AXELROD MEETS COURNOT:

OLIGOPOLY AND THE EVOLUTIONARY METAPHOR PART 1.

Huw David Dixon*

Steven Wallis**

Scott Moss**

July 1994
Revised November 1994

*Economics Department, York University, York YO1 5DD U.K.
(hdd1@york.ac.uk; tel:+44-1904-433788).
** Centre for Policy Modelling, Faculty of Management and Business, Manchester Metropolitan University, Aytoun Building, Aytoun Street, Manchester M1 3GH, U.K. (s.moss@mmu.ac.uk; s.wallis@mmu.ac.uk).

ABSTRACT.
This paper explores the implication of evolutionary models (replicator dynamics) in a simple Cournot duopoly model. A firm type is a linear decision rule in which the firm's output depends on the other firm's previous output. First we run an Axelrod Tournament between firm types. The champion firm is a near profit-maximizer. Secondly, we allow social evolution to occur using replicator dynamics. Here we find that there are very strong forces leading towards a collusive or near collusive outcome, so long as there is not too much "noise" in the dynamics.

We would like to thank Ken Binmore, Larry Samuelson, David Collard, Martin Cripps and Ernesto Somma for their comments, and seminar participants at Warwick, Manchester and the Industrial Economics Study Group meeting at York University in April 1994 for comments on an earlier draft of the paper. We would also like to thank the Nuffield foundation and SERC for their support.

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). It has also been suggested that firm's rules of thumb might evolve over time in a manner akin to Darwinian evolution (Alchian (1950), Nelson and Winter (1982), Winter (1971)). This is a special case of the notion that human behaviour and social conventions in general can be explained (to some extent) by an evolutionary perspective, which dates back to "Social Darwinism"[1]. More recently game-theorists have made the connection between conventions, bounded rationality, learning and evolutionary models (Binmore and Samuelson (1992,1993,1994), Canning (1992), Kandori Mailath and Rob (1993), Peyton-Young (1993), Selten (1991), Van-Damme (1994), Weibull (1994)). The key popularizer of this idea in the general social sciences in recent years has been Axelrod ((1984), (1986)).

The central notion is very simple, and is closely related to the notion of bounded rationality. In the complex and uncertain world in which we live, agents follow "strategies", or rules which tell them what to do[2]. Different agents try out different strategies (and/or the same agents try out different strategies). Some strategies are more successful than others. Over time, successful strategies will become more common, either through a form of *propagation*, or *imitation*. Hence strategies that lead to firms being more profitable will tend to predominate over time. We can then explain the strategies of firms as being the result of such a process of social evolution.

In this paper, we apply the evolutionary approach to the standard model

[1]see Chattoe (1994) for a comprehensive bibliography and review of the literature in economics, AI and cognitive psychology.
[2]We use the terms "strategies", "rules of thumb", "conventions" and "heuristics" as synonymous in the context of this paper.

of Cournot duopoly, which is a *strategic* environment. A firm "type" is a particular decision rule, which gives the firm's current output as a function of the other firm's output in the previous period. We first run an Axelrod-Tournament between the firm types: each firm type plays all of the firm types. We then allow the population of firm types to evolve according to replicator dynamics: firm types that follow decision rules that are more profitable become more common. We allow the replicator dynamics to be noisy, in the sense that there may be some random mutation[3]

This paper seeks to use this evolutionary approach explain the behaviour of firms (their decision rules) in the simplest possible Cournot environment. This is an important enterprise for two reasons. First, the Cournot model is one of the canonical models of oligopoly, and the most widely used in economics. It is important for us to know what the implications of the evolutionary approach are for this model. Secondly, it enables us to employ the evolutionary approach within the context of a well-known model with practical applications. The use of the evolutionary approach in game theory is often restricted to abstract and simple games which have no direct or obvious connection with the models used in standard economic theory. Most research (including that of Axelrod) has been done on the prisoner's dilemma: whilst this is a common model, it is also very special in that it possesses a strictly dominant strategy[4].

[3]Here the mutation takes the form of firms mutating to en existing firm type, not mutating into a *new* firm type.

[4]The main related literature on evolutionary models since Axelrod has been in modelling the prisoner's dilemma using the notion of finite automata. The idea here (following Abreu and Rubenstein (1988)) is that rational players select a finite automaton to play the repeated prisoner's dilemma (Binmore and Samuelson (1992), Linster (1992, 1994), Nachbar (1992), Probst (1992)). The use of the finite automaton is to capture the notion of *complexity* of the strategy employed: here we use linear decision rules with one period memory to model possible strategies. The Prisoners' dilemma is a very special game,

The fundamental question we seek to answer is: what types of behaviour would we expect to find duopolists following in a Cournot environment? The traditional reasoning of Cournot himself[5] or of game theorists using the model of "rationality with common knowledge" predicts the textbook Cournot-Nash outcome. Others have argued for a range of behaviour (conjectural variations), more collusive outcomes, and the possibility of strategic behaviour leading to leader follower outcomes as in the Stackelburg model. Axelrod himself argued that under certain conditions (see (1984 part 4), cooperation can and does evolve in a wide range of social and biological situations[6]. In this paper we are particularly interested in this *Cooperative Hypothesis*: namely that evolution will tend to promote cooperation. In the context of the Cournot duopoly model, cooperation means collusion and joint profit maximization.

The conclusions of the paper are really quite simple and quite far reaching. The first main conclusion relates to the type of firm which does best in the Tournament. This is an environment where there are lots of different types of firms, some quite unconventional. We find that the "Champion" firm in this context is a "near profit maximizer"[7]: indeed, the

which has non-cooperation as the *dominant* strategy: it is thus very unlikely that evolution will lead to cooperation. As Kandori, Mailath and Rob (1993), that any Darwinian process must lead to the dominant strategy having a population share of unity (see Theorem 2, p.43). In coordination games of common interest, where there are multiple pareto ranked equilibria, KMR show that the Darwinian process selects the equilibrium with the largest basin of attraction (Theorem 3). A direct comparison with our results is not straightforward, since KMR employ a "long-run" equibrium process with mutation grafted on to the Darwinian mechanics.

[5]See Leonard (1994) for a discussion of the relation of Cournot's original work to Nash's.

[6]For a discussion of what Axelrod showed, see Linster (1990, 1992), Nachbar (1992), Binmore (1991), Binmore and Samuelson (1993a).

