Computational Economics **20:** 139–156, 2002.

Axelrod Meets Cournot: Oligopoly and the Evolutionary Metaphor

HUW DAVID DIXON
Economics Department, York University, York Y01 5DD U.K.; E-mail: hdd1@york.ac.uk

STEVEN WALLIS and SCOTT MOSS
Centre for Policy Modelling, Faculty of Management and Business, Manchester Metropolitan University, Aytoun Building, Aytoun Street, Manchester M1 3GH, U.K.;
E-mail: 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.

Key words: evolution, Cournot, duopoly

It has long been argued that firms use rules of thumb for solving problems and making decisions,[1] and also that firm's rules of thumb might evolve over time in a Darvinian manner.[2] More recently game-theorists have made the connection between conventions, bounded rationality, learning and evolutionary models,[3] and the idea has been popularized in the social sciences by Axelrod (1984). The central notion is very simple. Firms follow 'strategies', or rules which tell them what to do: different 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*.[4] 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. The question which this paper seeks to address is what behavior will result form this process: competition or collusion?

This paper seeks to explore the evolutionary approach to explain the behavior of firms (decision rules) in the simplest possible Cournot environment using a simulation methodology. 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. Furthermore, we can employ the simulation based 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 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.[5] [6]

What types of behavior would we expect to find duopolists following in a Cournot environment? If we restrict firms' decision rules to the simplest case of constant outputs, then the Nash equilibrium will be selected under a wide range of modeling assumptions.[7] However, we consider the case of *linear decision rules*, where multiple Nash-equilibria exist. Hence the question is not whether a unique Nash equilibrium will be selected, but rather which equilibria out of a large number of Nash equilibria will be tend to be selected by the evolutionary dynamics? As such, the paper has more in common with evolutionary models of coordination games (Vega-Redondo (1993), Oechssler (1997, 1999)).

The conclusions of the paper are really quite sample and quite far reaching. Our first main conclusion relates to the type of firm which does best in the 'Tournament', where there are many different types of firms, some quite unconventional. In the Tournament, each strategy is played against all strategies. We find that the 'Champion' firm in this context which earns the highest profit after playing all strategies is a 'near profit maximizer':[8] indeed, the standard Cournot best-response function does very well. *Near profit maximizing behavior is a robust decision rule in an environment where a firm might meet a wide variety of firm types*. The intuition here is that if you can only choose one decision rule with which to play *all* possible strategies, then a strategy that ensures that you always play a near best response output to your opponent's output will do well overall. Of course, when playing a specific decision rule, this is not so: for example, if you are playing a standard best-response Cournot reaction function, then your best strategy is a decision rule that gives you the Stackelburg leader's output, which is not a best response to the follower's output. However, this result indicates that there is no single rule which strategically exploits all other decision rules: the best it appears you can do is to near-profit maximize.

Our second main conclusion concerns the evolution of cooperative behavior. If we allow a process of evolution to take its course, decision rules that yield low profits are weeded out. When this happens, 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 behavior towards cooperation*. The intuition follows from the fact that the collusive decision rule is one that imitates its competitor: when it meets itself it maximizes joint-profits.[9] In an economy with diverse decision rules, the collusive decision rule will do badly overall (since it copies poor decision rules). As evolution drives out poor decision rules, however, the imitative nature of the collusive decision rule means that it must do at least as well as its competitor and when it meets itself it earns the highest joint profit. Hence if enough collusive firm types prevail in an economy, they will tend to take over the population. Collusion pays, and in an evolutionary environment this means it can posper. Unlike the standard Prisoners dilemma, collusion is robust against invasion here: the imitative behavior of the collusive decision rule means that it will do as well against the invading strategy as the invader does against itself.[10]

This result can be set in the context of other papers. Where we have multiple Pareto ranked Nash equilibria, we can expect the Pareto dominant equilibrium to be selected more frequently, in the sense of it having a larger basin of attraction in some sense (although Pareto-inefficient equilibria cannot be ruled out). Experimental evidence certainly supports the ability of subjects to coordinate on payoff dominant outcomes much of the time: see for example Bacharach and Bernasconi (1997), Mehta et al. (1994a, b).[11]

However, as the replicator dynamics become noisier, we find that the evolutionary simulations move away from the collusive and towards the near profit maximizing behavior that performed best in the tournament. Evolution towards collusion requires that the population is able to evolve towards more cooperative decision rules. However, with more noise in the evolutionary process, there remains a diversity in decision rules which limits the extent to which it is possible for cooperative types to predominate.

In Section 2, we outline the basic Cournot duopoly model used' in Section 2 we describe the method employed to generate the decision rules and the resultant Tournament. In Section 3 we outline the evolutionary dynamics both with and without and present the simulation results.

1. The Cournot Model

There are no costs, and the market price P is a linear function of the two outputs x_i with a unit slope and intercept:

$$P = \max[0, 1 - x_1 - x_2] \tag{1}$$

The firms' profits are given by the payoff function $U_i : A \to [0, 1/2]$:

$$U_i(\boldsymbol{x}) = x_1 \cdot (1 - x_i - x_j) \tag{2}$$

and where $A \equiv \{\boldsymbol{x} \in [0, 1]^2 : 1 - x_i - x_j \geq 0\}$, the unit triangle.

