoly adopt a very simple imitative rule based on aspirations in their decision making: at every time t they compare their actual profits with an aspiration level which is endogenously determined on the basis of the average profit in their neighbourhood. If their profits are at least equal to this aspiration level they stick to the strategy used in the previous round of action, on the other hand, if their profits fall short of the aspiration level they revise their strategy by imitating the strategy of another, randomly chosen, firm of their neighbourhood.

We consider two variations to this simple learning rule. According to the first, we let the probability of revising the strategy for unsatisfied firms vary with the difference between their current profits and the average level of profits in the neighbourhood. In other words firms for which the difference between current profits and the average level of profits in the neighbourhood is substantial will be more likely to change their strategy with respect to firms that experience a lower difference. This rule may reflect the presence of inertia in the strategy revision procedure or some form of lock-in motivated by technological reasons or by the presence of switching costs.

We have also studied a perturbed version of our model. According to this "noisy" version unsatisfied firms may, with probability ϵ make a mistake in their imitation procedure and end up choosing a new strategy at random. Noise, as now customary in evolutionary models, is also introduced to test the robustness of the model.

The fact the firms can only imitate the strategies of a finite subset of the population of firms (their neighbourhood) reflects the presence of costs of gathering information about the strategies and the performance of firms located in distant locations. This "limitation" could also be explicitly desired by firms. In fact, especially in the case where we interpret the cells of the lattice as locations or markets in a geographical space, firms may be not interested in imitating the strategy of firms far-off on the lattice, because they could think that those strategies might not be very effective in their area due to spatial differences in the "fundamentals" of the economy.

Despite the localization of interaction, the behaviour of a firm, not only directly affects the outcome of its direct competitor in the duopoly, but can also indirectly propagate to more distant agents by successively influencing the choices of locally interconnected neighbours via both the local average mechanism and the imitation procedure.

Even though the model is quite simple, the nature of interaction among firms is very complex. There are, as we have seen, different levels of interactions among them: at a very local level (the location or duopoly level) firms interact with their competitors in the Cournot duopoly game, at an intermediate level they interact with the firms in their neighbourhood via the imitation mechanism and, in the case in which the average level of profits is computed over the whole "economic system" there is also a global level of interaction in which every firm indirectly interacts with all the others by way of the global level of average profits.

The complexity of the interactions makes extremely difficult, if not impossible, to analyze analytically our model (a problem common to many models of spatial interaction). For this reason the analysis of the dynamical system will rely mainly on computer simulations.

We will give full account of the results of our simulations in section 3, but we can anticipate here that in almost all the simulations we got convergence to the cooperative equilibrium characterized by all firms playing the joint profit maximizing strategy. This result is also robust to small levels of "noise".

The reminder of the paper is structured in the following way: in section 2 we introduce the formal model, in section 3 we report and discuss the results of the simulations, section 4 contains some concluding remarks.

2 The Model

2.1 The spatial structure

Consider an "economy"' of duopolies placed on a two-dimensional connected square lattice. Every cell on the lattice is the address of a pair of firms permanently tied to play a Cournot game.

We identify every duopoly by $d(x, y)$, where x and y are the horizontal and vertical coordinates of the cell on the lattice occupied by the duopoly. If the number of rows and columns is equal to K (a square lattice) we have that the duopoly space is $\mathcal{D} = \{1, \ldots, K\} \times \{1, \ldots, K\}$ and its cardinality $|\mathcal{D}|$ is, of course, equal to K^2.

We measure the distance $dist\,[d(x, y); d(u, v)]$ between any two cells or duopolies $d(x, y)$ and $d(u, v)$ using a modified version of the metric induced by the $\|x\|_\infty$ norm, that takes into account the fact that in our case distances are measured over a connected lattice (a torus). According to this metric we have that:

$$\begin{aligned}\text{dist}\,[d(x, y); d(u, v)] = \max\{ &\min\,[\mid x - u \mid, K- \mid x - u \mid]\,, \\ &\min\,[\mid y - v \mid, K- \mid y - v \mid]\}.\end{aligned}$$

In other words the distance between any two duopolies is given by the minimal number of duopolies that must be crossed (in all directions, including diagonals and allowing for border-crossing) in order to go from the first to the second. We define the neighbourhood of radius ρ of duopoly $d(x, y)$, $N^\rho_{x,y}$ as the set:

$$N^\rho_{x,y} = \{d(u, v) \in \mathcal{D} \mid \text{dist}\,[d(x, y); d(u, v)] \leq \rho\}\,,$$

that is the set containing all cells or duopolies located within ρ distance from $d(x, y)$[3]. This formalization of the neighbourhood relation is symmetric: for all $d(x, y), d(v, u) \in \mathcal{D}$, the fact that $d(u, v) \in N^\rho_{x,y}$ implies that $d(x, y) \in N^\rho_{u,v}$. We denote the number of duopolies in the neighbourhood of size ρ of $d(x, y)$ by $|N^\rho_{x,y}| = (2\rho + 1)^2$ and, since at each duopoly there are two firms we have that the total number of firms in $N^\rho_{x,y}$ is $2|N^\rho_{x,y}|$. Figure 1 shows two neighbourhoods characterized by two different values of the parameter ρ; the locations depicted as grey circles constitute the neighbourhood of the location depicted as a black circle. Every circle represents a duopoly.