[7]"Near profit maximizing" means that the firm has a decision rule in which its output is chosen to nearly maximize profits given the output of the other firm in the previous period.

standard Cournot best-response function does very well. *Near profit maximizing behaviour is a robust decision rule in an environment where a firm might meet a wide variety of firm types*. If we allow a process of evolution to take its course, we find a very strong tendency towards the collusive outcome: indeed with no noise in the replicator dynamics, joint profit maximization is a very common outcome (although not universal). *There are strong evolutionary forces driving behaviour towards cooperation*. However, as the replicator dynamics become noisier, we find that the evolutionary simulations move away from the collusive and towards the near profit maximizing behaviour that performed best in the tournament. We would interpret this result to be very much in support of the Cooperative hypothesis in the context of Cournot Oligopoly. The exercise we have undertaken in this paper is of its nature preliminary: the model can be generalized in many directions. However, we believe that it constitutes a vital first step in applying and developing the evolutionary metaphor in the context of oligopoly theory.

The plan of the paper is as follows. In section 1 we outline the basic Cournot Duopoly model, and briefly review the standard theoretical results in section 2; in section 3 we outline the method for generating firm types, and describe the Axelrod Tournament and its results; in section 4 we motivate and describe the replicator dynamics, and evaluate the results of evolutionary simulations; in section 5 we briefly relate our results to related literature.

1: The Basic Market Model

We will be examining the simplest model of Cournot duopoly. In this, there are no costs, and the market price P is a linear function of the two outputs x_i with the slope and intercept normalized to unity:

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

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:

$$U_i(\mathbf{x}) = x_i.(1-x_i-x_j) \qquad (2)$$

For practical purposes, this enables us to concentrate on the set A: $A \equiv \{\mathbf{x} \in [0,1]^2\colon 1-x_i-x_j \geq 0\}$. This is depicted in Figure 1. It is useful to remind ourselves the standard reference points on this.

|Fig 1 here|

The edge $1-x_1-x_2=0$ is the Walrasian set, where the outputs sum to unity yielding price 0 (which equals marginal cost). The Cournot-Nash outcome occurs at point C where both firms produce 1/3, and the price is 1/3. The monopoly point for firm 1 (2) is the point M_1 (M_2) where the firm produces the monopoly output 1/2, and the other firm produces nothing. The line segment $1-x_1-x_2 = 1/2$ is the set of Joint profit maximizing outputs. The 45° line represents symmetric outcomes, and M is the symmetric joint profit maximizing outcome. The points S_1 and S_2 represent the Stackelburg points for firm 1 and firm 2, where the leader produces 1/2 and the follower 1/4 (and hence the price is 1/4).

Firms play the duopoly game using a *decision* rule. A decision rule maps the information that the firm has onto its choice of output. In this paper we view the decision rule as a primitive, a "black box". It links action to information, and hence runs together the "strategy" of the firm and the process of reasoning leading to that strategy. 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. All of these restrictions could be relaxed. This paper is a first attempt to apply this methodology to duopoly, and we have opted for the simplest and most familiar. With all three restrictions, the model is not only simple, 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. Hence the appearance of the name "Cournot" in the title of the paper.

In our approach a *Firm type* is a decision rule, and under linearity can be represented as a pair of parameters, the *intercept* h_o and the *slope* h_1. A firm type i is thus characterized by a pair $\{h_{0i}, h_{1i}\}$. It should be noted that in our approach, a firm type (decision rule) is a primitive. This contrasts with much of oligopoly theory which *derives* the decision rule from some underlying hypotheses (for example profit maximization, conjectures about the other firms etc.). We can consider various standard types of decision rule, represented in Table 1:

Firm type	intercept	slope
Myopic Cournot Profit Maximizer	1/2	-1/2
Walrasian Duopolist	1	-1
Stackelburg Sticker	1/2	0
Cournot Sticker	1/3	0
Joint Profit Maximzer/Copy Cat	0	1

TABLE 1: Some Standard Decision Rules

The Myopic Cournot Profit Maximizer (MCPM) is derived under the hypothesis that the firm believes that the other firm's output in period t will equal that in t-1, and chooses its output to maximize (current) profits given that conjecture. The Stackelburg Sticker (SS) produces the Stackelburg Leaders output 1/2 every period. The Cournot Sticker (CS) produces the Cournot-Nash output every period. Clearly, the two "stickers" have a decision rule that commits them to a constant output, with no "reaction" to the other firms actions. The "Copy cat" reaction function is the analog of Axelrod's Tit-for-Tat strategy in his treatment of the prisoner's dilemma: firm i produces the other firm's output in the previous period. However, the space of decision rules is much larger. In principle, we can allow any decision rule that passes through the set A (hereafter the "unit triangle"), or indeed even ones that do not so long as they pass through the positive orthant (although to avoid trivialities, we rule out the latter in this paper).

In our analysis, we will want to talk about decision rules in economic terms in respect of their intercept and slope coefficients. Recall that the market game summarized by (2) is one where firms' outputs are *strategic substitutes*: the marginal profit of one firm's output is decreasing in the other firm's output (hence the MCPM is downward sloping). 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_0). 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 maximizing 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 *cooperative*: 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 cooperative behaviour. Firm types with 0 slopes will be noted *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 maximization, 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: Traditional Analyses of the Cournot Duopoly Problem

There have been several papers which have analyzed the notion of equilibrium in this sort of model. First, there are the well known results which consider subgame-perfect equilibria in the finitely or infinitely repeated Cournot Duopoly model. These approaches are based on the assumption that firms behave as if they were the "rational agents" of game-theory, assuming Common Knowledge, that firms have infinite memories, can choose strategies of arbitrary complexity, have unlimited powers of reasoning, and can at the commencement of any subgame change their strategy for the rest of the game. The results of this analysis are familiar: in a finitely repeated game, there is a unique SGP equilibrium, with both firms producing the Cournot-Nash output (1/3) in every period. There are two cases in the infinitely repeated case. With no discounting, any pair of feasible payoffs can be sustained by a SGP equilibrium, so long as each player receives its "security level" (zero in this case). With discounting, the matter is more complex: for a discount rate of 0 (i.e firms ignore the future) we are back to the Cournot-Nash outcome as the unique SGP equilibrium; in the limit, as the discount rate tends to 1, we are back with the no-discounting case (all outcomes are possible).

Secondly, there has been the analysis of dynamic reaction function equilibria (Friedman (1977), Stanford (1984a,b), Robson (1986). The key results are Stanford's: with decision rules like ours, the unique subgame perfect equilibrium in the infinite horizon game with discounting is the Cournot Sticker (intercept 1/3, zero slope). With no discounting, any output pair with strictly positive outputs earning positive profits is sustainable as a SGP.