Firms play the duopoly game using a *decision* rule for choosing output. Decision rules are linear, and give the output of firm i in period t as a function of the output of the other firm j in the previous period. A *Firm type* is defined by a decision rule, and can be represented as a pair of parameters $\{h_{0i}, h_{1i}\} \in \mathbb{R}^2$, the *intercept* h_o and *slope* h_1. We consider various standard decision rules in Table I.

The Myopic Cournot Profit Maximizer (MCPM) is the standard best-response function. The Stackelburg Sticker (SS) produces the Stackelburg Leader output 1/2 every period.[12] The Cournot Sticker (CS) produces the Cournot-Nash output 1/3

Table I. Some standard decision rules.

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 maximizer/copy cat	0	1

every period. The 'Copy cat' 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. We call this the Joint Profit Maximizer since it supports the joint profit maximum.[13]

2. The Method

Our method consists of three stages. First 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. Thirdly, we then run an evolutionary algorithm based on replicator dynamics. Let us now consider in detail these three steps in turn.

The way we have chosen to generate firm types is based on a result found in Hey and Martina (1988) and Klemperer and Meyer (1998).[14] If we take any point in the interior of the unit triangle A, we can generate a reaction function by taking the tangent to firm 1's isoprofit curve at that point. Let us consider some point $\boldsymbol{x}' \in \text{intA}$, at $\boldsymbol{x}'$ the payoff of firm 1 is $x_1(-x_1 - x_2)$. It is easily verified that the tangent to the isoprofit curve at $\boldsymbol{x}'$ is the ratio of the marginal profits $\partial U_1/\partial x_i$ at $\boldsymbol{x}'$ and is characterized by the slope and intercept terms:

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

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

Note that if firm 2 chose the decision rule thus generated, then the best profit that firm 1 could achieve is by choosing a decision rule passing through $\boldsymbol{x}' : \boldsymbol{x}'$ solves the problem $\max U_1(\boldsymbol{x})$ subject to $x_2 = h_0 + h_1 x_1$, where h_0 and h_1 are as in (3). A Nash-equilibrium in decision rules occurs at $\boldsymbol{x}'$ when firm 1 chooses as its reaction function the tangent to firm 2's isoprofit curve at $\boldsymbol{x}'$, and vice-versa. Such an equilibrium can be constructed to support any point in the interior of the unit triangle (Klemperer and Meyer, 1988). Our method of generating firm types 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 (the mapping represented by (3a, b) is 1-1).[15]

The algorithm for generating the firm types is implemented using a grid search on the unit triangle A. We specify the *granularity* of the grid, the distance between two adjacent points, and generate firm types from points in the interior of the unit triangle. The set of decision rules that are best responses to themselves are to be found on the 45°line ($x_1 = x_2$). 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, insofar as 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 strictly dominated.

Given the set of firm types, we then calculate payoff matrix by running an Axelrod Tournament:[16] each firm type meets each other firm type (including itself) to play the *constituent* duopoly game. In the *constituent* duopoly game the pair of decision rules $\{i, j\}$ defines the dynamic system:

$$\boldsymbol{x}_t = \boldsymbol{h} + \boldsymbol{H} \cdot \boldsymbol{x}_{t-1}$$

where $\boldsymbol{h} = [h_0 i, h_{0j}]$, and $\boldsymbol{H} = \begin{bmatrix} 0 & h_{1i} \\ h_{1j} & 0 \end{bmatrix}$. We have tried both *simulation* and *analysis* to evaluate the firms' profitability in the constituent game. We chose the analytical method: using the eigenvalues of $\boldsymbol{H}$ to classifty the dynamic properties of the system, and simulating only where unavoidable and sensible (for about 5% of constituent games).[17] Where the system was stable, we used the 'stationary' per-period payoff to measure profit;[18] where the system was unstable, we took an analogous measure. The exact algorithm used is described in the appendix.

Having run the full Axelrod Tournament with n types of firm, we then have the $n \times n$ payoff matrix $\boldsymbol{T} = [u_{ij}] : u_{ij}$ being the payoff of type i when it meets type j. The *Average Tournament Profit for firm i* (ATPi) is the arithmetic average of firm i's payoffs summed over $j = 1 \ldots n$:

$$\mathrm{ATP_i} \equiv (1/\mathrm{n}) \cdot \left[\sum_{\mathrm{j=1\ldots n}} \mathrm{u_{ij}} \right] . \tag{4}$$

