J Evol Econ (1997) 7: 219–254

Journal of
Evolutionary
Economics


# Using co-evolutionary programming to simulate strategic behaviour in markets

**Tony Curzon Price***

Energy Policy Group, Imperial College, 48 Princes Gardens, London SW7 2PE, UK
(e-mail: j.curzon-price@ic.ac.uk)

**Abstract.** This paper describes the use of a genetic algorithm (GA) to model several standard industrial organisation games: Bertrand and Cournot competition, a vertical chain of monopolies, and a simple model of an electricity pool. The intention is to demonstrate that the GA performs well as a modelling tool in these standard settings, and that evolutionary programming therefore has a potential role in applied work requiring detailed market simulation. The advantages of using a GA over scenario analysis for applied market simulation are outlined. Also explored are the way in which the equilibria discovered by the GA can be interpreted, and what the market analogue for the GA process might be.

**Key words:** Industrial organisation – Evolutionary programming – Genetic algorithms – Strategy selection – Learning

**JEL-classification:** C61; C72; D43; D44; L13; L94

## 1 Introduction

Evolutioary programming (EP)[1] techniques have been recognised by a number of economists[2] as being of potential use to the discipline. These researchers have tended to apply EP methods to relatively complex

* This work was in part supported by the ESRC and Scottish Nuclear Ltd. I wish to thank David Harbord, Nigel Lucas and two anonymous referees for helpful comments.

[1] EP comprises a number of techniques, the major ones being Genetic Algorithms (GAs) (Holland, 1975, 1992) and Genetic Programming (Koza, 1992).

[2] An overview is presented in Lane (1993).

problems that require advanced tools even for analytical solution[3]. In this paper, I describe the results of applying a simple genetic algorithm to several more standard models: Bertrand and Cournot competition, a vertical chain of monopolies, and a simple model of an electricity pool.

There are two reasons for applying EP to well known economic models. Firstly, most modellers would want a technique that is to be useful in complex settings to provide valid results in standard micro-economic games. The work presented here shows that EP generates interesting results in these simple settings, which is encouraging for further work on more complex problems[4].

The second reason for addressing well understood games is that EP has a potential vocation for applied simulation work, in which the underlying economic models are often quite simple, but the simulations complicated and richly detailed in important ways[5]. This potential was recognised by the precursors of EP in economics[6], but more recently it has been overlooked in favour of a more theoretical view of EP's role[7]. The advantage of EP in applied simulation is that it can inject plausible behavioural elements into models which either have no behavioural elements at all (equilibrium judgements are made "off-model"), or implausible ones (where the computational sophistication attributed to agents is incredible, or the data that is assumed available to them happens to fit just the right mathematically tractable forms[8]).

---

[3] A good example is Marrimon *et al.* (1990), who examine whether a simple classifier system can reproduce the results of a stochastic dynamic programming model. One successful application of GA simulation is in the study of repeated prisoner's dilemmas. See Marks (1992).

[4] Fudenberg and Levine (1996) make this point regarding learning rules in games: "... sensible rules should do reasonably well in a broad set of circumstances. Rules that do poorly even in simple environments are not likely to be used for important decisions."

[5] A good example of this sort of analysis is described in von der Fehr and Harbord (1996), where the issue is to examine "the potential for the emergence of effective competition in the interconnected, inter-state [Australian electricity] market, under various alternative scenarios for the horizontal market structure of generation". A simpler piece of analysis but in a similar vein is described in Lucas and Taylor (1993).

[6] See Nelson and Winter (1982) and Anderson (1994).

[7] Holland and Miller (1991), for example, write that "The artificially adaptive agent models [these are broadly speaking what I am calling EP models] complement current theoretical directions; they are not intended as a substitute. Many of the most interesting questions concern points of overlap between artificially adaptive agent models and classical theory." I entirely agree with this, but would only add that there are also many interesting questions at the overlap of EP and applied work. Arifovic (1996), Chen *et al.* (1996) and Noe and Pi (1995) also move in the direction of bringing EP closer to empirical work. These papers examine the extent to which EP techniques fit the game-playing behaviour of experimental subjects, and the results are encouraging.

[8] To keep applied game theory simulations tractable, modellers regularly have recourse to assumptions of continuously differentiable functions. For example, Green and Newbery (1992) use supply function equilibria to model the England and Wales electricity pool, and their equilibria are shown by von der Fehr and Harbord (1993) to break down under the discrete bid structures required by the market rules. It is very difficult to determine *a priori* whether results are robust to the relaxation of the simplification used.

## 2 Description of the model

GAs are a special sort of search algorithm. All such algorithms can be thought of as ways of exploring the space of possible solutions to a problem, and selecting one (or several) possible solutions as being optimal[9] (or just satisfactory). The GA uses a close analogy with Darwinian evolutionary search to select possible solutions: a number of solutions are evaluated for "fitness", and the fitter solutions reproduce, recombine, and possibly mutate. The average fitness of solutions tends to increase, and the algorithm stops searching either after a specified number of generations, or once some other externally defined criterion is satisfied. Thus, GA's are a search method which use an evolutionary process to generate increasingly good solutions to the problem posed. The driving idea behind their use is that natural evolution has solved some extraordinarily complex design optimisation problems; simulating this process may allow us also to solve complex optimisations.

### *2.1 A simple example: The price choice of a monopolist*

To get an idea of how this method can be translated into an economic context, here is a very simple application: using the GA to determine a monopolist's optimal pricing strategy. Take the simple analytic model defined below:

The monopolist faces:

- a linear demand curve:

  $Q = k - m\,P$

where $Q$ is quantity demanded, $k$ and $m$ constants, and $P$ is price charged.

- constant average cost, $C$.

  His profit function is therefore:

  $\pi M = Q(P - C)$

  where $P$ is his choice variable.

  Profit is maximised at $d\pi M/dP = 0$.

For $k = 32$ and $m = 0.5$, this entails $P = 36.5$.

How does the GA represent and solve this model[10]? The 6 steps involved are[11]:

[9] A thorough description on GAs is given in Holland (1975, 1992). A good introduction is Mitchell (1996). Kane (1996) has an interesting discussion of the "fitness landscape" metaphor in economics.

[10] The description provided here is not intended to be entirely general: GA implementations can vary greatly in how the detail is worked out. A thorough discussion of various types of implementation is provided in Goldberg (1988). Mitchell (1996) is a good introduction.

[11] Annexe A provides a standardised, tabular representation of the configuration of the GA in all the cases presented in this paper.

1. Construct a market simulation. In this case, the market simulation takes as input the choice variable (the monopolist's price) and yields as output his profit.
2. Develop a representation of strategies that can code for all possible strategies (in this case prices). A binary representation[12] is used here. We limit the monopolist's search domain to prices between 0 and 63, which allows us to represent all possible strategies as a six digit binary number. So, for example, a price of 31 is represented by 011111, and a price of 1 by 000001. A strategy thus coded can be compared to a gene, since it provides the instructions that react with the environment (the market simulation) to determine fitness (profit)[13]. In more complicated situations, for example, where both a price and a quantity have to be selected, the strategy is made up of a number of genes, and is analogous to a chromosome.
3. Create a large number of possible strategies (the "population"). In this example, 100 genes were created randomly by setting each digit in the gene to either 1 or zero using a random number generator.
4. Perform a large number of tournaments[14] in which a strategy is picked from the population at random and evaluated in the market simulation. In this example, there were 100 tournaments.
5. Pick a number of the "fittest" strategies[15] (i.e. prices yielding most profit) to breed and allow the least fit to disappear from the population. The breeding method used here resembles genetic recombination: two parent strategies are chosen, a crossover position[16] is chosen at random, and two separate offspring are created, one each for the two ways of sharing the parents' genetic information. For example, if prices 1 (000001) and 31 (011111) were chosen as parents, with a crossover position of 4, the two offspring prices would be 3 (000011) and 29 (011101). It is also possible in this step to include a mutation operation which randomly "flips" the value of a bit with some probability.

[12] Although binary coding is often used in GAs, there are no hard and fast rules for what the best coding is. Interesting discussions of alternative coding are given in Davis (1991).

[13] The gene has to be interpreted by the simulation to provide its interpretation. In the example given (and all but one examples below), the interpretation is the simple one of assuming that the binary digits represent integer numbers. But more indirect "mappings" from the genes to the choice variables are possible – six bits could also be used, for example, to represent numbers between zero and 1 in gradations of $1/64^{\text{th}}$.

[14] The term "tournament" is more applicable to the competitive simulations described below, but is kept for consistency. Moreover, in the simple case of the monopolist's price, it would be possible to evaluate every possible strategy in the population, rather than picking them at random. However, this quickly becomes infeasible in competitive games where the number of permutations of strategies becomes large (with two players each with a pool of 100 strategies, $100^2$ simulations would need to be carried out to test every strategy against every other).

[15] The algorithm is often sensitive to the number of strategies picked for breeding, and this parameter can be interpreted as the "single-mindedness" of the search. This is discussed in more detail below.

[16] The crossover position is defined as being "a number of positions (bits) from the left of the start of the gene.