Thirdly, there is the analysis of static "reaction-function" (or supply function) equilibria: see for example, Hart (1984), Hey and Martina (1988), Klemperer and Meyer (1989), Bone and Dixon (1992). These static models can be seen as equivalent to the Stanford case with no-discounting (although there is no formal proof of this at present). Again, any output pair with strictly positive outputs and profits can be supported as a Nash-equilibrium.

We draw two points of interest from the traditional literature. First, with linear decision rules like ours, the theory points to two possible scenarios: first, from Stanford's (1984a) paper with discounting we get the "trivial" outcome of the unique SGP equilibrium with Cournot outputs in each period; second, from the case with no-discounting and the static models we get the case of almost any outcome being supported by an appropriate decision rule pair.

3: The Method.

Our method consists of three stages. In stage one we generate a set of firm types. Second, given the set of firm types, we then run an *Axelrod Tournament*, by which we mean that every firm type plays every other firm type in a "round-robin" tournament. Thirdly, we then run an evolutionary algorithm based on replicator dynamics. Let us now consider in detail these three steps in turn.

3A: Generating firm types.

The way we have chosen to generate firm types is based on the HMBD result (Hey and Martina (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 $\mathbf{x'} \in \text{int}A$, 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 $\mathbf{x'}$. Hence the tangent at $\mathbf{x'}$ is characterized by the slope and intercept terms:

$$h_0 = 2x_1 + 2x_2 - 1 \qquad \text{(intercept)} \qquad (3a)$$

$$h_1 = (1 - 2x_1 - x_2) / x_1 \qquad \text{(slope).} \qquad (3b)$$

Fig.2 Here

This is depicted in Figure 2.

|Figure 2 here|

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 $\mathbf{x}'$. That is, $\mathbf{x}'$ solves the problem:

$$\max U_1(\mathbf{x}) \text{ subject to } x_2 = h_0 + h_1 x_1$$

where the slope and intercept are as in (3). A Nash-equilibrium in decsision rules occurs at $\mathbf{x}'$ when firm 1 chooses as its reaction function the tangent to firm 2's isoprofit curve at $\mathbf{x}'$, 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).

Our use of the HMBD analysis to construct an algorithm to generate firm types can be viewed in a variety of ways. First, if we consider the papers analyzing static reaction-function/supply-function equilibria to be useful, then this algorithm generates reaction functions that are potential candidates for such static equilibria. Secondly, one can simply view it as a more or less arbitrary method of generating linear decision rules that pass through the unit triangle. We remain to some extent agnostic, and the interpretation of the paper does not depend on it. However, there are two main advantages of the algorithm used. First, the algorithm is defined over a compact convex set. Other alternatives might be (for example) to do a grid-search over the intercept and slope parameters. Whilst the intercept might reasonably be restricted to [0,1], there is no obvious reason why the slope should be bounded: indeed, with our algorithm it is unbounded from above. Second, our method 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). We are currently experimenting with alternative algorithms for generating firm types, to see how sensitive the results are.

We believe that it is essential to have a large and diverse set of firm types. If we consider the original Axelrod Tournaments, there were relatively few different types (63). As Linster (1990) has shown, the results of Axelrod are very sensitive to the number of agents and the precise strategies. However, all our tests have shown that with a large number of firms the outcome of the Tournament is very robust: the outcome does not look much different if you have 1,000 firm types or 20,000.

The algorithm for generating the firm types is implemented using a grid search on the unit triangle. We specify the *granularity* of the grid, which is the distance between two adjacent points as you move horizontally or vertically around the grid (or "lattice"). We only generate firm types from points in the interior of the unit triangle. If the granularity of the grid is (for example) 0.01, then we generate approximately $100^2/2$ (less the 300 points on the edges) firm types. When reporting the results below, we will state both the granularity and the total number of firms.

Using this algorithm we can divide up the area of the unit triangle into subsets according to the slope (or intercept) that it generates. This is done in Figure 3:

|Figure 3 here|

Two things are worth noting. First, the set of decision rules that are best responses to themselves are to be found on the 45% line. Secondly, the decision rule generated by a point $a \in A$ is the best response to the rule generated by the point a' which is the reflection of point a in the 45% line. Thus insofaras the algorithm generates a set of points that are symmetric about the 45% line, each decision rule generated is a best response to another decision rule in the set, and no decision rule (however strange) is dominated.

3B: The Axelrod Tournament[8].

Given our set of firm types, we then run an Axelrod Tournament between them. That is, each firm type meets each other firm type (including itself) 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):

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

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

We have tried a couple of ways, *simulation* and *analysis* to evaluate the firms' profitability in the constituent game. One method is to simulate each game: assume a starting point, use (4), and have a rule for stopping: the payoff can then be the average over the whole or part of the simulation. This has various problems: the outcome maybe sensitive to initial positions, and using the last few periods profits might introduce an "endpoint" bias, and so on. In the end we went for an analytical method: we used the eigenvalues of the system (4) to classify the dynamic properties of the system, and used simulation only where unavoidable and sensible (for about 5% of constituent

[8]Clearly, our Tournament differs from Axelrod's, in that he did not use a population of strategies generated by an algorithm like ours. Rather, he invited strategies from real individuals. An approach more similar to Axelrod in an oligopoly environment was conducted by Fader and Hauser at MIT in 1985-6 (Fader and Hauser (1988)). Our approach differs from both in having a large number of strategies systematically generated

games). The exact algorithm used is described in the appendix.

Having run the full Axelrod Tournament with n types of firm, we then have an n×n matrix of payoffs: each element in the matrix corresponding to the payoff when the relevant row firm meets the relevant column firm. More formally, let u_{ij} be the payoff of firm i when it meets firm j. We can define the n×n Tournament payoff matrix **T** as $[u_{ij}]$ where i,j=1...n, and i denotes "row", and j "column". The *average Tournament Profit* for firm i is then simply the arithmetic average of firm i's payoffs summed over row i:

$$ATP_i \equiv \sum_{j=1..n} u_{ij} \qquad (5)$$

The Average ATP of the Tournament is simply the arithmetic average of all firms individual ATPs.