[3]The neighbourhood relation we have defined, is also known, in the theory of cellular automata, as Moore neighbourhood.

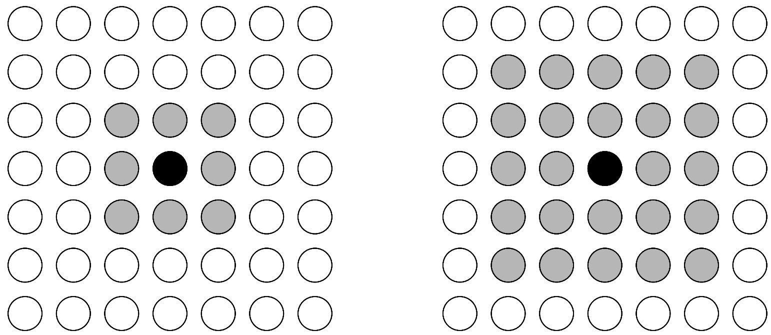

A neighbourhood of size $\rho = 1$. A neighbourhood of size $\rho = 2$.

Figure 1: The neighbourhood structure

We have, due to computational constraints, restricted the size of the square lattices to values of $K \leq 10$ (corresponding to a population of duopolies $N \leq 100$), however the results of the simulations have shown that the value of K, for K sufficiently large i.e. $K \geq 6$, does not affect significantly the final outcome, at least in terms of the limiting distributions of the evolutionary processes.

Even though geographical structures such as torii (chequerboards where the edges are pasted together) do not exist in real world, we have decided to adopt these structures on the grounds of their capability to avoid "boundary effects", i.e. the ability to maintain the number of duopolies in the neighbourhood of each duopoly constant independently of the position of the duopoly on the lattice.

2.2 The Cournot game

The set of all firms f in the economy is denoted by $\mathcal{F}$, since at each location or duopoly $d(x, y) \in \mathcal{D}$ there are two firms, the number of the elements of the set $\mathcal{F}$ is $|\mathcal{F}| = 2|\mathcal{D}|$.

We define the index function $g(f, x, y)$ as the function that returns 1 if and only if firm f operates at duopoly $d(x, y)$ and 0 otherwise:

$$g(f, x, y) = \begin{cases} 1 & \text{iff } f \text{ is at } d(x, y) \\ 0 & \text{otherwise.} \end{cases}$$

The set $f^N \subseteq \mathcal{F}$ is the set of all firms f in the neighbourhood $N^{\rho}_{x,y}$ of duopoly $d(x, y)$ and it is defined as:

Table 1: Cournot Duopoly: Reference Points

outcome	output	profit per firm
Cournot-Nash Equilibrium	$q^f = q^{f'} = \frac{1}{3}$	0.1111
Joint-Profit Maximum	$q^f = q^{f'} = \frac{1}{4}$	0.1250
Stackelberg	$q^f = \frac{1}{2}, q^{f'} = \frac{1}{4}$	0.1250, 0.0625
Walrasian	$1 - q^f - q^{f'} = 0$	0

$$f^N = \left\{ f \in \mathcal{F} | g(f, x, y) = 1 \text{ and } d(x, y) \in N^{\rho}_{x,y} \right\}.$$

We say that firm f' is a competitor of firm f at duopoly $d(x, y)$ if and only if $g(f, x, y,) = g(f', x, y) = 1$.

At each location or duopoly $d(x, y)$ there are two firms f and f' playing the simplest possible Cournot game that is assumed to be the same everywhere on the lattice. The market price $p_{x,y}$ at $d(x, y)$ is a linear function with slope and intercept normalized to 1 of the outputs q produced by the two firms f and f'. That is to say, dropping the location subscripts for convenience:

$$p = \max \left[0, 1 - q^f - q^{f'} \right],$$

where both q^f and $q^{f'}$ are non-negative. The profits of each firm f are therefore $\pi^f = q^f \ p$. We have not explicitly included costs in the demand function, but we can interpret the price as net of constant average production costs.

In this Cournot-Nash model, the outputs of the two firms are strategic substitutes (the best-response functions are downward sloping), and the payoff of each firm is decreasing (when positive) in the output of the other firm. In order to generate a finite strategy set for this model, we constructed a grid of outputs, so that each strategy is an output level. The (symmetric) $S \times S$ payoff matrix Π gives the payoffs $\pi_{i,j}$ to the firm producing output levels $i = 1 \ldots S$ when the other firm is producing output levels $j = 1 \ldots S$.