The Tournament ATP is the arithmetic average of the ATP$_i$. 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. Since we generate type i from the tangent to firm 1's iso-profit curve at $\boldsymbol{x}_i$, it follows that if firm 1 were faced with this decision rule, its highest profits will be its profits at $\boldsymbol{x}_i$, so that superfirm profits are:

$$\left(\frac{1}{n}\right) \sum_{i=1}^{n} x_{1i}(1 - x_{1i} - x_{2i}) .$$

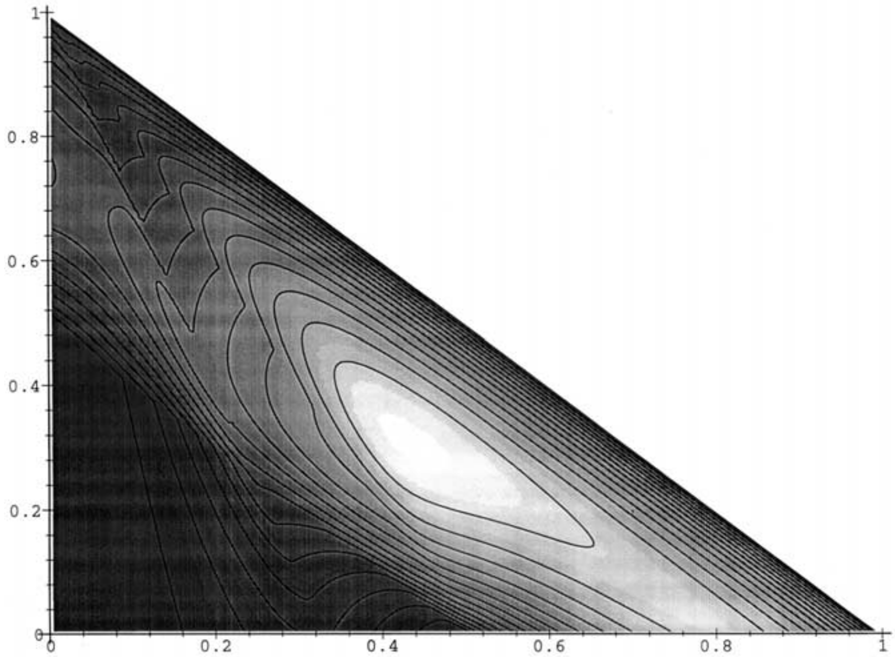

Figure 1. Payoffs in the large Axelrod tournament.

The ratio of actual so superfirm profits represents the cost to the firm type of not having the flexibility to tailor its decision rule to each different opponent type, and is hence an indicator of the cost of Bounded Rationality, or the dual notion of the value of flexibility.

We report the largest Tournament we have run: the grid search had granularity 0.005, generating a total of 19,702 firm types. All our computations were in douible digit precision, although we only report values to 4 s.f. unless stated otherwise. We also report a smaller Tournament which we use in the next section for the evolutionary algorithm: here the granularity was 0.02, generating 1,176 firm types.

In Figure 1, we have a contour map for the large Tournament: each point on the map represents a firm type, and the height of the contour it is on represents its Average Tournament Profit (ATP). 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 -0.414 and intercept 0.49. The ATP of the Champ is 0.07466, (which represents 88% of superfirm profits, and almost twice the average ATP). In Tables IIa, b we present the performance of the Champ's in the two Tournaments in comparison to the firms in Table I (or the closest firms to them in the grid, marked with an *).

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

Table IIa. The large Axelrod tournament.

	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
Stackelburg sticker	0.25	0.50	0.50	0	0.05203

Granularity 0.005; 19,702 firm types; superfirm profits 0.08375 (4 s.f.). Average ATP 0.3775.

earn low profits). The most striking feature of Tables II is to note that the MCPM does so well, being quite 'close' to the Champs in terms both of payoff and the intercept/slope coefficients (the profits of the MCPM are over 98% of the Champs). The other reference firms from Table I do badly: the JPM rule is even below the average ATP.

The fact that the MCPM does so well may seem at first sight something of a puzzle. However, the season why the Champ and the MCPM are so close is that the Champs 'near profit-maximizing' behavior is very robust. In the population, there is a great diversity of firm types: the Champ has to perform well across the board, no matter who it meets. Near profit maximizing behavior ensures at least that its output in each of the constituent games is almost a best response to the output of the other firm. Whilst other firm types might do very well against specific other types, they do not do well on average against all types. For example, the copy-cat behavior of JPM is not very good in this environment, since it can only prosper with rules that do well against themselves. Thus when JPM meets the Stackelburg Sticker, both firms produce output of 0.5 and earn nothing; MCPM produces 0.25 and may earn less than the Stackelburg Sticker, but more than JPMs nothing. The Stackelburg Sticker is designed to optimize against only one other strategy: MCPM. When it plays itself or JPM, it earns zero. MCPM is more robust: when it meets another strategy, it is able to earn strictly positive profits so long as the other firm produces less than 1 unit of output in equilibrium. It is the ability to do reasonably well against all decision rules which means that it does so well in the Tournament. *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.*

3. Evolution in the Duopoly Model

There has been much discussion of the evolutionary metaphor in economics in recent years.[19] In biology, successful species or genes tend to become more common because they give rise to more progeny. In the context of *social* evolution, mech-

Table IIb. The tournament used for the evolutionary dynamics.