6. With the new population created by step 5, check whether the termination criteria are met, and if not start the tournaments again (Step 4). In the economic games described below, the termination criterion was whether evolutionary stability seemed to have been reached[17].

   Each repetition of steps 4 to 6 are counted as one generation.

The search method described in these steps probably does not at first sight seem to have any obvious economic interpretation[18] (what, for example, is the analogue of strategy recombination within the firm's decision making process?), and the reader could justify wonder at this stage why the GA should be of any more interest than another optimisation algorithm. This is an important question, discussed at greater length below. Suffice it to say here that a number of possible analogical interpretations can be offered, from the naive (the strategies are actually tried in the market by firms) to the more sophisticated (strategy testing in the GA is analogous to the corporate world's beloved scenario analysis, in which case the GA replicates the firms' acquisition of knowledge about its environment). Each interpretation can yield interesting insights. But exploration of these questions must come after the description of results.

Figure 1[19] is a representation of the state of the population of strategies after 5, 10 and 15 generations[20]. The population is comprised of 100 possible prices each between 0 and 31, so the entire population at any point in time can be described as a frequency distribution of prices, as shown in the graph. The lozenges show the price frequency distribution of the population after 5 generations, the squares after 10 generations and the triangles after 30 generations.

---

[17] "Evolutionary Stability" is defined in Maynard Smith (1974). A strategy is evolutionarily stable if, under the relevant dynamic, it is resistant to mutant invasions. The analytical treatment of dynamic equilibrium notions such as this one has recently received considerable attention, although general results seem to be difficult to obtain. An excellent overview is provided in Fudenberg and Levine (1997, especially chapters 3 and 4). In section 4.9, Fudenberg and Levine consider the particular case of GA simulation, and make several points that go very much in the direction of this paper.

[18] Indeed, it is even obvious that the steps constitute a good optimisation procedure. Holland (1975, 1992) shows that the GA is a good all-purpose search method, navigating well the difficult route between exploration of new possibilities and exploitation of "sure bets". The proof is quite difficult (the population dynamics of the system are clearly complicated), but relies on the two facts that parent fitness is a good predictor of offspring fitness, and that the GA tends to increase the proportion of fit "sub-genes" (schema). An excellent overview of the proof is given in Mitchell (1996).

[19] For visual clarity, the "Price" axis shows only values in the region of the solution – between 25 and 53. The entire price range was between 0 and 63.

[20] Many engineering applications of Gas are essentially concerned only with the fittest elements of the population. However, the equilibrium and behavioural interpretations placed on the GA require an analysis of the population as a whole (this is the approach taken throughout section 3, where the result of a GA simulation is taken to be an entire set of population distributions). Schematically, we view the population as the reservoir of ideas from which the firm (eventually) comes to randomly pick its actual market strategy.

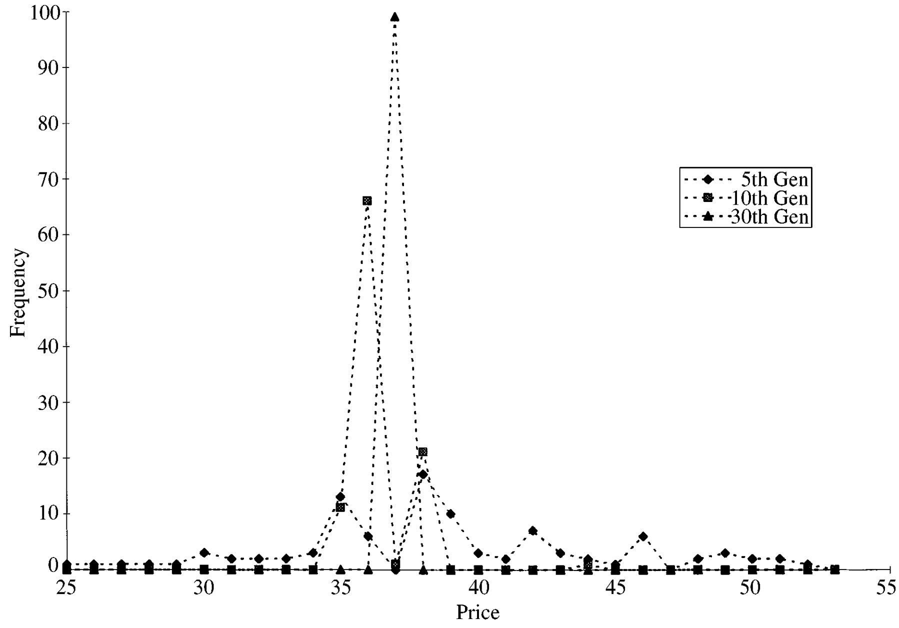


**Fig. 1.** Price frequency distribution for the simple monopolist

There are several things to note about the graph. Firstly, the GA this was run on allowed only integer solutions, so that the $30^{th}$ generation result with 100% of prices at 37 is optimal. This leads to the second observation: the monopoly problem with integer constraints is substantially more complex than the analytic one described above, and yet the GA has found on optimal solution very quickly. In fact, even by the $5^{th}$ generation, we see a remarkable amount of "organisation" in the population of strategies[21]. The $10^{th}$ generation distribution is entirely composed of prices between 35 and 38, and by the thirtieth generation, every non-optimal price has been driven out of the firm's set of possible strategies.

A naïve interpretation of this result would be that even a monopolist devoid of capacities for rational optimisation, but blindly following evolutionary rules to select strategies, would very soon be behaviourally indistinguishable from the rationally calculating monopolist usually encountered in economic theory. A slightly more sophisticated interpretation (see Section 4) is that the firm uses its own internal model of its market to explore strategy space and to hone its behaviour, which it puts into practice once equilibrium is reached. On this interpretation, a monopolist using a method analogous to the GA internally is entirely indistinguishable from the optimising monopolist of ordinary theory.

[21] Remember that the starting population is randomly generated, and therefore approximately uniform.

## 3 Description of the runs

The GA method described above was applied, with only small modification, to the more interesting set of models in which agents interact strategically. The models simulated were Bertrand competition (stiff price competition with no capacity constraints), Cournot competition (competition in quantities), a simple chain of monopolies model (a monopolist manufacturer sells to a monopolist retailer who sells on to consumers, in which the first two simultaneously choose prices), and two versions of an electricity pool model (represented by a sealed bid auction)[22].

The application to the monopoly described above involved only a single agent, and no strategic interactions in the market. All the runs described below involve market simulations in which the performance of one strategy depends on the other strategies present. In terms of the six steps involved in the simulation, not very much changes. A tournament is now set up by selecting (at random) a single strategy from each player, which determines a "possible market".

However, there is a considerable increase in the complexity of the mechanisms at work. Each player has a separate population of strategies. In each population, the "fittest" strategies (the ones that reproduce most) are fittest only *relative to the other populations*, so that the state of each population after a number of generations depends on the past and present state of all others. Thus, the optimal state of each population depends on the states of all other populations. An agent's strategies that tend to be good in one generation will affect the mix of strategies in the other agents' populations (by affecting their fitness, and therefore their chance of reproduction) in the next generations (which will affect the original agent's fit strategies in generations after that, et cetera ...). Mitchell (1996) characterises co-evolution as follows:

> " ... in nature [organisms] evolve defences to parasites that attack them only to have the parasites evolve ways to circumvent the defences, which results in the hosts evolving new defences, and so on in an ever rising spiral – a "biological race." " [23]

In the competitive game simulations described below, when one firm discovers a good strategy, the other is spurred to discover a better one, which encourages the first to find a riposte, et cetera. This is the essence of co-evolution[24].

[22] The analytic treatment of the first three is standard. A clear exposition is found in Tirole (1989). The last is a simple version of the model developed by von der Fehr and Harbord (1993).

[23] Mitchell (1996), p. 26.

[24] The process of co-evolution is thus already very close to many economists' views of competition (notably in Schumpeter's "gales of creative destruction").

### *3.1 Bertrand competition*

#### The Bertrand model

The textbook case of Bertrand competition occurs when two producers of identical goods face no capacity constraints, equal (constant) average costs, downward sloping demand, and compete on price by simultaneously offering the price at which they are prepared to supply. The market simulation can be represented as follows:

- market price is the lower of the two producers' prices;
- the low-price producer's revenue is the market price times the quantity demanded at the price;
- the high-price producer's revenues and costs are zero;
- in the case of a tie on price, each producer satisfies half the market.

#### The analytic solution

The Nash equilibrium of this game is fairly intuitive[25]: each producer wants to price below the other producer, as long as the price exceeds cost. The two firms are choosing prices without knowing what the other has offered (simultaneity), so each has to predict what the other will do. Each firm works out that the other will not price above cost, since that would lead to the easy riposte of pricing a minute amount above cost, capturing the whole market, and making a minute profit. Moreover, each firm predicts that the other will not price below cost, since that would entail a loss for one of the two firms (and would therefore be irrational). So the only coherent prediction seems to be that each firm prices at cost, and makes zero profits.