According to this specification of the demand function the Cournot-Nash outcome occurs when both firms produce 1/3 with a price of 1/3 and a corresponding profit per firm of 1/9. The joint profit maximizing outcome (JPM) occurs when both firms produce 1/4 and earn 1/8 profits. Table 1 shows some of the key reference points of the model.

In the simulations the set of outputs (or strategies) is generated by a grid search over the interval $[0.1, 0.6]$[4]. The granularity of the grid is 0.05

[4]We have decided to restrict the range of the grid to the interval $[0.1, 0.6]$ because this

resulting in 11 firm types; we have slightly perturbed the grid in order to include the Cournot firm (output 0.3333). In the remainder of the paper we will indifferently refer to the 0.125 strategy as the JPM or the cooperative strategy. The outputs used in the simulations are shown in table 2.

Table 2: Set of outputs used in the simulations

0.1	0.15	0.2	0.25[a]	0.3	0.3333[b]	0.4	0.45	0.5	0.55	0.6

[a]JPM output
[b]Cournot-Nash output

In some simulations we let the firms choose their strategies (outputs levels) from the closed interval $[0.1, 0.6]$. In this case the strategy space is $\mathcal{S} = [0.1, 0.6] \times [0.1, 0.6]$.

2.3 Aspirations and learning

We have modeled firms as boundedly rational agents that use a very simple behavioural rule based on aspirations in their decision making. Each firm at each duopoly has an aspiration level of profits: if it is earning at least this level of profits it continues to adopt the same strategy it was previously using, if it is earning below aspiration profits it "experiments" a new strategy by imitating the strategy adopted by one of the firms in its neighbourhood. In our model the aspiration level is endogenous and equal to the average level of profits in the neighbourhood.

According to this behavioural rule a strategy can be considered as a routine that is kept in use as far as it provides to the firm a satisfactory level of profits. It should be noted that the experimenting firm does not take into account any strategic consideration when revising its strategy, it simply observes at random one of the firms in its neighbourhood and imitates its strategy (that could be the same strategy it was using). Consequently the probability for a strategy of being imitated by an experimenting firm depends exclusively on the number of firms in the neighbourhood adopting it.

If we define the profits of each firm at $d(x, y) \in \mathcal{D}$ as $\pi^f_{x,y}$, we have that the average level of profits in the neighbourhood $N^\rho_{x,y}$ of duopoly $d(x, y)$ is therefore:

restriction speeded up the simulations by leaving out many zero-profit pairs of strategies. Furthermore the outcome of the simulations is not affected by this restriction.

$$\bar{\pi}_{x,y} = \frac{1}{2\left|N^{\rho}_{x,y}\right|}\left[\sum_{f\in f^N} \pi^f_{x,y}\right]. \tag{1}$$

When $2\rho + 1 = K$ that is to say that the size of the neighbourhood equals the size of the lattice i.e. the interaction is global, the average level of profits is given by:

$$\bar{\pi} = \frac{1}{2K^2}\sum_{f\in\mathcal{F}} \left(\pi^f_{x,y}\right). \tag{2}$$

At every time t each firm f carries out its strategy revision procedure by comparing its profits with the neighbourhood average. The outcome of this comparison determines what we have called the learning state of the firm. We say, in fact, that a firm f is in learning state 0 and does not experiment if it is fulfilling its aspirations (it is earning a profit at least equal to the local average), on the other hand, a firm f is in learning state 1 and it does experiment if it is not fulfilling its aspirations (it is earning below local average profits). According to the learning state of the firms we can also partition the set of all firms into two mutually exclusive partitions $\mathcal{L}_0$ and $\mathcal{L}_1$ whose elements are defined in the following way:

$$\begin{aligned} \mathcal{L}_0 &= \{f \in f^N | \pi^f \geq \bar{\pi}_{x,y}\} \\ \mathcal{L}_1 &= \{f \in f^N | \pi^f < \bar{\pi}_{x,y}\} \end{aligned}$$

All firms in the set $\mathcal{L}_1$ (below aspiration firms) will choose the strategy to adopt in the following period by means of a simple imitative rule: they randomly sample a firm in their neighbourhood and adopt its strategy; as a consequence the probabilities of switching from any strategy to any other will depend on the frequency of adoption of each strategy in the neighbourhood, the higher the proportion of firms adopting a certain strategy, the higher the probability for that strategy of being adopted by below aspiration firms. Firms in $\mathcal{L}_0$ (above aspiration firms) will "stay".