	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

Granularity 0.02; 1,176 firm types; superfirm profits 0.08500. Average ATP 0.03905.

Table III. NSS and ESS firm types.

Type	0.26	0.28	0.3	0.32	0.34*	0.36	0.38
U_{ii}	0.1248	0.1232	0.12	0.1152	0.1088	0.1008	0.0912

ESS firm type marked with *.

anisms 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*:[20] firms tend to imitate the more successful practices of other firms, as in ‘benchmarking’. There is also *learning*: firms will receive signals from the capital market and elsewhere about how they are performing relative to other firms: this will lead less successful firm types to adapt their behavior. 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.

Having run the Tournament, we proceeded to apply an evolutionary algorithm. The vector $\boldsymbol{Z}$ gives the proportions Z_i of each firm type i ($Z_i \in [0, 1]$ and $\Sigma Z_i = 1$). We start from an initial condition $\boldsymbol{Z}_0$ and then represent evolution using the *replicator* dynamics.

On a purely theoretical level, we can identify which strategies are *Evolutionary Stable* (ESS) and which *Neutral stable* (NSS).[21] Given the payoff matrix $\boldsymbol{T}$, it is easy to calculate these types.[22] A strategy i is NSS iff: $u_{ii} \geq u_{ki}$ for all $k = 1 \ldots n$, and if there exists some $k \neq i$ s.t. $u_{ki} = u_{ii}$, then $u_{kk} \leq u_{ik}$. A strategy i is ESS iff it is NSS and the second inequality is stict. In Table III, we list the 7 NSS strategies, one of which (the near Cournot Sticker) is ESS and is marked with an asterisk. In the top row, the firm type is identified by one ordinate (since all NSS are on 45°line), and the bottom row gives the payoff.

The *set* of all 7 NSS strategies itself is not an Evolutionary Stable (ES) set as defined by Thomas (1985).[23] The set of NSS types is Pareto-ranked: the payoffs descend from left to right, with the JPM (type 0.26) the Pareto dominant NSS strategy. Note that there are not more NSS strategies: decision rules generated for outputs below 0.26 (0.24 to 0.02) have positive slopes in excess of 1 and are unstable against themselves. Those generated by outputs above 0.4 are too competitive: although they are stable against themselves, they can be beaten by firms with large slopes which generate a cycle and over the cycle earn higher profits.

In this paper, we have modeled the evolutionary mechanism using the discrete time replicator dynamics: we chose this merely as a representative of the class of payoff monotone dynamics whose properties are well known. The basic replicator dynamics used were (following van Damme (1987) and Gale et al. (1995)):

$$Z_{is} = Z_{is-1}(1-d) + Z_{is-1}\left[(U_{is-1} - u_{s-1}\right] + (d/n) \quad (5)$$

where $d \in [0, 1)$ is a 'noise' parameter, U_{is} is the average payoff of firm i in iteration s ($U_{is} = \Sigma_{j=1\ldots n} Z_{js} \cdot u_{ij}$), and u_s is the average payoff of all firms in iteration s ($u_s = \Sigma_{j=1\ldots n} Z_{js} \cdot U_{is}$). For computational reasons, we allowed for extinction when Z_i became very small.[24]

Without noise, the change in proportions depends on the absolute difference in profits from the population average. 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).[25] Assuming that there exists an attractor(s) the replicator dynamics converge to a situation where all surviving firms have equal average profits: $\boldsymbol{Z}^*$ is such that $U_i = U_j$ for all $i, j = 1 \ldots n$ when $\boldsymbol{Z}_{i,j} > 0$. We can interpret firm types as 'pure strategies' in a game: if the replicator dyanmics converge, then $\boldsymbol{Z}^*$ is a Nash-equilibrium given the set of strategies that were present with non-zero proportions in the initial vector $\boldsymbol{Z}_o$.[26]

With noise $d > 0$, each firm type i loses $d \cdot Z_i$, and these firms are randomly allocated across all firm types, each firm type gaining d/n. this sort of 'noise' can be interpreted as random 'mutation' (as in Linster (1994, pp. 348–353 and Gale et al. (1995)).

The fact that a strategy is ESS or NSS is neither necessary nor sufficient for it to be an attractor of the replicator (or other) dynamics. However, it does give us some idea what to expect from the simulations: the key question is which amongst these NSS strategies will the replicator dynamics select?

3.1. RESULTS: DIFFERENT NOISE LEVELS

First, we describe the evolutionary simulations we have run for different values of the noise parameter d from the initial vector $Z_{io} = 1/n$, which are summarized in Table IV. We discuss simulations with different initial positions in Table V. The second column of Table IV gives the *mean firm* at the end of the simulation, thus we weight the location of each firm type by its proportion: $\Sigma_{i=1\ldots n} Z_i \cdot \boldsymbol{x}_i$. The third column gives the *modal* firm at the end of the simulation. The next four columns

Table IV. Evolutionary results for different values of noise d.

d	Mean	Mode	Int.	Slope	Prop.	MATP	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					0.1177 [a]		<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

Int. = intercept of modal firm; Slope = slope of modal firm; MATP = 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 cycle (for $d = 0.001$).
All simulations started from the initial proposition $Z_{i0} = 1/n, i = 1, \ldots, 1, 176$.