#### The GA simulation

Bertrand competition was modelled with the GA in the following way:

1. The market simulation was based on a demand curve given by $Q = 32 - 0.5\ P$ where $Q$ is quantity demanded, and $P$ is market price. The market price was determined in a tournament as being the lower of the two prices selected.
2. Strategies (prices) were coded as strings of six binary digits, and interpreted in the simulation as integers between 0 and 63.
3. Two populations of strategies were created, one for each firm, containing 100 genes each, and initially set to random values.
4. Tournaments were created by randomly choosing one strategy from each population, which were evaluated in the market simulation. Each generation was comprised of 200 tournaments.

[25] This is described more formally in Tirole (1989), pp. 210–211.

5. Birth, death and mutation was carried out for each population separately, as described above.
6. The simulation was stopped after an equilibrium has been reached[26].

Figure 2 shows the evolution of market price over the 300 generations the GA ran for. By the end of the run (in fact, by the 150$^{th}$ generation), the market price in tournaments had fallen to 1. As usually presented, the Nash equilibrium of the game has price falling to average cost, and profits falling to zero. The discrepancy between this and the GA result arises from the fact

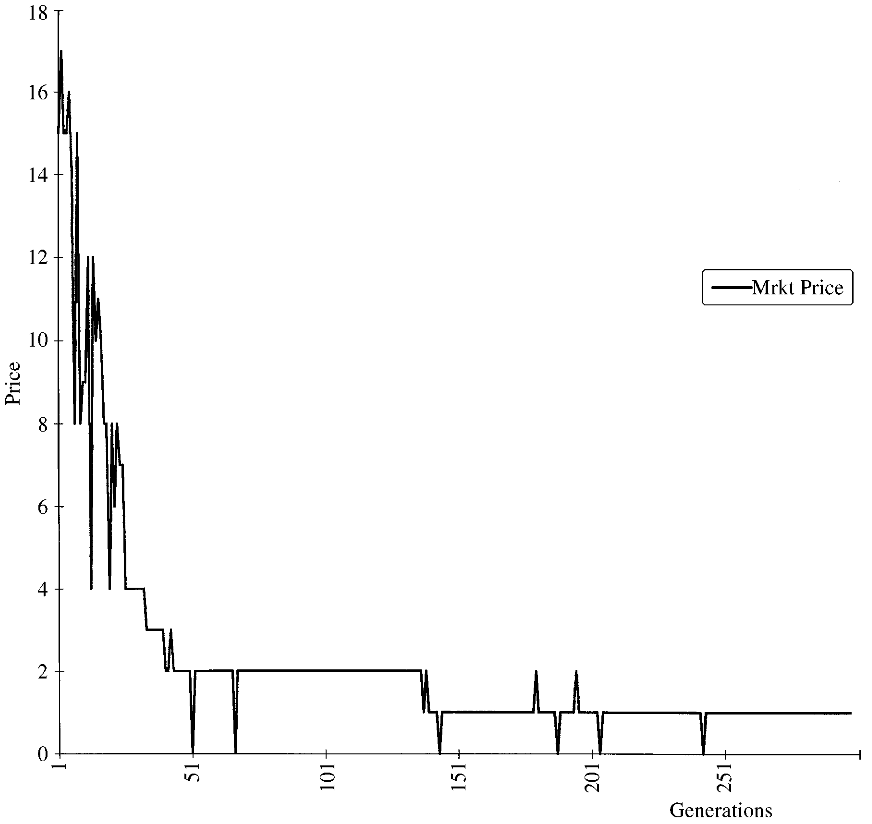


**Fig. 2.** Evolution of price over run

---

[26] In all but the second part of section 3.4, the approach taken to recognising states as equilibria is pragmatic. Three questions can be asked of the system: first, are the population distributions stable? If yes, can we understand why these are equilibria? If yes, do we return to recognisably similar points in all similar runs? When all three are true, we judge that an equilibrium has been reached. As noted in footnote 17, some analytical work on the dynamics of systems similar to GAs has been done. However, these results do not have the generality to be able to make equilibrium predictions about the GA dynamic in general, so some pragmatic procedure for judging equilibrium imposes itself.

**Table 1**

| Price | Frequency Population 1 (%) | Frequency Population 2 (%) |
|---|---|---|
| 0 | 1 | 0 |
| 1 | 99 | 100 |

that the GA as set up only allows integer prices to be offered. Thus, each population has a tendency to undercut the other population as long as price exceeds 0, which is at a price of 1[27]. The GA result is thus a Nash equilibrium given the integer constraints.

Table 1 shows the frequency with which prices are encountered in each population by the end of the run. Almost all surviving price are at $p = 1$, and the one that is not, $p = 0$, we can safely assume has arisen out of mutation and will not survive (since if both set price at 1, they share the market and make a small profit (of 15.5) whereas at $p = 0$, no profit is made at all).

Figure 3 shows the average fitness (profitability) of each firm's population of strategies. In the early part of the run, average fitness is high. However, co-evolutionary competition soon drives profits down for both in a seemingly hap-hazard way. This is followed by periods of relative stability, punctuated by rapid change (for example after generations 129 and 257).

### *3.2 Cournot competition*

#### The Cournot model

Cournot competition arises in the following sort of setting[28]:

> "Two producers work in isolation preparing the quantity they bring to market. These quantities are decided upon with a foreknowledge of this market structure and with knowledge of the characteristics of demand, but neither side gets to see how much the other is producing. Each producer brings the quantity it produced to a central market place, where it is handed over to a "state sales commission" that sets price so that the total amount brought to market is just demanded at the set price."

In other words producers compete over the quantity produced. When this occurs in a market with ordinary cost and demand functions, it is easy to show that when one agent increases quantity, the other should reduce[29].

[27] When both price at 1, they share the market.

[28] This is taken from Kreps (1990), p. 443.

[29] Quantity is said to be a strategic substitute. See Tirole (1989), pp. 218–220.

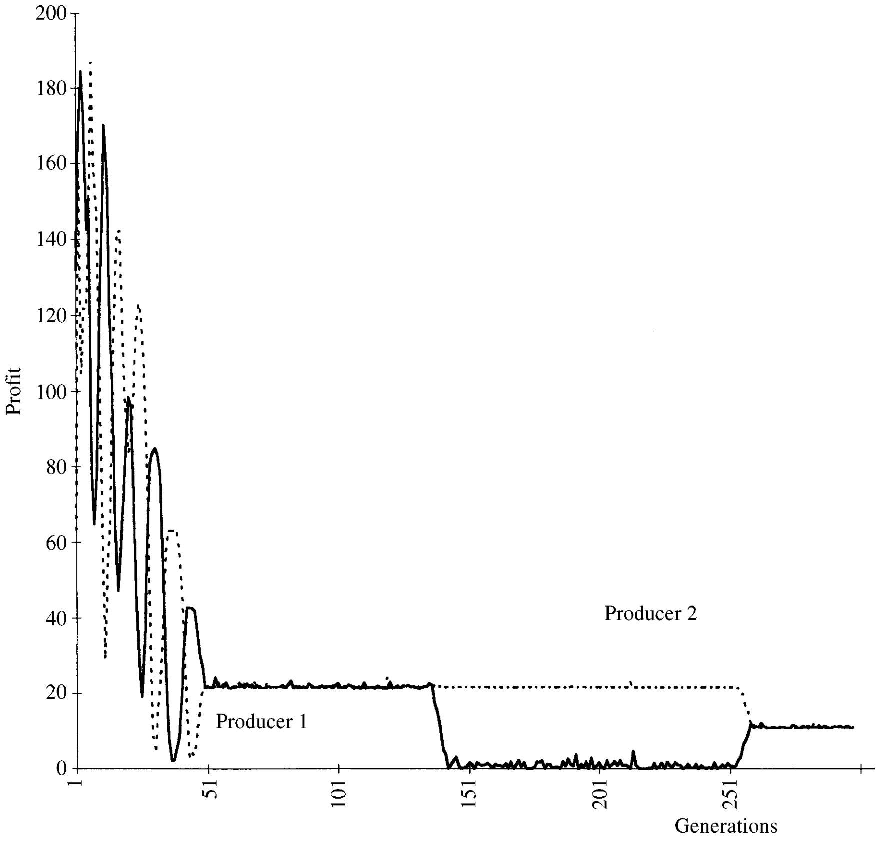


**Fig. 3.** Average population profit

## The analytic solution

The market simulation used can be analytically described as follows:

- Demand is defined as:
  $P(q_1, q_2) = 62 - 2(q_1 + q_2)$, where $P$ is price, $q_1$ is the first producer's output and $q_2$ the second's.
- profits are therefore
  $\pi_1 = q_1(P(q_1, q_2) - 8)$, where $\pi_1$ is the first producer's profit, and 8 is the average cost of production. The second producer's profit function is $\pi_2 = q_2(P(q_1, q_2) - 8)$.
- Nash equilibrium requires that $d\pi_1/dq_1 = 0$, and $d\pi_2/dq_2 = 0$ which occurs when $q_1 = q_2 = 9$.

## The GA simulation

The GA was set up as in the Bertrand game, except that strategy genes are now coded as five binary digit strings and interpreted as being quantities (interpreted as integers between 0 and 31). The payoffs in the market simulation are as described above.