In order to compute the switching probabilities let's first define the function $I(s,t,x,y)$ that returns the number of firms f in the neighbourhood $N^{\rho}_{x,y}$ of duopoly $d(x,y)$ playing strategy $s \in S$ at time t. Now we can calculate the proportion $P_t(s)$ of firms in neighbourhood $N_{x,y}$ adopting strategy $s \in S$ at time t as:

$$P_t(s) = \frac{I(s,t,x,y)}{2|N^{\rho}_{x,y}|}. \tag{3}$$

We are now in a position to define the probability $\mu^{0,t}_{x,y}(s,s')$ for a firm $f \in \mathcal{L}_0$ adopting strategy s at time t of switching to strategy $s' \in S$ at time $t+1$. This probability is given by:

$$\mu^{0,t}_{x,y}(s,s') = \begin{cases} 1 & \text{if } s' = s \\ 0 & \text{otherwise.} \end{cases} \tag{4}$$

On the other hand, the probability $\mu^{0,t}_{x,y}(s,s')$ for a firm $f \in \mathcal{L}_1$ adopting strategy s at time t to switch to strategy $s' \in S$ at time $t+1$ is given by:

$$\mu^{1,t}_{x,y}(s,s') = P_t(s'). \tag{5}$$

Note that when $2\rho + 1 = K$, i.e. the size of the neighbourhood coincides with the size of the lattice, firms take their decisions on the basis of global average profits and can, if unsatisfied, imitate the strategy adopted by any other firm in the "population"; in this case our model looses its locational features becoming a model of global interaction.

As in Kirchkamp (1996) and Hoffmann and Waring (1996), we have that the localization of action which happens at the duopoly level, differs from the localization of learning, which happens at the neighbourhood level. In fact, it is not clear why they should coincide: in many economic situations agents apply locally strategies which are chosen on the basis of global information; learning in this model is "social" rather than strategic.

In section 3 we also explore the case where firms compute their aspiration levels of the basis of global average profits, but restrict the "pool" of firms to imitate to the local neighbourhood. This particular specification of the width of the neighbourhoods can, in a way, be used to model situations where there are some institutions, (financial markets, for instance) that centrally process and distribute information about the performance of agents. This information is then used by agents in order to assess the effectiveness of their strategies, but given the presence of costs in observing the strategies adopted by other agents, their imitation capabilities are restricted to a local level.

2.4 Differential switching

We have also explored the implications of a different behavioural rule according to which firms earning below average profits do not carry the strategy revision procedure on with certainty. According to this *differential* rule unsatisfied firms experiment with a probability proportional to the distance between their current payoffs and the aspiration level. The fact that below average firms have only a probability rather than the certainty of switching to a new strategy may reflect the presence of difficulties by part of the firms in the evaluation of their performance (relatively to the average level of profits) and/or the presence of switching costs. We have modeled this probability as a decreasing linear function of the current profits of the firm: each firm f in the neighbourhood $N^{\rho}_{x,y}$ revises its strategy at time t with a probability $\alpha_t(\bar{\pi}_{x,y}, \pi^f)$ given by (dropping the firm superscripts):

$$\alpha_t(\bar{\pi}_{x,y}, \pi^f) = \max\left[0, 1 - \frac{\pi^f}{\bar{\pi}_{x,y}}\right]. \tag{6}$$

Thus the function that generates the switching probabilities $\nu_t(s, s')$ from strategy s to $s' \in S$ at time t takes the following form:

$$\nu_t(s, s') = \begin{cases} 1 - \alpha_t + \alpha_t P_t(s') & \text{if } s' = s \\ \alpha_t P_t(s') & \text{otherwise.} \end{cases} \tag{7}$$

2.5 Imitation with noise

The last learning rule that we consider is a simple modification of the imitative learning rule of section 2.3 that allows for some random experimentation by part of below average firms. In order to take into account this possibility into the model we have only to replace formula (5) by the following:

$$\mu^{1,t}_{x,y}(s, s') = \begin{cases} P_t(s') & \text{with probability } 1 - \epsilon \\ \frac{1}{S} & \text{with probability } \epsilon. \end{cases} \tag{8}$$

Note that only below aspiration level firms can randomly switch to a new strategy, above average firms will continue to adopt the strategy used in the previous period with probability one. This "noisy" version of our basic imitative rule should take into account errors by part of the firms in observing the strategies adopted by other firms.