Table V. 107 simulations for $d = 0$ with different initial positions.

Frequency	Modal firm	Profit
66	(0.28, 0.28)	0.1232
29	(0.26, 0.26)	0.1248
12	(0.32, 0.32)	0.1152

describe the modal firm: its intercept, slope, and its proportion and average profits (MATP) 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 in that the end-of-simulation weights are used. 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.). All simulations converged, except in the case of $d = 0.001$, which generated in regular cycle.

When $d = 0$, we have a very clear result which supports the cooperative hypothesis: the one surviving firm type JPM* is the Pareto Dominant NSS from Table II, JPM* earns about average profits in the Tournament. 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 such firms.

When there is noise, all firm types survive so that the mean and modal firm types differ. The second and third row of Table IV obviously represent very low levels of noise, and the resultant outcome is still very close to the cooperative outcome, as measured by the MATP. As we increase the level of noise, MATP gradually

Table VI. The three most common firms when $d = 0.001$.

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

decreases, until when $d = 0.1$, the MATP is not so far from the ATP of the Champ in Table IV. As d increases, the modal firm moves away from the 'cooperative', and becomes closer to the 'near profit maximizing' firm types. This reflects the fact that as d increases, there is a larger presence of all firm types, bringing the environment nearer to that of the Tournament.[27] Only in the case of $d = 0.001$ did we find that the replicator dynamics did not converge to a static population proportions. Instead, we found a cycle emerged with 3 firm types being most common. These three firms are given in Table VI, with the last column giving their maximum proportion over the cycle, which lasts about 12,000 iterations.

The average profits of all three firms are almost the same, but small differences generate a long cycle: profits of the three firms fluctuate over a range 0.1161 to 0.1187, averaging 0.1177 over the cycle.

3.2. RESULTS: NO NOISE, DIFFERENT INITIAL POSITIONS

In order to investigate the importance of the initial position, we ran 107 simulations for the case of $d = 0$, with each simulation starting from an initial position chosen by an algorithm. Since the initial vector $Z_{0i} = 1/n$ is in the middle of the unit simplex, we decided to use an algorithm that picked extreme initial positions favoring small clusters of types.[28] All of the simulations were run for at least 250,000 iterations and until the modal firm was stable with a share of at least 0.999, being reported in Table V. The three NSS types with the highest profits are the attractors here, not the ESS type.[29] *Whilst JPM* is clearly not a global attractor for the noiseless replicator dynamics, the outcomes generated by these simulations indicate that there is a strong tendency towards cooperation*.

One possible explanation of this tendency is provided if we note that ignoring stability problems the pure JPM weakly dominates all other strategies *in pairwise contests*: if JPM plays itself, it earns 0.125: if it plays type j, it earns u_{jj}. This is similar to the 'secret handshake' story employed by Binmore and Samuelson (1994) to explain the possibility of cooperation in the PD. There are two main differences, however. First, our structure has a rich population of strategies, and none is dominant and possibly none are dominated. Second, in the PD cooperation is not a Nash equilibrium: in our model it is Nash and NSS. Hence we are observing

Table VII. Simulations for $d = 0$ with different initial positions and only the 7 NSS types.

Type	0.26	0.28	0.30	0.32	0.34	0.36	0.38
Number	1218	90	40	21	12	7	5
(%)	(87.4)	(6.5)	(2.9)	(1.5)	(0.9)	(0.5)	(0.4)

the selection of the Payoff dominant Nash equilibrium most of the time, not the selection of a non-Nash outcome.

Lastly, with no noise, we looked at the case where only the NSS firm types were present.[30] The resultant 7×7 payoff matrix T_7 is included in the Appendix, Table A.I. All of the pure strategies are ESS within this restricted set: whilst there exist mixed Nash equilibria, none of these are NSS (see proposition in appendix). We ran 1393 simulations for different starting points (the algorithm is described in the appendix), and the results are shown in Table VII (all simulations converged to one of the types having a share of 1).

Clearly, when we restrict ourselves to the 7 NSS strategies, the tendency to select JPM* is even stronger than when all 1,176 strategies are present, since we have excluded the rules whose imitation leads to low payoffs.

4. 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 when the evolutionary dynamics are not too noisy. The evolution of collusion in the setting of the replicator dynamics requires that evolutionary selection is able to weed out the non-cooperative types of firm before evolution can take hold. The presence of a sufficient level of noise means that this is unable to occur. When there is a high level of noise, the myopic Cournot profit maximizer does very well, and is shown to be a robust decision rule that does reasonably well against all possible firm types it meets.

We can obviously generalize 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 generalize 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.