Figure 4 shows the most frequently occurring strategy in each population, and Figure 5 shows the evolution of profits for the two players over the 110 generations for which the model was run. By the end of the run, $q_1 = q_2 = 9$ (the most frequently occurring quantities, and the analytic solution) represent over 90% of all strategies. The system is thus clearly attracted to the Cournot Nash equilibrium.

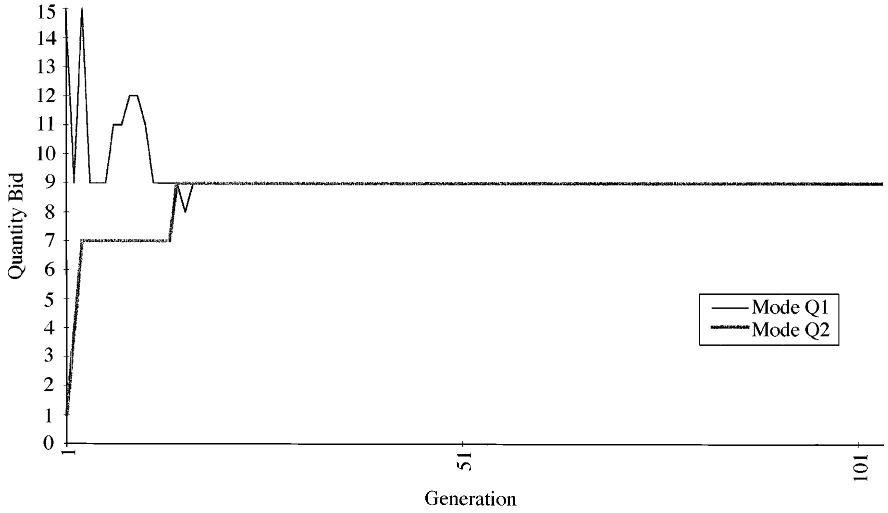


**Fig. 4.** Mode of player's quantity bits

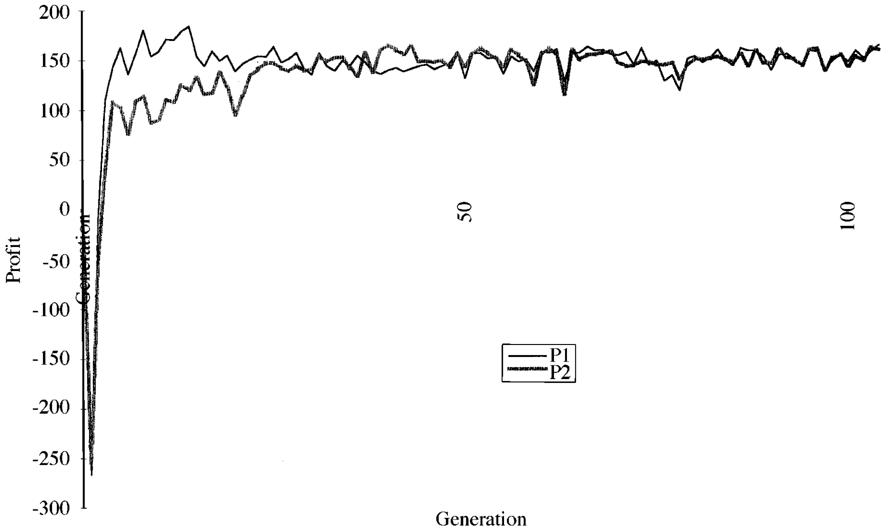


**Fig. 5.** Average profit in the cournot game

### *3.3 The chain of monopolies*

The chain of monopolies model

In the chain of monopolies model[30], we assume a monopolist producer is selling to a monopolist retailer, who sells on to a consumer represented by a demand curve known to both players. Each chooses price simultaneously. This model is mainly interesting in industrial organisation because it provides the simplest demonstration of the slightly paradoxical fact that a chain of monopolists supply a smaller quantity at a higher price than an integrated monopolist would. In terms of GA modelling, the example is interesting in that it provides a setting for heterogeneous players in the simulation: the producer's payoffs depend on, but are structurally different from, the retailer's payoffs.

The analytic solution

The market simulated can be described analytically as:

- Demand is given by $Q(p_1, p_2) := 32 - 0.5(p_1 + p_2)$, where $Q$ is demand, $p_1$ is the producer's price and $p_2$ the retailer's price.
- The manufacturer is assumed to have average costs of 1 and the retailer of 8, so the profits of each are given by:

  $$\pi_1 = Q(p_1, p_2)(p_1 - 1)$$

  $$\pi_2 = Q(p_1, p_2)(p_2 - 8)$$
- The solution requires that $d\pi_1/dp_1 = 0$ and $d\pi_2/dp_2 = 0$, which occurs at $p_1 = 19\ 1/3$ and $p_2 = 26\ 2/3$.

The GA simulation

The GA simulation was set up as in the Bertrand competition case, with the only difference being the payoffs in the market simulation are now determined by the profit functions given above.

Figure 6[31] shows the price frequency distribution of strategies after 292 generations. The GA solution is very close to the analytic solution: 97% of prices at 19, and 91% at 27 (the GA was constrained to integer solutions). A very similar picture is already apparent after 125 generations. Thus, the system seems to behave similarly to the Cournot case in this game. The exogenous market simulation parameters were identical in this run to the simple monopolist described above. We can thus easily see the GA reproducing the standard "double wedge" result.

[30] The vertical chain of monopolists is described in Tirole (1989), p. 174.

[31] For visual clarity, the "Price" axis has been truncated.

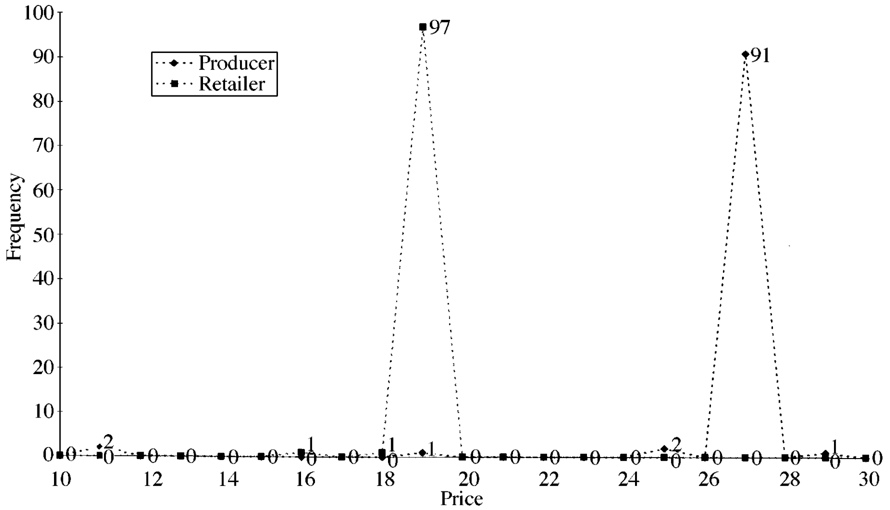


**Fig. 6.** Price frequency distribution for serial monopolists

### *3.4 A simple electricity pool*

The model used here is a simple version of that developed by von der Fehr and Harbord (1993) to represent the structure of the UK electricity market. In the UK, power producers bid their generating plant into a pool, where a market price is determined as the bid of the last producer required to satisfy demand. All producers who have bid lower than this receive the market price for their output[32].

#### The pool model

At a first level of simplification, this market is modelled by assuming that neither of two producers has sufficient capacity to satisfy market demand, but that both together have more than enough[33]. Therefore, if the two bid different prices for their output, the higher priced producer manages to sell less than his full capacity, whilst the lower priced producer sells his full capacity at the high price. We assume that in the case of a tie, the market is shared equally, and that there is a maximum price beyond which demand is zero[34].

In this market simulation, the pay-offs can be characterised as follows:

- $\pi_1(p_1, p_2)$, the profit of generator 1, is a conditional function of $p_1$ and $p_2$, the bid prices of the two generators:

[32] This is a great simplification on the operations of the England and Wales pool. Newbery (1995) contains an introduction to the intricacies of the market rules.

[33] This is only one of the several cases considered by von der Fehr and Harbord (1993).

[34] von der Fehr and Harbord (1993) justify this "ceiling price" by arguing that there is an electricity price above which the regulator would intervene to redress behaviour.

$10p_2$ if $p_2 > p_1$
$8p_1$ if $p_1 > p_2$
$9p_1$ if $p_1 = p_2$
$10p_1$ if $p_2 > 45 > p_1$
$0$ if $p_1 > 45$

- $\pi_2(p_2, p_1) = \pi_1(p_1, p_2)$

The first two conditions in the pay-off function simply say that if generator 1 is the low bidder, then he sells 10 units at the bid of the high bidder, but that if he is the high bidder, he sells only 8 units at the price that he bids (average cost is assumed to be zero). In other words, each producer knows that he will be a monopolist over the residual quantity. The third condition deals with the case of a tied bid, in which case the market is shared. The fourth condition deals with the case in which one of the two bids is higher than 45, in which case the lower bidder sells his entire capacity at his bid price. The last condition states that a bid above 45 yields zero profits for the bidder in question.