3 Simulation design and discussion of the results

In this section we report the results of our simulations. In particular, we explore and compare the results of the simulations of the learning processes

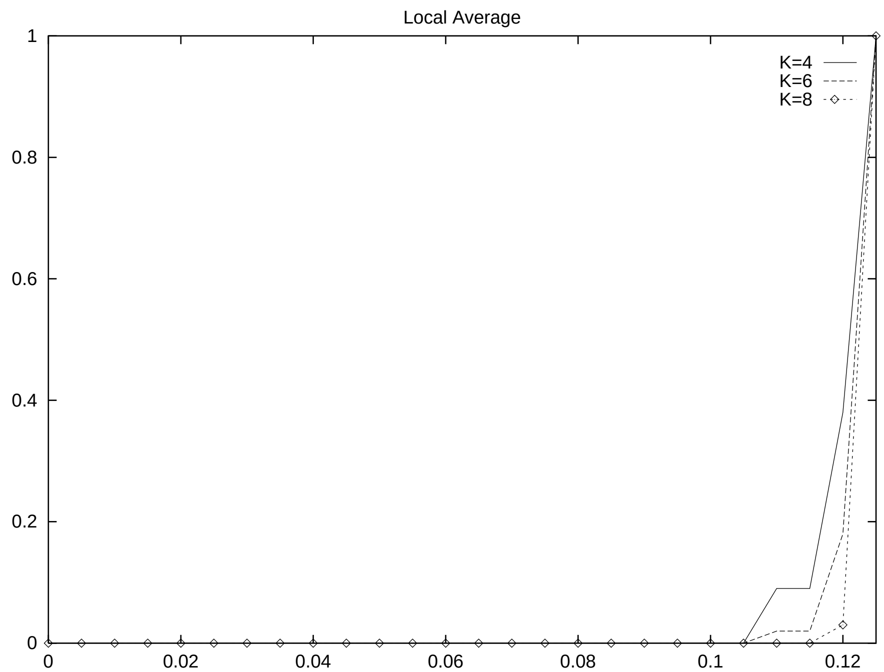

Figure 2: Cumulative distributions of final average profits. Evolution by imitation, discrete strategies, $\rho = 1$.

induced by the three behavioural rules defined in sections 2.3–2.5. For each of the learning rules we consider both the case where firms choose their strategies from a discrete set and the case where firms are allowed to choose strategies continuously form the interval $[0.1, 0.6]$. Then, for each of the resulting combinations of learning rules and strategy space type we run simulations for different values of the size K of the lattice. We have started all simulations from an initial random distribution of firms' strategies over the strategy space. In other words at the beginning of each simulation each firm at each duopoly on the torus chooses at random a strategy from the set of available strategies. At this point we let the computer carry on the evolutionary process until all the firms on the lattice satisfy their aspiration level requirement and we record the final average output which we can consider as a measure of cooperativeness among firms . For each of the considered combinations of learning rules, strategy space type and lattice size we computed 100 simulations in order to get a cumulative distribution of the final average profits.

3.1 Imitation without noise.

Figure 2 shows the cumulative distributions of final average profits for different values of the size of the lattice and for a neighbourhood radius $\rho = 1$. We can observe that the cumulative distributions tend to collapse toward the 0.125 value (the level of profits associated to the joint profit maximizing strategy) as the size of the lattice increases[5]. In other words the degree of cooperation among firms is positively related to the size of the lattice (and to the number of firms). When, for instance, the size K of the lattice is equal to 8, resulting in $2(8^2) = 128$ firms, we get perfect cooperation at all locations in 97% of the simulations. This percentage is sensibly higher than the one, 62%, that we get when $K = 4$.

The reason behind the relationship between neighbourhood size and degree of cooperation is fairly simple to understand. A sufficient condition for the cooperative strategy to spread over the lattice is that at time $t = 0$ there is at least a duopoly $d(x, y)$ where both firms produce the cooperative quantity $q^f = q^{f'} = 0.25$; this, in fact, ensures that both f and f' at $d(x, y)$ receive the maximal level of profits $\pi^f = \pi^{f'} = 0.125$ achievable in our simple Cournot model. As a consequence both firms f and f' at $d(x, y)$ are in learning state 0 and will never leave this state (and change their strategies) since their profits cannot, being maximal, fall below average[6]. On the other hand, all firms in the neighbourhood of duopoly $d(x, y)$ that earn below average profits can, with a strictly positive probability, imitate the cooperative strategy used by the firms at $d(x, y)$. Some of them, over time, will do it increasing, in turn, the number of cooperative duopolies in the neighbourhood, the probability of adopting the cooperative strategy, and the local average level of profits. This increase in the average profits will force some of the firms who earned profits higher than the local average, but still not cooperative to switch from learning state 0 to learning state 1. This evolutionary process will continue until all duopolies in the neighbourhood of $d(x, y)$ adopt the cooperative (JPM) strategy; finally, since all neighbourhoods are interconnected the cooperative strategy will propagate over the entire lattice. It is now easy to understand why the size of the lattice can affect the degree of cooperation. In fact, as the size of the lattice (and the number of duopolies) increases, the probability of having at time $t = 0$ a duopoly where both firms play the JPM strategy,

[5]Since the evolutionary process ends when all firms satisfy their aspirations, final average profits coincide with the final profits of each and all firms. In fact, at the end of the evolutionary process all firms end up choosing the same strategy and earning the same level of profits.