A useful statistic for evaluating the profits of the various firm types is to compare it to the largest possible profits that could be earned if a "Superfirm" was able to choose its decision rule optimally for each individual firm it meets. In the Tournament, each firm type presents the *same* decision rule to all comers. In fact the calculation of Superfirm profits is very simple with our algorithm for generating firms. Since we generate type i f take the tangent on firm 1's iso-profit curve at $\mathbf{x}_i$, it follows that if firm 1 were faced with this decision rule its highest profits will be its profits at $\mathbf{x}_i$, if we assume that the profits are evaluated at the intersection of the decision rules[9]. The equation for the superfirm profits is therefore: $(1/n)\sum_{i=1}^{n} x_{1i}\,(1-x_{1i}-x_{2i})$. The ratio of actual to superfirm profits represents the cost to the firm type of not having the flexibility to tailor its decision

[9]An alternative would be for the Superfirm to solve the dynamic optimisation problem given the the other firm's decision rule. This sort of problem was analyzed by Friedman (1978).

rule to each different opponent type, and is hence an indicator of the degree of Bounded Rationality.

We have run many Tournaments. We will report only the largest Tournament. This was run with a grid search of granularity 0.005, which generated a total of 19,702 firm types (there are 200 intervals along each of the axes of the unit triangle, and the unit triangle contains $200^2/2 = 20,000$ firms: the number of firms is smaller because our search algorithm left out most of the boundary points). Clearly, this is a massive computational task, involving almost 200 million individual market games between oligopolists.

This size of Tournament generates a massive amount of information, but the results can be simply represented thanks to our algorithm generating firm types. In figures 4a,b we have have a contour map: each point on the map represents a firm type, and the height of the contour it is on represents its Average Tournament Profit (ATP).

\Figures 4a,b here\

It is worth spending some time on this figure. Note that the isoprofit contours are smooth. The "champion" of the Tournament which earns the highest ATP is in fact located at the point (0.435, 0.310). This corresponds to the decision rule with slope −0.409 and intercept 0.48. The ATP of the Champ is 0.07466. This represents 88% of superfirm profits, which is surprisingly high given the wide variety of firms in the population. Furthermore, the Champ earns almost twice the average ATP.

If we look at the black areas in Fig 4b, we can see two areas: one is the southwestern black valley, the other is the narrow strip close to the Walrasian edge. The southwestern black valley is populated by firms with large positive slope coefficients: the reason they do so badly is that most of their duopoly games are unstable, and often lead to zero profits.

In table 2 we present the performance of the Champ in comparison to the firms in table 1 (or the closest firms to them in the grid, market with a *):

	X_1	X_2	intercept	Slope	ATP (4 s.f.)
Champion	0.435	0.310	0.49	-0.414	0.07466
MCPM	0.5	0.25	0.5	-0.5	0.07342
JPM	0.25	0.25	0	1.0	0.03588
Walrasian Duopolist*	0.5	0.495	0.99	-0.99	0.005066
Cournot Sticker*	0.335	0.335	0.34	-0.0149	0.06164
Stackelberg Sticker	0.25	0.50	0.50	0	0.05203

TABLE 2: THE LARGE AXELROD TOURNAMENT

Granularity 0.005; 19,701 firm types; superfirm profits 0.08375 (4sf).

Average ATP 0.03775.

The most striking feature of table 2 is to note that the MCPM does so well, and that it is quite "close" to the Champ, in terms of payoff and the intercept/slope coefficients (the profits of the MCPM are in fact 99.00% of the Champs). The other reference firms do very badly in comparison: in particular the JPM rule is less than half of the superfirm profit, and even below the average ATP.

The reason for the fact that the Champ and the MCPM are so close is, we believe, that the Champs near "profit maximizing" behaviour is very robust. In the population, there are lots of very weird and wonderful firms. To do well, the Champ has to perform well with all of them. The near profit maximizing behaviour of the Champ ensures at least that the output it ends up with in each of the constituent games is almost a best response to the output of the other firm. The copy-cat behaviour of JPM is not very good in this environment: if it meets a weird firm, it will imitate it and hence not do very well. JPM can only succeed with rules that do well against themselves. *We would therefore conclude that in an environment with a wide variety of firms, near profit maximization (as represented by the Champ) is a very robust decision rule for the Cournot duopolist.*

In table 3, we summarize the range of behaviour that is represented by the top 147 firms (the top 0.75%). The first thing to note is the range of behaviour represented by the intercept and the slope. Whilst the top firms can all be reasonably called "near" profit maximizers, they clearly tend to have less negative slopes than the MCPM: their behaviour is more cooperative.

	min	max
X_1	0.4	0.475
X_2	0.27	0.35
Intercept	0.44	0.55
Slope	-0.5109	-0.317
ATP	0.0740	0.07466

TABLE 3: THE RANGE OF VALUES FOR THE TOP 147 FIRMS.
Granularity 0.005; 19,701 firms; Superfirm 0.08375.
All decimals to 4 s.f.

We have completed the first stage of our analysis. We have generated a wide range of firm types, and conducted an Axelrod Tournament amongst the firms. We have calculated the ATP, which is based on the notion that all firm types are equally well represented. However, the range of ATP is very large: from almost nothing, to the profits of the top firms in table 3. The next stage in our argument is to allow for successful decision rules to become more common, and the less profitable to wither away.

5:Evolution in the duopoly model.

There has been much discussion of the evolutionary metaphor in economics in recent years[10]. The notion of evolutionary dynamics in biology is based very much on the notion of reproduction: successful species or genes tend to become more common because they give rise to more progeny. In the context of *social* evolution, such mechanisms of propagation might also be present: successful firms grow and diversify, their managers circulate, good firms take-over bad, unsuccessful firms go bust. However, in social evolution, there is also the mechanism of *imitation*: firms tend to imitate the more successful practices of other firms (as in the practice of "benchmarking"). The actual processes involved are very complex, and we make no attempt to develop new theoretical results here. We *do* seek to explore the implications of an existing model of social evolution in a new context. This section is divided into three parts. In part A, we describe and discuss the evolutionary dynamics; in part B, we describe he results; in part C we interpret the results.

5A: The Evolutionary Algorithm.

Having run the Tournament, following Axelrod we proceeded to apply an evolutionary algorithm. First, consider the Tournament Payoff matrix **T**. The vector **Z** gives the proportions Z_i of each firm type i (clearly, $Z_i \in [0,1]$ and $\Sigma Z_i=1$). We start from an initial condition $\mathbf{Z}_0$ and then represent evolution

[10]From the game-theoretic perspective, see Selten (1991) for an excellent discussion, and indeed the other papers in that special edition of *Games and Economic Behaviour*. Binmore (1992, chapter 9) and Binmore and Samuelson (1992a) are also excellent and aimed at the general reader. Chattoe (1993) adopts a wider perspective on this approach.

using *replicator* dynamics. In particular, we consider the evolution to be run over *iterations* s=1....S. Note that the iteration s can be thought of as time periods, but they are not the same as the periods t within the market simulation which generated the Tournament payoff Matrix **T**. We can think of the time-periods within a simulation occurring within one round of evolution. Thus the story underlying the model is that in each iteration individual firms are randomly matched (in line with population proportions), and play the duopoly game with their partner. In the next round this occurs again, but with new proportions, and so on.