## The analytic solution

In the analytic solution to this game, there are two pure and one mixed strategy Nash equilibria. The pure strategy equilibria are fairly intuitive: there is a price so low that if player 1 knows player 2 will play it, then player 1 prefers to bid high; therefore, if player 2 bids this price, he knows that player 1 will bid the maximum. With the payoffs given above, this occurs at $p_2 < 36$, since player 1 prefers to sell only 8 units for 45 (making a profit of 360) than 10 units at less than 36 (making a profit smaller than 360). Thus, one player bidding the maximum while the other bids sufficiently low are the two pure strategy Nash equilibria[35].

In a mixed strategy equilibrium, each producer chooses to play a number of strategies with some probability. The producer picks the probability for each strategy such that the expected pay-off (i.e. before play) is equalised across all possible strategies his opponent might adopt. If both players are choosing randomising probabilities in this way, then the outcome is by *construction* a Nash equilibrium (if she is playing such that I am indifferent, then I could play anything; but unless I play such that she is indifferent, her best response will not be to play such that I am)[36].

The indifference property of the mixed strategy suggests a way of computing the randomising probabilities. Only when the opponent is playing an optimal mixed strategy is the player indifferent between all choices; hence setting the probabilities that the opponent uses such that the player's expected payoffs are equalised will determine the opponents optimal mix. Mixed strategy equilibria can thus be determined by solving a set of simultaneous equations (of which there will be the number of strategies

[35] Strictly speaking, there are many more pure strategy Nash equilibria, since the "low" bidder is indifferent to any price bids below the critical price, but these are all equivalent and interchangeable.

[36] A clear introduction is given in Binmore (1992), chapters 6 and 7.

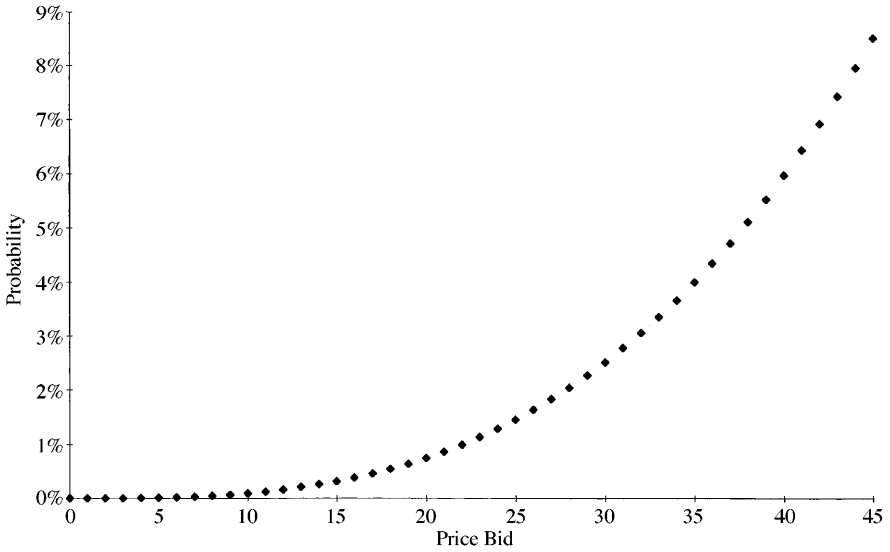


**Fig. 7.** Mixed strategy equilibrium for simple pool game

minus one independent equations, and a constraint that probabilities sum to one).

The equilibrium mixing probabilities (allowing discrete integer price bids between 0 and 45) for the von der Fehr and Harbord game were computed as described above[37], and are shown in Figure 7. If both players mix their strategies according to this probability distribution, the expected profit is approximately 364 for each player.

Figure 7 shows prices on the horizontal axis and the probability with which they should be played on the vertical axis (so a price of 40 should be played approximately 6% of the time). The randomising probabilities are seen in Fig. 7 to be increasing in price. This is a reflection of the fact that for the pay-offs described, the penalty to being the high bidder is relatively small (sales of 8 units rather than 10), so the penalty for both players bidding low is relatively large[38].

## The GA simulation

The GA was set up as for the Bertrand simulation, with the replacement of the payoff function for that described above.

Figure 8 shows the distribution of prices in the two populations after 320 generations. This corresponds exactly to one of the pure strategy Nash

[37] The Mathematica model used to compute these can be examined at URL:http://is.eu-net.ch:80/Customers/curzon/hmix22.ma

[38] The mixing probabilities were re-computed for a payoff function in which the high bidder sold only 2 units, against 10 for the low bidder. The mixing strategy is then falling in prices.

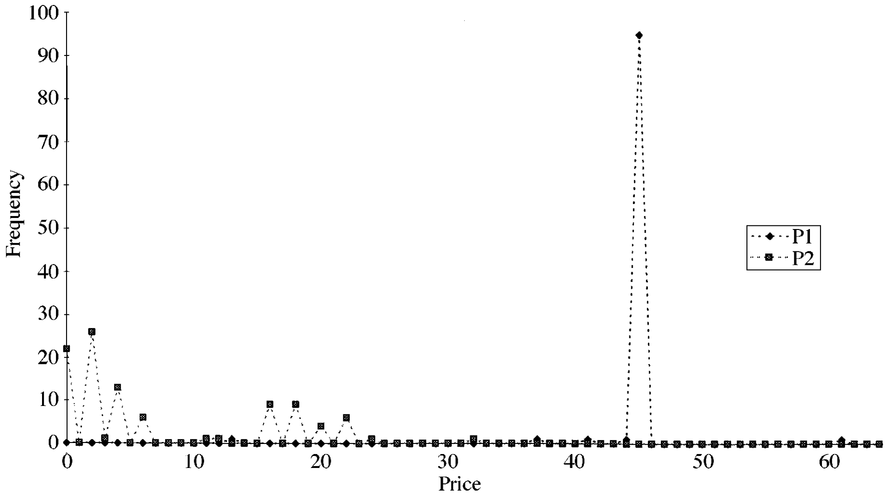


**Fig. 8.** Price frequency distribution after 320 generations

equilibria, with producer 1's strategy population consisting entirely of the maximum price bid (45), and producer 2 bidding indifferently anywhere below 36.

Which producer ends up with the higher profits (lower bids) is entirely a matter of chance. To see in more detail how this occurs, Fig. 9 shows the evolution of profit streams for the two producers over the length of the run. The high/low profit outcome appears soon after the tenth generation, then seems to stabilise for about 130 generations. However, between generations 150 and 180 (approximately), the identities of the high and low bidders switches around, and produces a new period of apparent stability to the end of the run.

To examine in more detail the dynamics of this switch, we need to look at the evolution of the frequency distributions over the course of the run. This is shown in Fig. 10, which maps, for each player, the frequency of price occurrences in the populations every five generations, from the start to generation 320. The zero'th generation is an entirely homogenous shade, showing the near uniform distribution of prices. We see that already by the $20^{th}$ generation, producer 2's strategies are bunching around high values, whilst producer 1's are still relatively undifferentiated. From the evolution of profit figure (Fig. 9), we see that neither producer is obviously more profitable than the other at this point. In fact, producer 2 sees the greatest increase in profits by tending towards the high pricing strategy. This is because, with neither strategy populations very differentiated, there is a good chance that market price will turn out low, so that influencing the probability of a high price is a good strategy for one player, even if some of the benefits accrues to the other.

Producer 1 continues to have a very low frequency of strategies above a price of 36 until about two fifths of the way through the run, where a cluster of prices between 36 and 45 can be seen appearing. With a sufficiently high