[6]The pair of strategies $q^f = q^{f'} = 0.125$ is, in fact, the global attractor of the dynamic process in the space of pairs of strategies (see Dixon, 1998)

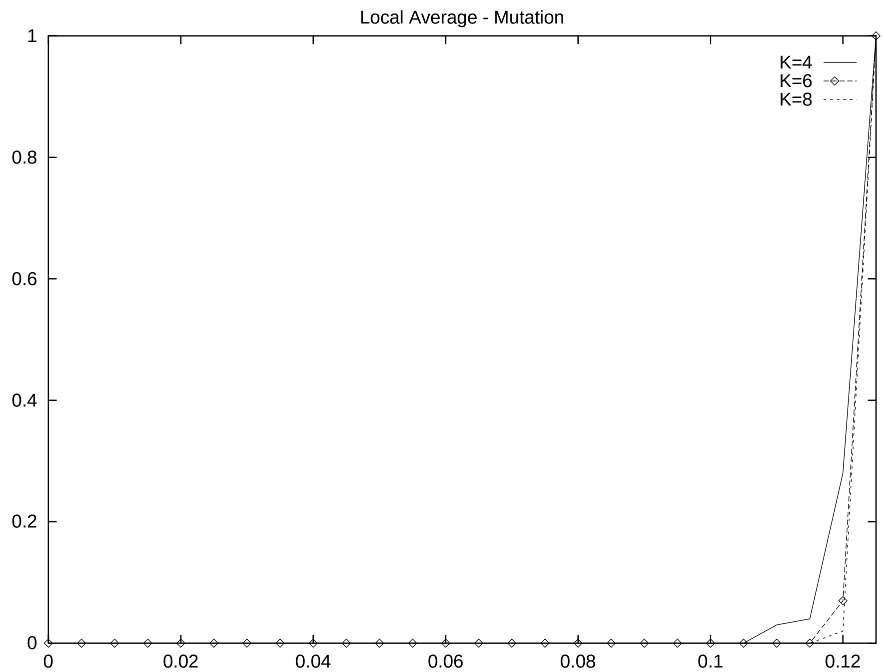

Figure 3: Cumulative distributions of final average profits. Evolution by imitation with noise, discrete strategies, $\rho = 1$.

and in general the probability than in a certain neighbourhood more than one firm plays the cooperative strategy also increases.

Global cooperation can also be achieved when there are no duopolies where both firms play the cooperative strategy at time $t = 0$. In this case, however, the cooperative outcome is less likely to emerge and its realization depends on the initial number of *firms* adopting the cooperative strategy. In this case, in fact, since the JPM strategy do not perform very well against other strategies[7] it is very likely that it will "perish" during the initial phases of the evolutionary process. But if the JPM manages to survive, because it is able to provide above average profits to at least some of the firms f adopting it, then it is possible that, if their competitors f' earn below average profits, they will imitate the JPM strategy establishing a new cooperative duopoly. It is also possible that two below average firms at duopoly by chance end up imitating the cooperative strategy adopted by one firm in their neighbourhood. The likeliness of all these events is directly relate to number of firms adopting the JPM strategy at the beginning of each simulation which is proportional to the size of the lattice.

[7]Against the JPM strategy, in fact, every other strategy $q \geq 0.25$ has a better performance as long as the market price is positive.

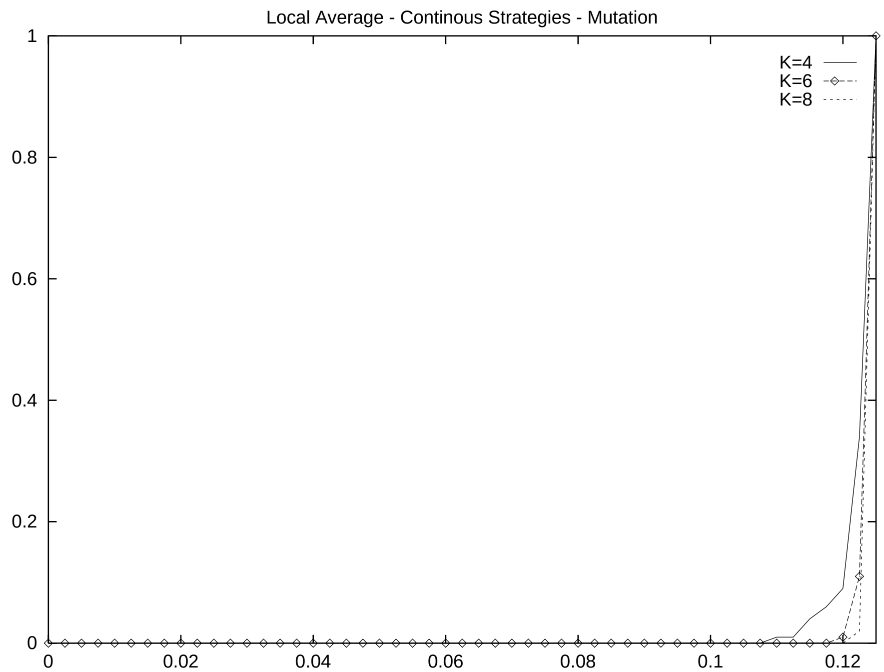

Figure 4: Cumulative distributions of final average profits. Evolution by imitation, continuous strategies, $\rho = 1$.