Appendix

(A) THE CALCULATION AND SUMULATION OF THE DUOPOLY PAYOFFS

We outline the method used- the precise programme can be obtained from Steven Wallis. We ignore the non-linearity created by the non-negativity of output constraint, so the stability of $\boldsymbol{H}$ can be diagnosed by the eigenvalues $\lambda_i = \pm(h_{1i} \cdot h_{1j})^{1/2}$. Let us define the point of intersection of the two decision rules in $\mathbb{R}^2$ as $\boldsymbol{x}^*$. We consider 4 cases:

Case 1: Stability $(h_{1i} \cdot h_{1j}) < 1$. 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 $h_{1i} \cdot h_{1j} > 1$. Here initial output matters. If the output of either firm is above the point of intersection $\boldsymbol{x}^*$, 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 $x_j = \max[h_{oj}, 0]$, where 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 $\min(x_i^*, 1) \cdot \min(x_j^*, 1)$).

Case 3: Instability $h_{1i} \cdot h_{1j} < -1$. 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: $h_{1i} \cdot h_{1j} = 1$. 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: $h_{1i} \cdot h_{1j} = -1$. 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 cases 3, 4, 5.

4.1. THE 7 × 7 CASE

Search Algorithm with 7 NSS types.

The following algorithm was felt to allow for each type to start from a range of initial positions. Consider the unitary vector for each type i, σ_i which has a 1 in the ith row, and zeros elsewhere, and σ_A which has each firm with 1/7. For each type, we take the set of initial positions formed by the convex combinations $\lambda\sigma_i 1 + (1-\lambda)\sigma_A$, where λ takes values on a grid on [0,1]. The set of initial positions is the union of the sets generated by each type $i = 1 \ldots 7$. This set is symmetric across firm types, and uniform in terms of the distance from the centre σ_A. The

Table A.I. The payoff Matrix T_7. All figures to 4 d.p.

	0.26	0.28	0.30	0.32	0.34	0.36	0.38
0.26	0.1248	0.1231	0.1199	0.1151	0.1007	0.0911	
0.28	0.1243	0.1232	0.1199	0.1148	0.1082	0.1001	0.0904
0.30	0.1223	0.1228	0.1200	0.1150	0.1082	0.0997	0.0898
0.32	0.1184	0.1210	0.1196	0.1152	0.1085	0.0999	0.0896
0.34	0.1128	0.1173	0.1177	0.1147	0.1088	0.1004	0.0900
0.36	0.1053	0.1114	0.1138	0.1127	0.1082	0.1008	0.0908
0.38	0.0959	0.1031	0.1073	0.1083	0.1059	0.1001	0.0912

only parameter is the size of the grid. For Table VI, we used granularity of 0.005, and excluded the extreme values {0.1}, yielding 199 initial positions for each type, 1393 overall).

The 7×7 game restricted to the 7 NSS pure strategies.

In the matrix T_7 above, each strategy (firm type) is represented by the point generating it.

There exist mixed strategy Nash equilibria. For example, consider the two pure strategies corresponding to points 0.26 and 0.28: the 2×2 payoff matrix is:

$$\begin{bmatrix} (0.1248,\ 0.1248) & (0.1231,\ 0.1243) \\ (0.1243,\ 0.1231) & (0.1232,\ 0.1232) \end{bmatrix},$$

where the first row/column is 0.26 (JPM*) and the second 0.28. There are clearly 2 strict Nash equilibira: the symmetric pure strategy equilibria. There is also a mixed equilibrium with the probability 0.1667 of playing 0.26, and 0.8333 of playing 0.28. Many such mixed equilibria exist in the full 7×7 game. However, it is easy to show that only pure strategy equilibria can be NSS.

This proof relies on the structure of the payoff matrix. Order the strategies using u_{ii} as the criterion (as depicted in Table A.I). For each i, we can define the $(7 - i)$ submatrix consisting of all of the $7 - i$ strategies j for which $u_{jj} \leq u_{ii}$. This submatrix has the property that u_{ii} is the strictly largest payoff ($u_{ii} > u_{kj}$ for all $k, j = i \ldots 7$ except $k = j = i$). As a preliminary, let us define mixed strategy payoff-function $E\Pi(Z, Y)$, which gives the payoff to an agent playing Z against Y, where $Z, Y \in \Delta$ (Weibull (1995, pp. 8–9).

PROPOSITION. *Let Z be Mixed-strategy Nash equilibrium. If there are at least $2i \in S$, such that $z_i > 0$, the Z is not NSS.*

Proof. We consider mixed strategies $Z, Y \in \Delta$. It is fundamental property of a mixed equilibrium that all pure strategies in the support earn the same payoff.

Define Z_i to be the pure strategy where $z_i = 1$ and $z_j = 0$ $j \neg i$. The payoff of each strategy i in the support of Z is:

$$E\Pi(Z_i, Z) = E\Pi(Z, Z) = \sum_{j=1}^{7} z_j u_{ij} .$$

Now, consider the pure strategy j with $z_j > 0$, and for which $u_{jj} > u_{ii}$ for all i such that $z_i > 0$. Since, $E\Pi(Z, Z) = E\Pi(\sigma_j, Z)$, Z is NSS only if $E\Pi(\sigma_j, \sigma_j) = u_{jj} \leq E\Pi(Z, \sigma_j)$. Since $u_{jj} > u_{ii}$, it follows that $u_{jj} > u_{ij}$ for all i. Since $E\Pi(Z, \sigma_j)$ is simply a convex combination of u_{ij}, $u_{jj} > E\Pi(Z, \sigma_j)$, and Z is not NSS.