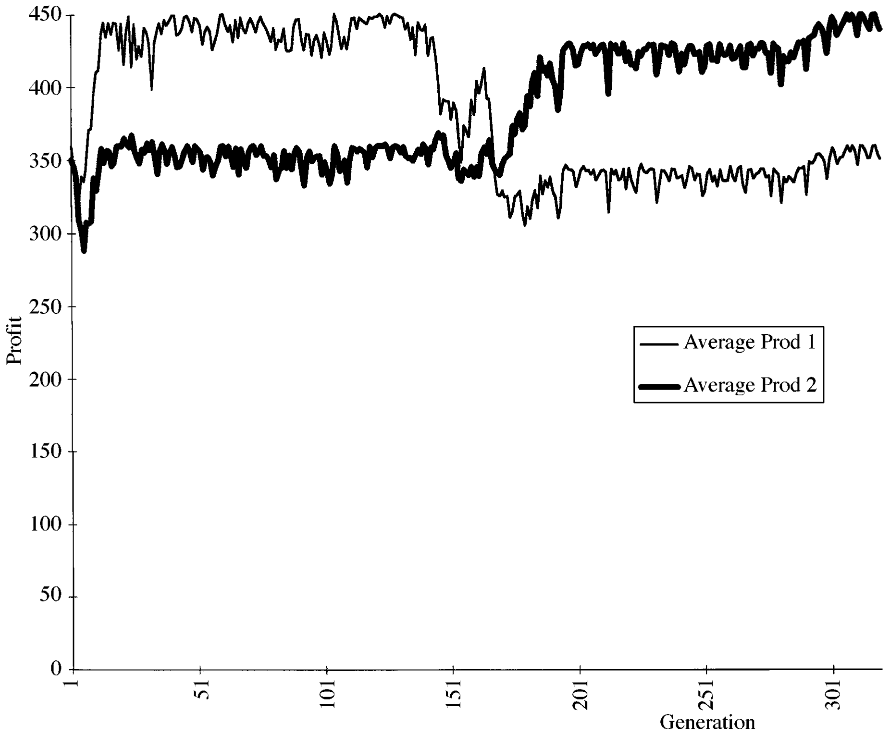


**Fig. 9.** Profits in the simple pool game

probability of prices being bid over 36, producer 2 finds advantage in undercutting producer 1, as starts to happen. The system eventually stabilises again, with producer 2 bidding low and producer 1 bidding high. The switch is thus the result of the appearance of a cluster of price bids above 36 by the previously "low" bidder. There is always the possibility that these will arise: the "high" bidder never has 100% of strategies at the maximum price (partly because mutation, partly because some strategies may, by chance, just not yet have been selected for tournament); similarly, the "low" bidder is unlikely to have no strategies above 36. During a generation, the vagaries of tournament selection may pit unusually "low" bids from the high bidder, against unusually "high" bids from the "low" bidder. Such "trembles"[39] may become self-reinforcing, leading ultimately to a switch of positions.

## The GA on mixed strategy equilibria

The GA was run with a number of different crossover rates to examine whether any would yield the mixed strategy equilibria. Although some parameter settings destroyed the pure strategy equilibrium described above (see section 4.1 for an interpretation of this result), none produced an

[39] See Selten (1983).

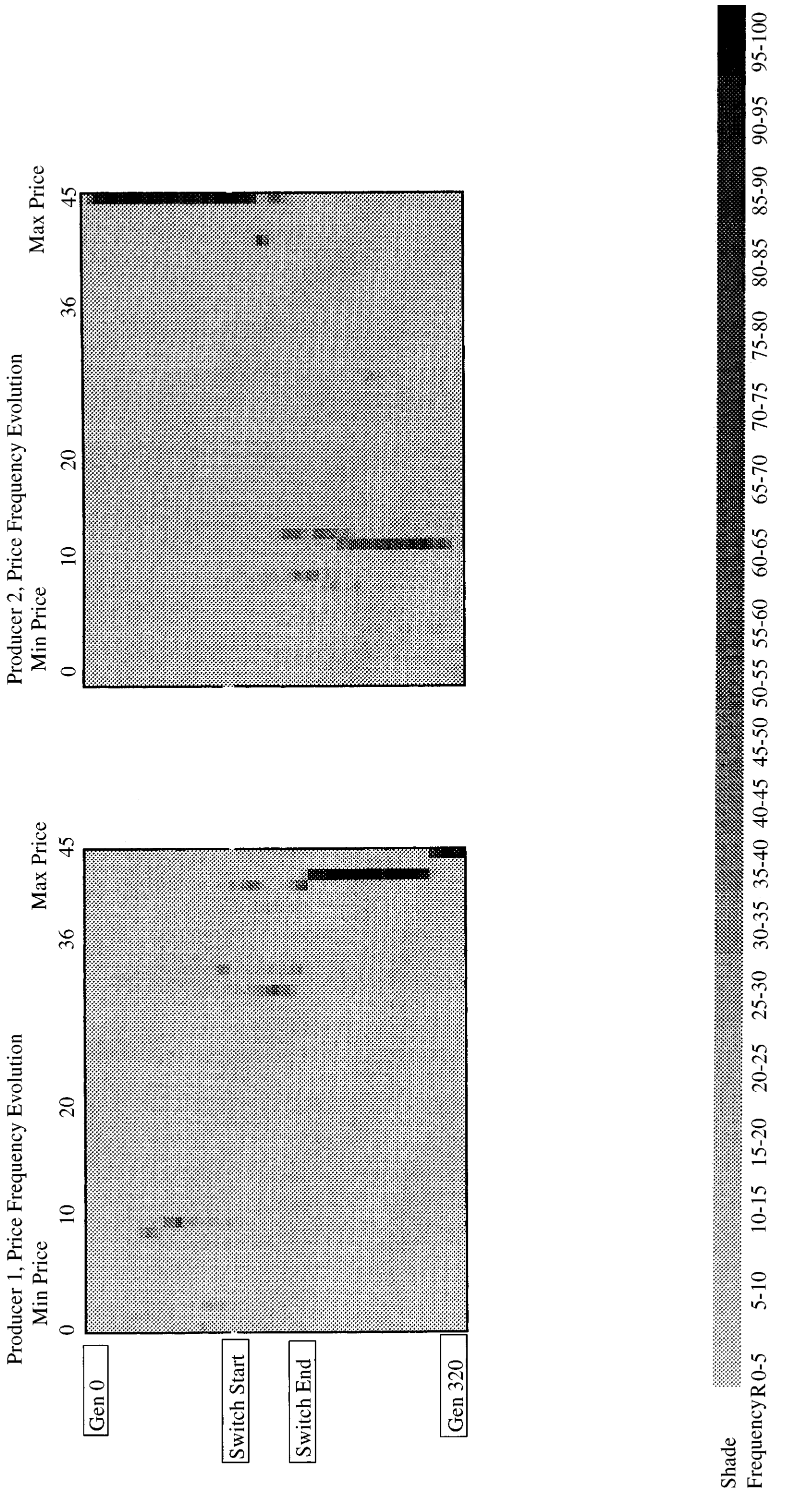

Producer 1, Price Frequency Evolution
Min Price
Max Price
0
10
20
36
45
Gen 0
Switch Start
Switch End
Gen 320
Producer 2, Price Frequency Evolution
Min Price
Max Price
0
10
20
36
45
Shade
FrequencyR 0-5
5-10
10-15
15-20
20-25
25-30
30-35
35-40
40-45
45-50
50-55
55-60
60-65
65-70
70-75
75-80
80-85
85-90
90-95
95-100


**Fig. 10**

outcome that was "recognisably" the mixed strategy equilibrium. This suggests that the mixed strategy equilibrium is too "knife-edged" to be the outcome of this heuristic process, or at least that the version of the GA used is too "blunt" to stabilise on it[40].

The second suggestion is much easier to test than the first. All that is needed is to "fix" the starting population at the mixed strategy outcome and see whether it is preserved. The first suggestion requires the testing of a large number of combinations of parameters over a large number of long runs to see whether the populations are ever "pulled" towards the mixing distributions[41].

The suggestion that the GA is too "blunt" to stabilise on the MSE was tested by fixing the populations at the start of the run to reflect the mix of strategies that would correspond to the mixed strategy equilibrium. So, for example, a price of 40 was encoded by 6% of the strategies (see Fig. 7). The GA was run to examine whether it diverged from the mixed strategy equilibrium. If the GA is given the mixed strategy equilibrium and yet does not stabilise on it, then *a fortiori*, the mixed equilibrium cannot be an equilibrium position for the GA[42].

The parameter values that yield the pure strategy equilibria very rapidly "destroy" the initial equilibrium. This can be seen in Fig. 11, which shows, for a number of reproduction and crossover rates, the evolution of the absolute deviation from the mixed strategy equilibrium. The "$R = 18$, Pop1" and "$R = 18$, Pop2" summarise the evolution of the deviation for a crossover rate of 18% . The rapid divergence of the two series corresponds to the rapid appearance of the pure strategy equilibria, with one high and one low bidder (the switch in the positions of the two series corresponds to a switch of the identities of the "high" and "low" bidders). In general, divergence of the absolute deviation paths was taken to be a sufficient condition for the disappearance of the MSE[43]. The GA was run to discover the

[40] The difficulty of defining an (analytic) dynamic learning process that actually settles on the right sequence of play to asymptotically yield the mixed strategy is noted by Fudenberg and Levine (1996) in the context of the "fictitious play" dynamic. There is no obvious reason to suppose that the GA dynamic should be particularly better suited to the task. The example cited by Fudenberg and Levine [taken from Fudenberg and Kreps (1993)] to illustrate the difficulty is a dis-coordination game that is structurally rather similar to the Pool game: in both cases, the pure strategy Nash equilibrium involves adopting a strategy *not* adopted by the other player.

[41] A very improbable succession of precisely the right tournaments, crossovers and mutations *could* yield the mixing distributions in any one generation; if the GA parameters were not such as to destroy the distribution (as in the second hypothesis), the GA would then have produced the MSE. However, the question is not one of possibility (could an improbable succession of outcomes … et cetera?) but of tendency (are there settings for which the GA usually produces the MSE?).

[42] Arifovic (1996a) uses a very similar modelling technique to show that an equilibrium predicted by an analytical model is not stable under GA dynamics.