It must be pointed out that if no firm adopts the JPM strategy initially, then, unless some form of mutation or random switching is introduced into the model, there is no possibility for the cooperative strategy to establish and propagate.

3.2 Imitation with noise.

Figure 3 shows the case where random experimentation or mutation is introduced in the model. In particular it shows the outcome of the simulations where below average firms have a probability $\epsilon = 0.1$ of choosing a new strategy at random. The cumulative distributions of final average profits do not differ very much from the distributions relative to the no mutation case, but we can observe that when the size of the lattice is small $K = 4$ the distribution is more skewed to the right revealing an higher probability of getting the cooperative outcome at the end of the learning process. The fact that experimenting firm can randomly introduce in the "pool" of strategies (like the JPM) strategies that at the beginning of the simulation were not sufficiently represented or not represented at all or strategies that during the evolutionary process went "extinct" can increase to a great extent the probability of establishing a cooperative duopoly.

In Figure 4 we can observe what happens when we allow firms to choose strategies continuously from the interval $[0.1, 0.6]$. In this case too, the fact that at time $t = 0$ are "generated" $S = 2(K^2)$ different strategies and that some new strategies can come to life by random switching seems to increase tendency toward cooperation in small populations of duopolies.

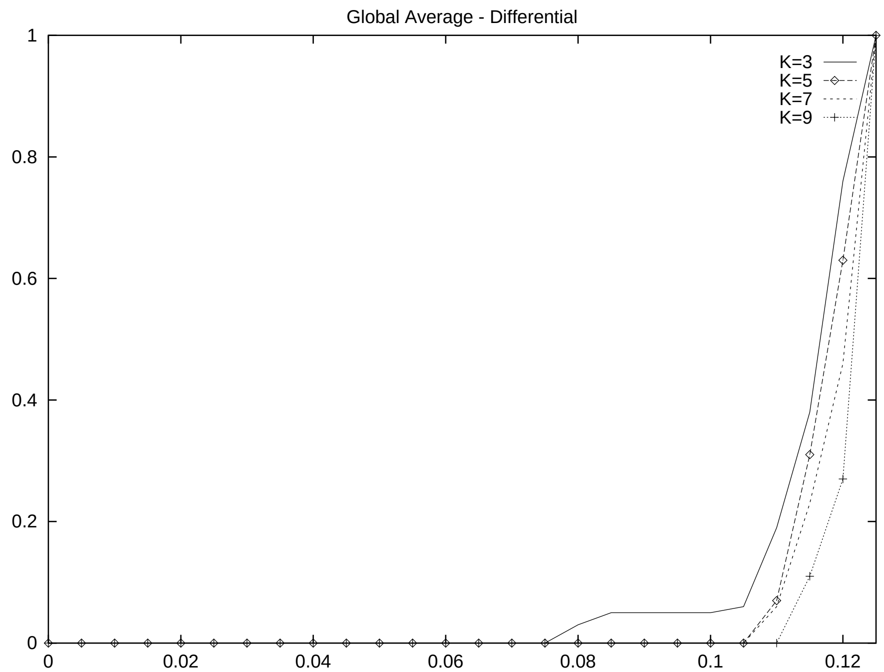

Figure 5: Cumulative distributions of final average profits. Evolution by imitation, differential rule, $\rho = 1$.

3.3 Evolution by imitation, differential rule.

Figure 5 reports the results of the simulations relative to the differential switching rule. As we can immediately observe, this rule, by introducing inertia in the strategy revision process of the firms reduces to a great extent the final degree of cooperation among firms. In fact, according to this rule below average firms can continue to adopt the same unsatisfactory rule with a positive probability. The effect of inertia is amplified when the size of the lattice is relatively small; the reduced rate of imitation do not allow to more cooperative strategies to spread over the lattice increasing the number of duopolies where both firms play sub optimal-strategies[8]

[8]For the simulation of the evolutionary processes characterised by differential switching we adopted a slightly modified stopping rule. According to this rule, the evolutionary

ρ	mean	std dev
1	100.40	26.91
2	164.95	57.88
3	287.81	200.40

Table 3: Average converge time of the evolutionary processes.