Acknowledgements

The original version of this paper is Dixon et al. (1995). We would like to thank Ken Binmore, David Collard, Martin Cripps, Bipasa Datta, Helen Robinson, Larry Samuelson, Avner Shaked, Ernesto Somma, seminar participants at UCL, Warwick, Manchester, Leicester, the 1995 Econometric Society World Congress (Tokyo), the Industrial Economics Study Group meeting at York University in April 1994 for comments on this and earlier drafts. We would also like to thank the Nuffield foundation and ESRC for their support, and the helpful comments of the referee. Despite all of this help, the remaining faults are surely our responsibility.

Notes

[1] e.g., Hall and Hitch (1939), Cyert and March (1963), Simon (1947).

[2] Alchian (1950), Nelson and Winter (1982), Winter (1947).

[3] e.g., Binmore and Samuelson (1992, 1993), Canning (1992), Kandori Mailath and Rob (1993), Young (1993), Selten (1991), Oechssler (1997,9).

[4] Gale and Rosenthal (1999) for an analysis of imitation and experimentation with boundedly rational agents in the context of stochastic stability analysis.

[5] The main related literature on evolutionary models since Axelrod has been in modeling the prisoner's dilemma using the notion of finite automata. The idea here (following Abreu and Rubinstein (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), Miller (1993)). 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.

[6] If a game possesses a strictly dominant strategy, then one need form no hypothesis about the other player's behaviour or beliefs to know that the dominant strategy is the best one to choose, and the decision is essentially non-strategic. Whilst few applications will possess this property, many are 'dominance solvable'. However, *iterated deletion* requires some suitable notion of common knowledge and rationality, being inherently strategic reasoning.

[7] Weibull (1995).

[8] '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.

[9] In our model, joint profits are maximized with equal profits per firm.

[10] As we show below, Table III, the collusive decision rule JPM is an NSS strategy. Collusion is a Nash-equilibrium here: the choice between collusion and competition a choice between equilibria as in the coordination game, not between equilibrium and non-equilibrium as in the prisoner's dilemma.

[11] Whilst the experimental evidence has been subject to debate (Crawford (1991), Van Huyk et al. (1990, 1997) and the debates summarized in Kagel and Roth (1995, chapter 3)), this has been about the *extent* of coordination failure. It is clearly incorrect to say that people will *always* focus on the payoff dominant outcome as some theories predict (e.g., Harsanyi and Selten (1998)).

[12] Since the Cournot reaction function is linear here, the Stackelburg output is equal to the monopoly output.

[13] There is a technical problem that the output is indeterminate when JPM meets itself. However, there exist decision rules arbitrarily close to the JPM which yield the symmetric joint profit maximum.

[14] We have used the same method in related papers, Moss et al. (1995) and Bone and Dixon (1999).

[15] We have experimented with different algorithms for generating firm types. So long as the algorithm generates a rich population of decision rules, much the same results hold. However, if an algorithm generates a restricted set of decision rules (for example, only rules with negative slope coefficients) then the ouitcome of the Tournament and the evolutionary simulation can be very different. A 'rich' population is one that includes the standard types (at least slopes between -1 and $+2$).

[16] Clearly, our Tournament differs from Axelrod's, in that he did not use a population of strategies generated by an algorithm like ours.

[17] Simulation involves assuming a starting point and having a rule for stopping: the payoff can then be the average over the whole or part of the simulation. Here the outcome might be sensitive to initial positions, and using the last few periods profits might introduce an 'endpoint' bias. Simulation is used as a last resort here.

[18] This is, of course, equivalent to assuming that there is no discounting.

[19] From the game-theoretic perspective, Selten (1991) for an excellent discussion, the other papers in that special edition of *Games and Economic Behaviour*, and Binmore and Samuelson (1992).

[20] Schlag (1998), Weibull (1995).

[21] Weibull (1995) chapter 2 for definitions and a detailed discussion of these two stability criteria.

[22] We are restricting our attention to pure strategies. There may also exist mixed NSS.

[23] This is because for each strategy i is that is NSS, there is a strategy k which is not NSS, such that $u_{ki} = u_{ii}$, and $u_{kk} = u_{ik}$. Weibull (1995, p. 51) for definitions.

[24] Without noise, the replicator dynamics imply that types never go extinct, but stay strictly positive but near zero. On a computational level, you need to specify a 'cut off' point at which you treat the proportion as zero. This can lead to 'underflow' problems highlighted by Nachbar (1992). In our simulations, we recalculated payoffs for the 'extinct' types every ten iterations (thus allowing for them to make a comeback).

[25] Gale, Binmore and Samuelson (1995) 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 current proportion. The same story could apply here.

[26] See for example Nachbar Proposition 1, p. 313 (1992).