[43] The MSE of the initial population were symmetric. Divergence between populations in the deviations from the MSE thus indicate a pull towards an asymmetric outcome. However, divergence cannot be taken to be a necessary condition, because the GA might evolve towards a symmetric stable state that is not in the neighbourhood of the initial MSE.

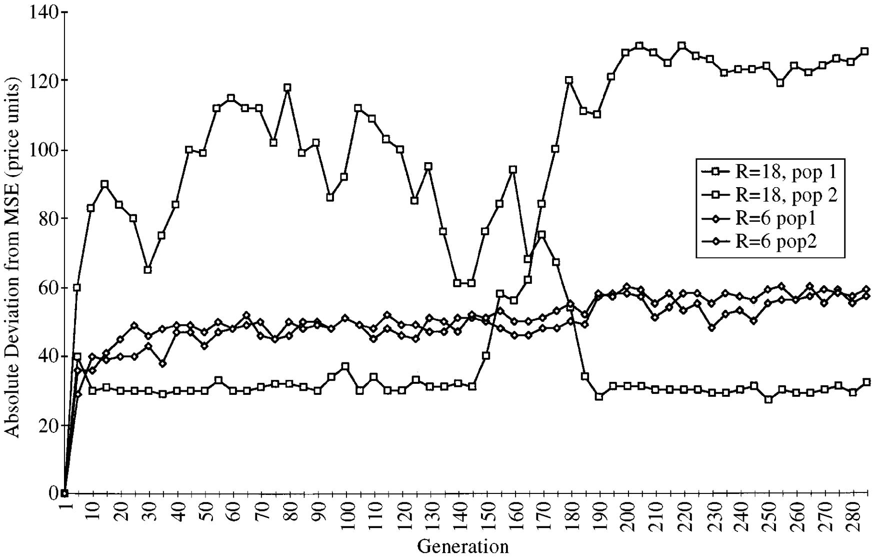


**Fig. 11.** Evolution of deviation from MSE

lowest crossover rate that avoided divergence of the two populations, which was found to be 6% ("$R = 6$ ..." in Fig. 11). There is thus a wide range of crossover rates for which the initial MSE is clearly unstable under the GA dynamic.

Whilst is has not been shown that the GA cannot settle on the MSE, these experiments suggest that, here at least, the GA is strongly predisposed to the pure strategy equilibria of the game[44]. For crossover rates that yield pure strategy equilibria, the GA tended to destroy the MSE, and the populations eventually reverted to a pure strategy equilibrium. The significance of this parameter sensitivity is discussed in section 4.1.

### *3.5 The pool model with irreversible capacity costs*

The original von der Fehr and Harbord model was used to show that as long as the situation of residual monopoly pertained, the electricity pool rules would not lead to competitive outcomes. This immediately poses the question of what productive capacity levels producers in such a market would choose: do they tend to invest only so much as to endow each with residual monopoly, or is sufficient capacity built to ensure a competitive market? The GA was run to answer this question.

[44] von der Fehr and Harbord (1993) consider a version of the pool model in which there are no pure strategy equilibria. Generators commit to price bids for 48 demand periods, between which demand can fluctuate widely. Thus, there is never the certainty that they will be residual monopolists. Work in progress with the GA suggests than in this setting mixed strategies persist (although it has not yet been shown these are equilibria).

The model

The simulation described above was modified to state that if either producer could satisfy the entire market, then the lowest bid would be chosen. Capacity choice was added as a strategic variable, with capacity commitments incurring a constant average cost. Average variable cost was maintained at zero. Thus, players offered both price and quantity bids, and the quantity offered incurred a cost whether or not the capacity was actually used to satisfy any demand.

The payoff function is defined as follows:

$$\pi_i(k_i, q_i, p_m) = q_i(p_m - c) - k_i C$$

Where

- $k_i$ is the capacity offered to market (or availability)
- $q_i$ is the quantity actually called on to produce
- $p_m$ is the market price
- $c$ is the average operating cost
- $C$ is the cost of making capacity available (the cost of the option to participate in the market)

This payoff function can be read as saying that the producer earns profits equal to net average revenue, $(p_m - c)$ times quantity called on, minus the quantity offered, $k_i$, times the cost of offering it, $C$. The determination of $q_i$ and $p_m$ are described below.

In the two player game, $k_i$ and $k_j$ are the capacity availability bids of the two players, and $p_i$ and $p_j$ are the price bids. We distinguish three regimes of availability bids, price bids and demand combinations for determining $p_m$, $q_i$ and $q_j$. In all cases, $Q$ is the market demand at market prices below $P$, and market demand at prices above $P$ is 0.

*Regime 1:*

$$k_i + k_j \leq Q \text{ and } \max\{p_i, p_j\} \leq P$$
$$p_m = \max\{p_i, p_j\}$$
$$q_i = k_i$$
$$q_j = k_j$$

This regime describes a state of unsatisfied demand[45] – both producers have bid below $P$, and the availability offered by both together is insufficient to meet market demand. In this case, each is asked to produce the full capacity bid, and market price is set at the higher of the two bids.

*Regime 2:*

$k_i + k_j > Q$ and $(k_{[S\{p_i,p_j\}]} < Q$ and $\max\{p_i, p_j\} \leq P$ where $[S\{p_i, p_j\}] = i$ if $p_i = \min\{p_i, p_j\}$, and $j$ otherwise.

[45] In the England and Wales pool, price paid to generators increases as the probability of demand being partially unsatisfied rises. Patrick and Wolak (1996) have hypothesised that this may be an important channel of strategic bidding. The GA version of the pool presented here does not consider this opportunity.

$$
\begin{array}{ll}
p_m = \max\{p_i, p_j\} & \\
q_i \ = k_i & \text{if } p_i < p_j \\
q_i \ = Q - k_j & \text{if } p_i > p_j \\
q_i \ = Q/2 & \text{if } p_i = p_j
\end{array}
$$

This regime is a generalised version of the game considered in section 3.4. It describes the conditions necessary for the high bidder to be a residual monopolist[46].

*Regime 3:*

$$p_i < \min\{p_i, P_j\}, \text{ and } [(p_j > P) \text{ or } (p_j < P \text{ and } k_i > Q)]$$

$$p_m = p_i$$

$$q_i = \min\{k_j, Q\}$$

$$q_j = 0$$

This regime describes the case when bidder *i* is the low price bidder, and either *i* can satisfy the whole market, or *j* has bid "out of bounds" (above *P*). Market price is then *i*'s bid, and quantity called on to produce is the lower of his quantity bid and market demand.

We consider the game where price and availability bids are restricted to integers between 0 and 63, where $Q = 18$ and $P = 45$, $C = 30$, $c = 0$.

## The GA simulation

The GA was used to find which regime (if any) would be selected, and how producers divided up the market within that regime. Table 3.5 of Annexe A gives the parameter and representation details of the GA for this model. The addition of a strategic variable adds no technical complications to the GA itself – the bit strings are operated on just as before[47]. However, they are interpreted differently by the simulation model. In this case, the first six positions of each bit string was interpreted as an (integer) price between 0 and 63, whilst the next five were interpreted as quantities between 0 and 20 (in step increments of 20/31$^{\text{st}}$ 's).

Capacity availability bids were modelled as being made simultaneously with price bids, and were made at every generation. This might be thought to exclude from the commitment choice the essential element that irreversible capacity costs make post-commitment decisions a function of the pre-commitment decision even when this was ex post sub-optimal. It is

---

[46] The essential condition is the first stated, that both capacity bids together are greater than demand, whilst the low bidder alone has bid less than demand. The (clumsy-looking) added condition ($k_{[S\{p_i,p_j\}]} < Q$) just ensures that it is only the low price bidder not being able to satisfy demand that is important.

[47] Some changes could be envisaged within the GA. For example, during crossover, should one pick a single point on the "price-and-quantity" string, or should one treat the price and quantity strings as separate, and allow each offspring to be a recombination of her parents' price and her parents' quantity strategies? Both methods were tried but did not seem to affect outcomes in this game.

certainly true that this feature of commitment is interesting, and a fertile source of path dependent behaviour[48]. The model considered here is simpler. Interest in it can nevertheless be justified on two fronts. The first flows from an interpretation of the GA as representing the firm's hypothetical exercises before coming to a market decision[49] (see section 4). The second changes the interpretation of the simulation. The quantity bids could be seen as offers to make existing capacity available for generation at a given time[50], which itself incurs irreversible costs[51] – it can be seen as the cost of the option to participate in the market.

## The GA results

Figure 12 shows the average profits earned by the two producers over the course of the run. We find the high profit/low profit equilibrium emerging rapidly, with producer 1 at first maintaining the high profit position. However, there is considerable instability within the pattern up to the sixtieth generation. This is best seen in Fig. 13, which shows the modal (most frequent) value of the strategic variables in each generation for each player.