3.4 Neighbourhood size effects

In order to explore the effect of neighbourhood size on the time necessary for the evolutionary process to achieve convergence we have run simulations for different values of the parameter ρ. In particular, for a lattice size of $K = 10$, we have run 100 simulations of the imitative evolutionary process of paragraph 2.3 (no noise) for three different neighbourhood sizes: $\rho = 1$, $\rho = 2$ and $\rho = 3$. In Table 3 are reported the mean and the standard deviation of the average times (computed over the batches of 100 simulations) of convergence for the different neighbourhood sizes. As we can clearly see from the figures, there exist an increasing relationship between the interaction horizon (given by the parameter ρ) and the time necessary to evolutionary process to convergence. The explanation of this phenomenon lies in the fact that a lower number of firms in the neighbourhood facilitates the coordination toward the cooperative strategy and its propagation over the lattice.

4 Conclusions

In this paper we extend the analysis of Dixon and Lupi (1997) by including a spatial structure conditioning the learning interactions of the firms. We consider a locally interactive evolutionary model where agents (duopolists) adopt simple aspiration-based behavioural rules in their decision making. In our setting at each location on a torus there are two firms that repeatedly play a Cournot game and revise their strategies according to some imitative rule.

We find that under the general imitative rule and all its variations we study, the limiting distributions is characterized by (almost) all firms playing the cooperative strategy (the strategy that maximizes joint profits). We also find that the size of the learning neighbourhood does not affect this result: both global and local imitation lead to cooperation among firms. Local interaction only affects the time path towards equilibrium by sensibly accelerating it. We also observe that the degree of cooperation among firms

processes ended when no firm for 50 consecutive periods changed its strategy.

is affected by the size of the lattice, with cooperation increasing with size of the lattice and the number of firms. These results are in line with the theoretical results of Dixon (1998) and the numerical results of Dixon and Lupi (1997).

A Appendix

All simulations were run on a Sun SPARCstation 5. All computer code[9] was written by the author in Gauss ver. 3.2.18 for Solaris 2.4. In almost all cases the output of Gauss was post-processed with Octave ver. 2.0.5 and then plotted with gnuplot ver 3.5.

[9]Programs are available from the author upon request.

References

Anderlini, L. and Ianni, A. (1996). Path dependence and learning from neighbours. *Games and Economic Behavior*, 13(2):141–177.

Axelrod, R. (1984). *The evolution of cooperation.* Basic Books, New York.

Axelrod, R. and Hamilton, W. D. (1981). The evolution of cooperation. *Science*, 211:1390–96.

Blume, L. E. (1993). The statistical mechanics of strategic interaction. *Games and Economic Behavior*, 5:387–424.

Canning, D. (1992). Average behavior in learning models. *Journal of Economic Theory*, 57:442–472.

Dixon, H. D. (1998). Keeping Up with the Joneses: Competition and the evolution of collusion in an oligopolistic economy. Research Papers in Economics 1810, Center for Economic Policy Researh.

Dixon, H. D. and Lupi, P. (1997). Learning with a known average: a simulation study of different learning rules. Discussion Papers in Economics 97/18, University of York. Presented at the III International Conference on Computing in Finance and Economics. Stanford, June 1997.

Ellison, G. (Sep 1993). Learning, local interaction, and coordination. *Econometrica*, 61(5):1047.

Hoffmann, R. and Waring, N. (1996). The localization of interaction and learning in the repeated prisoner's dilemma. Technical report, Santa Fe Institute.

Jonard, N. and Yildizoglu, M. (1997). Thecnological diversity in an evolutionary industry model with localized learning and network externalities. *Mimeo University Louis Pasteur – Strasbourg.*

Kandori, M., Mailath, G. J., and Rob, R. (1993). Learning, mutation, and long run equilibria in games. *Econometrica*, 61(1):29–56.

Kirchkamp, O. (1996). *Evolution and Learning in Spatial Models.* PhD thesis, University of Dusseldorf.

Marimon, R. (1997). Learnin from learning in economics. In Kreps, D. M. and Wallis, K. F., editors, *Advances in Economics and Econometrics: Theory and Applications*, volume I of *Econometric Society Monographs*, chapter 9, pages 278–315. Cambridge University press, Cambridge UK.

Marimon, R. and McGratten, E. (1995). On adaptive learning in strategic games. In Kirman, A. and Salmon, M., editors, *Learning and Rationality in Economics*, chapter 3. Blackwells.

Siegel, S. (1957). Level of aspiration and decision making. *Psychological review*, 64:253–262.

Simon, H. (1947). *Administraive Behaviour: a Study of Decision-making Processes in Administrative Organizations.* MacMillan.

Simon, H. (1996). *The Science of the Artificial.* MIT press, Cambridge, Massachussets, third edition.