The basic replicator dynamics we adopted was (following Van Damme (1987) and Binmore et al (1994)):

$$Z_{is} = Z_{is-1}(1-d) + Z_{is-1}\left[(U_{is-1} - u_s)/ u_s\right] + (d/n) \qquad (7)$$

where $d \in [0,1]$, U_{is} is the average payoff of firm i in iteration s:

$$U_{is} = \sum_{j=1}^{n} Z_{js}.u_{ij}$$

and u_s is the average payoff of all firms in iteration s:

$$u_s = \sum_{j=1}^{n} Z_{js}.U_{is}$$

We also have the constraint that the Z_i's are non-negative, so we allowed for extinction and appropriately renormalized $\mathbf{Z}_s$ so that it sums to unity in each round. The parameter d is a "noise" parameter. Let us first consider the case with no noise, d=0.

Without noise, the change in proportions depends on the absolute difference of profits from the population average[11]. This is a plausible specification from our perspective: better (worse) rules become more (less) common; the fact that the extent to which they change depends on their current proportion captures the notion that firm types with larger population proportions are more "visible" and likely to be imitated (if good), or avoided (if bad)[12]. Assuming that there exists an attractor(s) the replicator dynamics converges to a situation where all surviving firms have equal average profits: $\mathbf{z}^*$ is such that $U_i=U_j$ for all i,j=1..n when $\mathbf{z}_{i,j}>0$. We can interpret firm types as "pure strategies" in a game: if the replicator dynamics converge, then $\mathbf{z}^*$ is a Nash-equilibrium given the set of strategies that were present with non-zero proportions in the initial vector $\mathbf{z}_o$[13]. One possible interpretation of our process is that we have a duopoly playing mixed-strategies over firm types, and the replicator dynamics is a way of finding mixed-strategy equilibria.

When d>0, we have noisy replicator dynamics. In each period, a proportion d of firms randomly switch strategies: thus each firm type i looses $d.Z_i$. For simplicity, we assume that the these firms are randomly allocated across firm types, so that each firm type gains d/n (alternative assumptions are possible). This sort of "noise" can be interpreted as random "mutation": however, note that this mutation does not generate new firm types - it merely

[11]Note that this is identical to the equation $Z_{is}= Z_{is-1}\left[U_{is-1}/ u_S\right]$ in which form it appears in Linster (1992) and Nachbar (1992).
[12]Binmore, Gale and Samuelson (1994) tell a specific story of social evolution that gives rise to replicator dynamics. Every period, a certain proportion of firms change their strategies if those strategies are below an exogenous (random) aspiration level. Firms which change from their own strategy adopt a new strategy from the set of existing strategies in proportion to each strategy's proportion.
[13]See for example Nachbar proposition 1, p.313 (1992).

serves to randomly allocate amongst existing firm types (as in Linster (1994, pp.348-353), and Binmore et al (1994)).

5B: The Results of a the evolutionary simulations.

Two practical issues arise when we implemented the replicator dynamics. First, we were unable to use the payoff matrix of the n=19,702 Tournament which generated Fig.4 (this lies beyond our current computational power). Therefore, we adopted a smaller sample of firms: 1,176 firms generated by a grid of granularity 0.02. This is the Tournament represented in Table 4:

	X_1	X_2	intercept	Slope	ATP (4 s.f.)
Champion	0.44	0.30	0.48	-0.409	0.07516
MCPM*	0.5	0.26	0.52	-0.520	0.07415
JPM*	0.26	0.26	0.04	0.846	0.04064
Walrasian Duopolist*	0.5	0.48	0.96	-0.980	0.01796
Cournot Sticker*	0.34	0.34	0.36	-0.059	0.06428
Stackelberg Sticker*	0.26	0.50	0.52	-0.077	0.05594

TABLE 4: THE TOURNAMENT USED FOR THE EVOLUTIONARY DYNAMICS.
Granularity 0.02; 1,176 firm types; superfirm profits 0.08500
Average ATP 0.03905

Superfirm profits and the average ATP are a little higher in the smaller Tournament. This is partly due to the fact that we have excluded the edges from both grids, so that the larger Tournament covers a larger area, and includes firms closer to the edge (firms closer to the edges tend to earn very low profits).

First, we will describe the evolutionary simulations we have run for different values of d from the initial vector $Z_{i0} = 1/n$, which are summarized in Table 5 (we discuss simulations with different initial positions below). Let us first define some of the terms. The second column gives the *mean firm* at the end of the simulation, thus we weight the location of each firm type by its proportion: $\sum_{i=1}^{n} Z_i.\mathbf{x}_i$. The third column gives the *modal* firm at the end of the simulation. The next four columns describe the modal firm, its intercept, slope, and its proportion and average profits at the end of the simulation. SF gives the superfirm profits at the end of the simulation, and is defined as in the Tournament, except with the end-of-simulation weights. The last column reports for information the number of iterations we ran: when d=0 the simulation ran until surviving firms all earned the same profits (to 16 s.f.); when d>0, the simulations ran until the firms proportions were constant (to 16 s.f), or we pulled the plug (these are really just for information). All simulations converged, except the case of d=0.001, which we discuss in detail below. Let us discuss the cases in turn.