[27] This is reflected in the growing gap between the superfirm ATP and MATP.

[28] 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 use 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.

[29] Clearly, (0.34, 0.34) will have some basin of attraction, as may other types: however, these results indicate that the basins of attraction must be small relative to the three NSS firm types with the highest payoffs.

[30] We would like to thank a referee for suggesting this be done.

References

Abreu D. and Rubinstein, 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.

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

Bacharach, M. and Bernasconi, M. (1997). The variable frame theory of focal points: An experimental study. *Games and Economic Behavior*, **19**, 1–45.

Binmore, K. and Samuelson, L. (1992). Evolutionary stability in repeated games played by finite automata. *Journal of Economic Theory*, **57**, 287–305.

Binmore, K. and Samuelson, L. (1994). An economists perspective on the evolution of norms. *Journal of Institutional and Theoretical Economics*, **150**, 45–63.

Bone, J. and Dixon, H.D. (1999). Stability and equilibrium decision rules: An application to Duopoly. York Discussion Paper 99/33.

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

Crawford, V. (1991). An evolutionary interpretation of Vanhuyk, Battallio and Beil experimental results on coordination. *Games and Economic Behavior*, **3**, 25–59.

Cyert, R. and March, J. (1963). *A Behavioral Theory of the Firm*. Englewood Cliffs, NJ. Dixon, H. (2000). Keeping up with the Joneses: Competition and the evolution of collusion. *Journal of Economic Behavior and Organization*, **43**, 223–238.

Dixon, H., Wallis, S. and Moss, S. (1995). Axelrod meets Cournot: Oligopoly and the evolutionary metaphor part 1. York Discussion Paper 95/8.

Friedman, D. (1996). Equilibrium in evolutionary games: Some experimental results. *Economic Journal*, **106**, 1–25.

Gale, J., Binmore, K. and Samuelson, L. (1995). Learning to be imperfect: The ultimatum game. *Games and Economic Behavior*, **8**, 56–90.

Gale, D. and Rosenthal, R.W. (1999). Experimentation, imitation and stochastic stability. *Journal of Economic Theory*, **84**, 1–40.

Hall, R. and Hitch, C. (1939). Price theory and business behavior. *Oxford Economic Papers*, **2**, 12–45.

Harsanyi, J. and Selten, R. (1988). *A General Theory of Equilibrium Selection in Games*. MIT Press, Cambridge, Mass.

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

Klemperer, P. and Meyer, M. (1988). Consistent conjectures equilibria. *Economics Letters*, **27**, 111–115.

Mehta, J., Starmer, C. and Sugden, R. (1994a). The nature of salience: An experimental investigation of pure coordination games. *American Economic Review*, **84**, 658–673.

Mehta, J., Starmer, C. and Sugden, R. (1994b). Focal points in pure coordination games. *Theory and Decision*, **36**, 163–185.

Miller, J.H. (1993). The coevolution of automata in the repeated Prisoners Dilemma. *Journal of Economic Behavior and Organization* (forthcoming).

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 Compatition and Cooperation*. Ph.D. 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.

Moss, S., Dixon, H. and Wallis, S. (1995). Evaluating Competitive Strategies. *Intelligent Systems in Accounting, Finance and Management*, **4**, 245–258.

Nachbar, J. (1992). Evolution in the finitely repeated PD game. *Journal of Economic Behavior and Organization*, **19**, 307–326.

Nelson, R. and Winter, S. (1982). An Evolutionary Theory of Economic Change. H.U.P., Cambridge, Mass.

Oechssler, J. (1997). Decentralization and the coordination problem. *Journal or Economic Behavior and Organization*, **32**, 119–135.

Oechssler, J. (1999). Competition amongst conventions. *Computtational and Mathematical Organization Theory*, **5**, 31–44.

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.

Schlag, K. (1998). Why imitate, and if so how? An boundedly rational approach to the multi-armed bandit problem. *Journal of Economic Theory*, **78**, 130–156.

Selten, R. (1991). Evolution, learning and economic behavior. *Games and Economic Behavior*, **3**, 3–24.

Simon, H.A. (1947). *Administrative Behavior: A Study of Decision-Making Processes in Administrative Organizations*. MacMillan.

Thomas B. (1985). On evolutionry stable sets. *Journal of Mathematical Biology*, **22**, 105–115.

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

Van Huyk, J.B., Battalio, R. and Beil, R. (1990). Tacit coordination games, strategic uncertainty and coordination failure. *American Economic Review*, **80**, 234–248.

Van Huyk J.B., Cook, J. and Battalio, R. (1997). Adaptive behavior and coordination failure. *Journal of Economic Behavior and Organization*, **32**, 483–503.

Vega-Redondo, F. (1993). Competition and culture in an evolutionary process of equilibrium selection: A simple example. *Games and Economic Behavior*, **5**, 618–631.

Vega Redondo, F. (1997). The evolution of Walrasian behavior. *Econometrica*, **65**, 375–384.

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

Weibull, J. (1995). *Evolutionary Game Theory*. MIT Press.