TABLE 5 HERE

d=0:

Here we have a very clear result. There is one surviving firm, and in terms of our grid of firm types, it is as close to the JPM firm as we can get and hence appears as JPM* in table 4. Note also that it is an optimal response to itself (average profits equal superfirm profits). As can be seen from the intercept and the slope, this firm is highly cooperative, and indeed its profits are within 0.0002 of the real JPM. We do not claim that this firm type is a global attractor, and discuss the role of initial position below. However, this result clearly supports the Cooperative Hypothesis, in that the replicator dynamics clearly lead to an almost perfectly collusive outcome. From table 4, it can be seen that the survivor JPM* does not do well in the Tournament: its profits are about average. JPM* essentially behaves by imitating its opponents, and in the Tournament this involves imitating lots of weird firms that earn very low profits. In order to prosper, JPM* has to wait until the replicator dynamics have eliminated the weird firms, and it can eventually come to the fore.

d=0.00001,d=0.0001:

When there is noise, all firm types survive with at least a share of (d/n) from (7), and the mean and modal firm types differ. These two levels of d obviously represent very low levels of noise. The first observation to make is that the average profit is slightly lower than in the no-noise case, but still very closer to the JPM profit relative to the Cournot payoff. In order to measure the degree of collusion, note that between the Cournot payoff 0.1111 and the JPM payoff of 0.125 there is a difference of 0.0139. We can express the average profit of the modal survivor by its position on the "collusion scale" with Cournot at 0 and JPM at 100. With d=0.00001 average profits of the modal firm are at 98.56%: with d=0.0001 they are at 87.05. By any judgment, these represent very collusive outcomes. However, as we move to d=0.0001, note that the modal firm is becoming less collusive.

d=0.001: In this case we found that the evolutionary simulation did not converge, but with over 2,000,000 iterations we found a regular cycle. There are three firms that predominate over the cycle, but their proportions fluctuate. These three firms are given in table 5, with the last column giving their maximum proportion over the cycle, which lasts about 12,000 iterations. Figures 5a,b represent the cycle over 12,000 iterations well into the simulation (the scale does not start at time 0): the time scales are the same for both figures.

Location	Intercept	Slope	Maximum Proportion
0.3,0.3	0.2	0.333	0.6296
0.28,0.3	0.16	0.5	0.3818
0.38,0.3	0.36	-0.158	0.3183

TABLE 6: THE THREE MOST COMMON FIRMS WHEN d=0.001.

FIGURES 5a,5b HERE.

First note that the average profits of all three firms are always almost the same: they differ only very slightly. These small differences are enough to generate the very long cycle. Secondly, the profits of these three firms fluctuate over a range 0.1161 to 0.1187: the average over the cycle is 0.1177 (note that the vertical scale in Fig 5b starts at 0.1160). In terms of the degree of collusion this represents a point that is intermediate between the Cournot and the JPM outcomes: in terms of the 0-100 collusion scale introduced above, the average profits yield a score of 47.48% (and range from 35.97% to 54.68% over the cycle). This is not surprising, since the two most common firms are both fairly cooperative. Only the third (and least common) firm has a negative slope (although still well above -0.5).

d=0.01;d=0.1:

As the level of noise increases to 0.01 and 0.1, we can see that the outcome is much more competitive. At d=0.01, the modal firm is still cooperative, but its proportion of the final population has fallen to a little over 50%, and its average profits are almost at the Cournot level. When the level of noise reaches 0.1, the outcome has flipped over to a situation that is similar to the Tournament: the modal firm is quite similar to the champion firm in the Tournament represented in table 4 (the key difference being in the larger intercept). The level of superfirm profit is quite close to the levels in the evolutionary tournament. The reason for this is that with d=0.1, a lot of firm types survive with populations above 1/n. The evolutionary process does not lead to one type of firm predominating, but rather a wide range. This means that the rules that do well are closer to the "near profit maximizers", and need to be robust, rather than the cooperative firm types that predominated in the simulations with less noise. The outcome of the evolutionary process for d=0.1 is represented in Fig 6:

Fig 6 here

Whilst there is a clear "peak" around the modal firm, there is a long tail of firms which are more competitive, and lead to the lower overall profits.

Initial Position:

All of the simulations in table 5 were generated from an initial vector where all firms had an equal proportion (d/n). There is an obvious question as to how the initial position affects the final point of convergence: is the result typical? In order to investigate this issue we ran 107 simulations for the case of d=0, with each simulation starting from a different initial position. There are various ways of doing this, and we decided to use a deterministic algorithm for generating different initial positions in the 1,176 dimensional unit simplex. Since the initial vector $Z_{0i}=1/n$ is very much in the "middle" of the unit simplex, we decided to use an algorithm that picked extreme initial positions that systematically favoured small groups of firms. The algorithm was the following: starting from the first point in the grid, pick the first 10 firm types; give these a weight of 1/100 each (total 0.1), and the rest of the firms an equal part of the remaining weight (0.9 divided by 1166). The next firm is missed out, and then the next ten are picked and the new initial proportions calculated as before, and so on. This gives 107 initial proportion vectors which we used for simulations: these initial vectors represent a very wide diversity, since they take a highly skewed distribution. Furthermore, firm types from all over the grid are given a chance to start off with a large initial proportion.

All of the simulations were run for at least 250,000 iterations: by this stage in all but 7 out of the 107 simulations the modal firm type had a proportion no less than 0.999. We ran these 7 for 10,000,000 iterations. The results are quite startling:

Iter: 250K	Iter: 10,000K	Modal Firm	Profit
66	66	0.28,0.28	0.1232
28	29	0.26,0.26	0.1248
6	12	0.32,0.32	0.1152
4	0	0.30,0.30	0.1200
2	0	0.36,0.36	0.1008
1	0	0.34,0.36	0.1017

TABLE 7: THE RESULTS OF 107 SIMULATIONS FOR d=0 WITH DIFFERENT INITIAL POSITIONS

Average profit over 250K iterations = 0.1224; over 10,000K = 0.1227.

In the fourth column, we report the modal firm type. The first two columns give the frequencies we observed each modal-firm outcome (as a number out of 107). The 250K iterations column gives the modal proportions after all simulations had been run for 250K iterations: the 10,000K column gives the outcomes after the 7 simulations with modal proportions below 0.999 had been run for 10,000K. The last column gives the average payoff of the surviving firms. As can be seen, the most common outcome was for the firm (0.28,0.28) to emerge: this is only slightly less cooperative than (0.26,0.26), as can be seen from the profit column. In only 3 of the simulations at 250K iterations did we observe a level of profits below the Cournot level of 0.1111. The profits averaged over all 107 simulations after 250K iterations were 0.1224, which is 81.52% on the collusion index. *Whilst JPM* is clearly not a global attractor for the replicator dynamics, the outcomes generated by these simulations indicate that there is a strong tendency towards cooperation.*

Once we allowed the 7 simulations with a modal proportion of less than 0.999 to run to 10,000K, the outcomes became more cooperative. In fact the outcomes remained of the "mixed" variety, as depicted in Table 8:

0.32,0.32	0.34,0.32	028,0.3	0.4,0.3	0.28,0.28	0.36,0.28	Av. Prof
0.55966	0.29898	0.12617	0.01518			0.1152
0.56094	0.30074	0.12582	0.01250	-	-	0.1152
0.57846	0.29577	0.11598	0.00979	-	-	0.1152
0.59440	0.28248	0.11397	0.00915	-	-	0.1152
0.60908	0.27389	0.10891	0.00812			0.1152
0.65048	0.24421	0.09683	0.00848			0.1152
-	-	-	-	0.99767	0.00233	0.1232

TABLE 8: THE 7 CASES WHERE THE MODAL FIRM HAS A PROPORTION LESS THAN 0.999: THE PROPORTIONS AFTER 10,000,000 ITERATIONS.

Of these 7 mixed cases, 6 appear (in terms of payoff) equivalent to a "pure" (0.32,0.32) outcome: the last is obviously almost identical to the pure (0.28,0.28) outcome. We found no evidence of cycles in any of the 107 simulations, and the proportions appeared to be stable, although we would not place much emphasis on these 7 cases (they may be due to rounding errors, for example).

5C: Interpretation of results.

We would argue that there is a very simple interpretation of these results. They show that in the absence of significant noise, there is a strong tendency for the evolution of cooperation in this framework. In Cournot oligopoly, collusive behaviour arises through decision rules that lead to imitation. In an environment in which there is a wide variety of firm types, as represented by the Tournaments, or sufficiently "noisy" replicator dynamics, what is required to survive and thrive is not imitative behaviour, but rather near profit maximizing behaviour, or at least decision rules which accommodate the output of the other firms (as represented by a slope of around -0.5). The replicator dynamics however, lead to the decline of unprofitable rules, and their decline to near zero population shares (if noise is low). This leads to the possibility of more imitative, and hence collusive behaviour to emerge. As such, our simulations lend considerable support to the Axelrod hypothesis that cooperation will evolve in many social settings.

Another point to note is that the modal firm lies on or just off the 45% line in most cases in tables 5 and 7. This is not surprising. By the nature of the algorithm for generating firm types, all the points on the 45% line generate decision rules that are best responses to themselves. If the replicator dynamics results in a modal firm which dominates the population with a large share, then the firm type would have to perform well against itself. It is thus natural to think of the replicator dynamics acting as an *equilibrium selection* mechanism: picking out one of the possible symmetric equilibria along the 45% line. There are of course also "mixed" equilibria: however, the results of our simulations in table 7 suggest that these are "selected" only rarely.

As we move down "mode" column in Figure 5, we gradually move upward along the 45% line in Figure 1 from the collusive towards the Cournot outcome (which is produced when the firm generated by point (1/3,1/3)- alias the Cournot Sticker- plays itself). However, when we move to d=0.1, the situation has changed: the modal firm type has a share of only 2.63%, and the outcome is very much more competitive. Thus the role of noise is to change the best decision rule from a collusive and imitative decision rule towards a near profit-maximizing decision rule, as found in the Tournament result.

Conclusion.

In this paper we have applied the evolutionary approach to the most commonly used oligopoly model: Cournot has met Axelrod (or possibly Darwin). We have been able to reach some quite strong conclusions, and in particular have found that there is a strong tendency towards collusion. We can obviously generalise the approach in two directions. First, staying within the confines of Cournot oligopoly, we can introduce different cost conditions and also differentiated products. Secondly, we can generalise the approach to different oligopoly models (Axelrod can meet Bertrand and Hotelling, to mention but a few). This will enable us to see how robust the results we have found here are when we move outside the simplest Cournot environment.

On the broader front, there is the question of how we interpret these results in general terms. Clearly, the evolutionary approach is more of a metaphor than a model (although no less abstract than many economic models). Whilst we have treated the evolutionary process as a generally plausible story of how successful behaviour spreads, we could clearly try to tailor the story to provide a more realistic account of oligopolistic behaviour. Modelling the dynamics of the selection process in more detail would allow us to move beyond the borrowed replicator dynamics. This is clearly an important and urgent topic for current and future research.

Appendix: the calculation and simulation of the Duopoly payoffs.

We have tried different methods of calculating the payoffs of the constituent duopoly games. In the end, we believe that the exact method is not crucial: the obvious methods tend to give similar results for stable equilibria, and usually differ only in the case of unstable equilibria and unit roots. Since unstable games generally give rise to very low payoffs for one or both firms, it makes little difference how they are calculated (since we are interested in the ATP of the best). Unit roots are of their nature exceptional given our algorithm. We decided to use the eigenvalues to diagnose the properties of the system: first, it is neater and more precise; second, it is computationally less demanding. We outline the method used: the precise programme can be obtained from Steven Wallis.

If we ignore the non-linearity created by the non-negativity of output constraint, the stability of (4) can be diagnosed by the eigenvalues $\lambda_i = \pm(h_{1i}.h_{1j})^{1/2}$. Let us define the point of intersection of the two decision rules in $\mathbb{R}^2$ as $\mathbf{x}^*$. We consider 4 cases:

Case 1: Stability $(h_{1i}.h_{1j})^2<1$.
In this case, the roots are within the unit circle, and the system is stable. The equilibrium payoff is calculated as the payoff where the two decision rules intersect (allowing for the non-negativity constraint on output where this binds).

Case 2: Instability $h_{1i}.h_{1j}>1$. Here initial output matters. If the output of either firm is above the point of intersection $\mathbf{x}^*$, then the total output goes to infinity, and profits are zero. If the output of both firms is below the point of intersection, then output of both firms falls. Due to the non-negativity constraint, the result is that: x_i = max $[h_{oi},0]$ or x_j = max $[h_{oj},0]$. One of these will always be zero. A positive payoff can only occur when the initial outputs are both below the intersection point. So, we simply assume that the initial outputs are uniform on the unit triangle, and multiply the above payoff by the probability that the outputs are below their intersection values (i.e. the product $\min(X^*_i,1).\min(X^*_j,1)$).

Case 3: Instability $h_{1i}.h_{1j}<-1$. Here, because of the non-negativity constraint, the system converges on a 4 cycle, in which one of the firms has a zero in two consecutive periods (without the non-negativity constraint, the system would explode). We computed the profits over the 4 cycle.

Case 4: Positive unit root. $h_{1i}.h_{1j}=1$. This outcome is very unlikely, and the outcome depends on initial outputs. We took a range of initial outputs and averaged over the different outcomes (note that in many cases the outputs will both go to infinity, and hence profits are zero).

Case 5: Negative unit root $h_{1i}.h_{1j}=-1$. Again, this is very unlikely. In this case there are 4-cycles: we start from a range of initial positions, and average.

With this algorithm, we only need to simulate the system in case 3 for , and the rarely encountered cases 4,5.

Bibliography.

Abreu D, Rubenstein A. (1988): The structure of Nash equilibria in repeated games with finite automata, *Econometrica*, 56, 1259-1282.

Alchian, A.A. (1950): Uncertainty, evolution and economic theory, *American Economic Review*, 58, 211-222.

Arthur, B.(1993): On designing economic agents that behave like human agents, *Journal of Evoluti0nary Economics*, 3, 1-22.

Axelrod R. (1984): *The Evolution of Cooperation*, Penguin, London.

Binmore K. (1990): *Essays on the Foundation of Game Theory*. Blackwell, O

Binmore K. (1992): *Fun and games: a text on game theory*. D.C.Heath

Binmore K, Samuelson L (1992): Evolutionary stability in repeated games played by finite automata, *Journal of Economic Theory*, 57, 278-305.

Binmore K, Samuelson L. (1993a): An economists perspective on the evolution of norms. Mimeo, U.C.L.

Binmore K, Samuelson L. (1993b): Complexity and noise in evolutionary equilibrium selection, mimeo U.C.L.

Binmore K, Gale J, Samuelson L.(1994): Learning to be Imperfect: the Ultimatum Game. Mimeo, U.C.L.

Binmore K, Samuelson L, Vaughan R. (1994): Musical Chairs: Modelling Noisy Evolution. Mimeo U.C.L.

Bone J, Dixon, H.D.(1992): Oligopoly outcomes as Nash-equilibria in decisionrules: some further results. Mimeo, York University.

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

Chattoe, E (1994): The use of evolutionary algorithms in economics: metaphors or models of social interaction? Mimeo, Oxford Institute of Economics.

Cyert R, March J (1963): *A behavioural theory of the firm*, Englewood Cliffs, NJ.

Friedman J.W.(1977): *Oligopoly and the theory of games*, North-Holland.

Hall R, Hitch C (1939): Price Theory and Business Behaviour, *Oxford Economic Papers*, 2, 12-45.

Hart, O.D.(1985): Imperfect competition in general equilibrium: an overview, *Frontiers of Economics* ed. Arrow K and Honkaphoja S. Blackwells, Oxford.

Hey J, Martina R (1988): Reactions to reactions and conjectures about conjectures, *Scottish Journal of Political Economy*, 35, 283-290.

Klemperer P, Meyer M (1989): Supply function equilibria in oligopoly under uncertainty, *Econometrica*, 57, 1243-1277.

Leonard R.J.(1994): Reading Cournot, Reading Nash: the creation and stabilization of the Nash equilibrium, *Economic Journal*, 104, 492-511.

Linster B.(1990): *Essays in Competition and Cooperation*, PHD Thesis, University of Michigan.

Linster, B.G.(1992): Evolutionary stability in the infinitely repeated PD game played by 2 state machines, *Southern Economic Journal*, 58, 880-903.

Linster, B.G.(1994): Stochastic evolutionary dynamics in the repeated Prisoners Dilemma, *Economic Enquiry*, 32, 342-357.

Marks, R.E.(1992): Breeding hybrid strategies: optimal behaviour for Oligopolists, *Journal of Evolutionary Economics*, 2, 17-38.

Maynard-Smith J (1982): *Evolution and the theory of games*, C.U.P, Cambridge.

Maynard-Smith J, Price G.(1973): The logic of animal conflict, *Nature*, 246, 15-18.

Moss S, Dixon H, Wallis S. (1993): Evaluating Competitive Strategies, mimeo, Manchester Metropolitan University.

Nachbar, J(1992): Evolution in the finitely repeated PD game, *Journal of Economic Behaviour and Organisation*, 19, 307-326.

Nelson R, Winter S (1982): *An evolutionary theory of economic change*, H.U.P., Cambridge mass.

Peyton-Young H. (1993): The Evolution of Conventions, *Econometrica*, 61, 57-84.

Probst D. (1992): Evolution, automata and the repeated Prisoner's Dilemma, mimeo, L.S.E.

Robson, A (1986): The existence of Nash-equilibria in reaction functions for dynamic models of oligopoly, *International Economic Review*, 27, 539-544.

Selten R. (1991): Evolution, learning and economic behaviour, *Games and*

Economic Behaviour, 3, 3-24.

Simon, H.A, (1947): *Administrative behaviour: a study of decision-making processes in administrative organisations,* MacMillan.

Simon H.A. (1981): *The Science of the Artificial*, M.I.T. press.

Stanford W. (1986): Subgame perfect reaction functions in discounted duopoly games are trivial, *Journal of Economic theory*, 39, 226-232.

Stanford W. (1986): On continuous reaction function equilibria in duopoly supergames with mean payoffs, *Journal of Economic theory*, 39, 233-250.

Van-Damme E (1987): *Stability and Perfection of Nash Equilibria*, (2e) Springer-Verlag.

Van-Damme E, .(1994): Evolutionary Game Theory, *European Economic Review*, 38, 847-858.

Weibull J (1994): The "as if" approach to game theory: three positive results and four obstacles, *European Economic Review*, 38, 868-881.

Winter, S.(1971): Satisficing, selection and the innovating remnant, *Quarterly Journal of Economics*, 85, 237-261.

EVOLUTION RESULTS

d	Mean	Mode	Int.	Slope	Prop.	Av.Pr.	SF	Iterations
0		0.26,0.26	0.04	0.846	1	0.1248	0.1248	58,500
0.00001	0.2611,0.2596	0.26,0.26	0.04	0.846	0.9956	0.1248	0.1250	>2,500,000
0.0001	0.284,0.278	0.28,0.28	0.12	0.571	0.9814	0.1232	0.1237	216,000
0.001	see below					0.1177a		<2,000,000
0.01	0.358,0.304	0.3,0.32	0.24	0.24	0.5042	0.1099	0.115	52,900
0.1	0.434,0.334	0.38,0.4	0.56	-0.421	0.0263	0.0744	0.0897	2,594

TABLE 5: EVOLUTIONARY RESULTS FOR DIFFERENT VALUES OF NOISE d.

Int. - intercept of modal firm.
slope. - Slope of modal firm.
Av. Prof. - Average profit of modal firm.
SF - Superfirm profits when playing against all survivors (all firms when d>0.
Iterations- Number of evolutionary iterations performed (as a guide).
a - average profits over cyle (for d=0.001)
All simulations started from the initial position Z_{i0}=1/n, i=1..1,176.