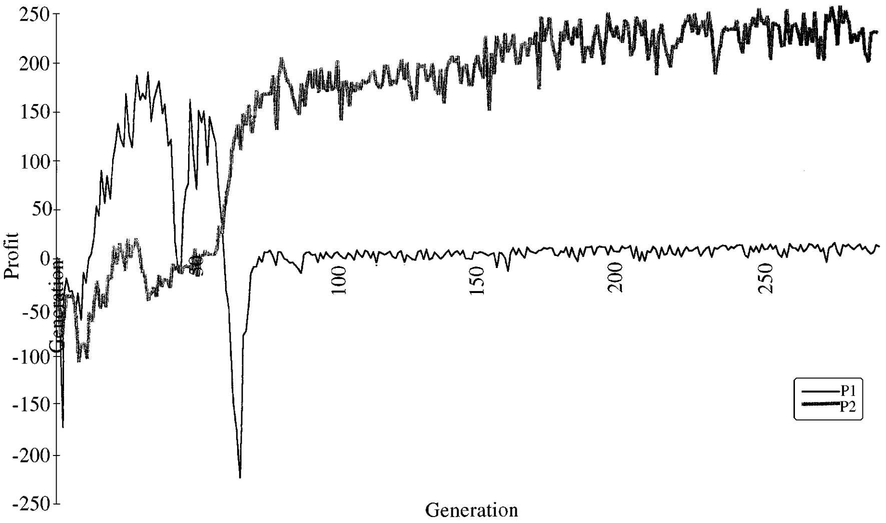


**Fig. 12.** Average profit in the pool game with capacity choice

---

48 Indeed, path dependency and irreversible commitment might even be seen as the same phenomenon viewed form different perspectives: the first looks back, and the second forward in time.

49 This is very similar to the justification for the standard "backward induction" method of solving multiple-stage games. See Kreps (1990), ch. 12.

50 In the English power pool, generators are required to declare availability in each of 48 half hourly price-setting periods one day ahead.

51 For example, a power station declared available needs to be fully manned whether or not it generates any electricity.

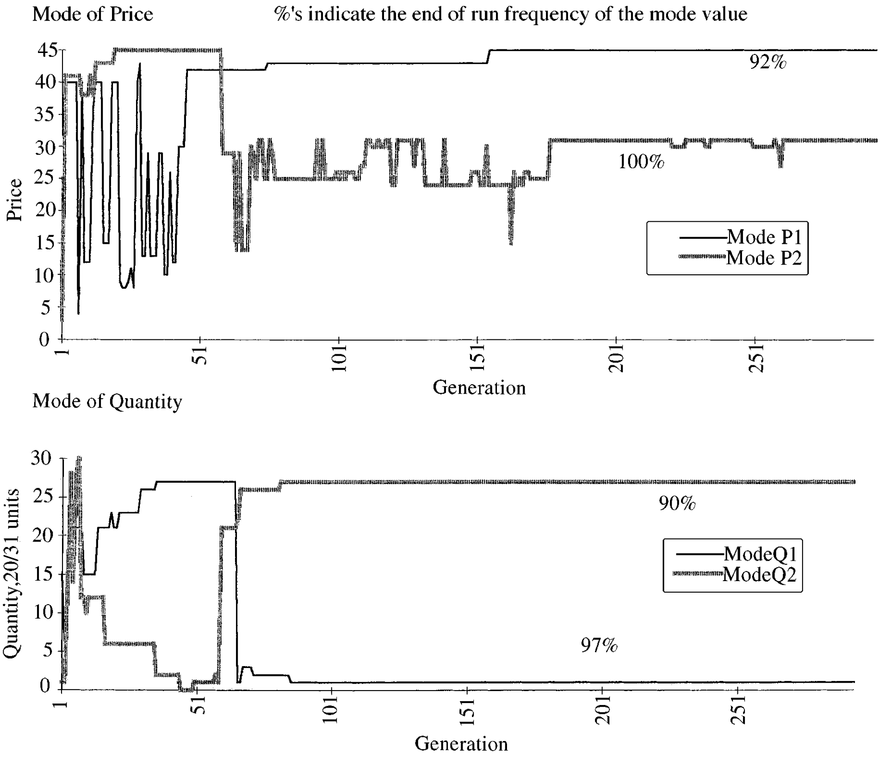


**Fig. 13.** Price and quantitity modes for each player

Up to the sixtieth generation, we see modal quantities (in the bottom half of the figure) separating into a clear "high" and "low" quantity pattern. However, the price strategies are less differentiated, with a high price often being modal for both players. The reason for the lack of early separation in price is that the frequency of the high price is low enough for both players to try to influence the probability of a high price outcome (actual frequencies in each generation are not shown in Fig. 13 except for the last generation).

After generation 60, a switch of positions is seen, and producer 2 emerges as the more profitable participant. By generation 300, producer 1's strategies are to price at 45 (92% frequency) and offer a quantity of a single 20/31$^{st}$ (i.e. 0.6 units, 97% frequency), whilst producer 2's strategies are to price below 31 (100% frequency) and offer a quantity of 27 20/31$^{st}$ 's (i.e. 17.5 units, 90% frequency). Thus, by the end of the run, the strategies can be characterised as one producer offering the lowest quantity possible and bidding it at the maximum price, and the other producer offering the highest residual quantity at a low price[52]. The equilibrium of the system

[52] There is no guarantee, of course, that another switch will not occur, as was seen in the simple pool game.

thus seems to yield residual monopoly (just as in the case in which capacity choice is exogenous). The firms evolve towards opting for capacity levels that avoid Bertrand outcomes[53].

## Analytical results

von der Fehr and Harbord (1993a) do not address the commitment problem in this way. They consider a much more realistic 2-stage capacity commitment problem: players can invest in one of a number of types of capacity (trading-off capital and operating costs), and the question they ask is whether different post-commitment regulatory structures (i.e. pricing rules) give rise to incentives that lead not only to the optimal quantity of investment, but also to the optimal mix of investment. The results of this model are not directly comparable, because the simple case considered above has no capacity-type decision (and no reason to prefer different capital and operating cost trade-offs).

However, they identify three ways in which the commitment decision is affected by the pool game.

1. A consumer price effect: restricting output increases price, and thus, ceteris paribus, there is under-investment[54];
2. A multiple-technology firm effect: this effect flows from the multiple technology types considered, and leads to over-investment.
3. A non-competitive spot price effect: if prices paid to capital commitment exceed average costs, there is a super-optimal incentive to invest.

In the simple game considered here, only effects 1 and 3 are at work. The GA converges to a state in which there is no over-commitment, suggesting that effect 3 is counterbalanced (at the equilibrium) by effect 1. Is this what an analytical treatment would suggest?

A simple thought experiment suggests that it is. Imagine that two producers are in regime 2, with $k_i + k_j = Q$ and $p_i = \min\{p_i, p_j\}$. An increase of $k_j$ by a single unit leads to lower profits for $j$, since, being the high bidder, the amount he is asked to produce $(Q - k_i)$ does not change (in fact, it is independent of his quantity bid), and yet he incurs the commitment cost $C$. Now imagine that $k_j$ is increased by a single unit. This does lead to $i$'s profits increasing, since each extra unit of commitment is actually called on to produce, and it earns supernormal profits $((p_m - c)$, which is greater than $(C + c))$. This, in fact, is effect 3. However, at this point, $k_i + k_j = Q + 1$ and $q_j = k_j - 1$. In other words, there is excess capacity, and all of it is held by $j$. $j$'s best response to this is to reduce commitment until $k_j = q_j$, at which point $k_i + k_j = Q$ once again (for the reasons given above). The analysis

[53] The same result is found analytically, in a slightly different game, in Kreps and Scheinkman (1993).

[54] This is the essence of the Kreps and Scheinkman (1963) result.

thus shows that effect 1 dominates over effect 3, so that the outcome follows the classic Kreps and Scheinkman (1983) result. Moreover this incentive remains *as long as j remains the marginal supplier*, which implies $q_i = Q - 1$[55] – the non-marginal supplier likes to increase quantity sold, as long as this does not lead him to being the sole supplier.

Thus, the logic of the game is that the low bidder would want to increase capacity commitments up to the point at which an increment of capacity would make him the market supplier at the low price. This is precisely the outcome we observe in the GA simulation, where the high price bidder finds himself "squeezed" into supplying the minimum non-zero quantity of 20/31 units.

## 4 Discussion of the results

Are the results presented in section 3 interesting because they are created by a particular algorithm, or (simply) because they offer an advance on widely used techniques of modelling markets? In this section, I tentatively propose a reason why the first might be the case, and argue that in some circumstances the second is the case. If the results of EP simulations are to be interesting in themselves, an analogue for the processes embodied by the algorithm must be found. If EP simulations are to be just useful tools, they (only) need to better solve certain problems than widely used techniques.

### *4.1 What evolves?*

The central question we need to answer if we are to treat EP simulations as worthy of interest *per se* is whether firms (or people) can be thought of as actually following a similar process in their decision-making. Axelrod (1984), justifying his use of an evolutionary process in modelling human decision-making, writes:

"The motivation [...] is to discover which kinds of strategies can be maintained by a group in the face of any possible alternative strategy. If a successful alternative strategy exists, it may be found by the "mutant" individual, through conscious deliberation, or through trial and error, or through just plain luck. If everyone is using a given strategy and some other strategy can do better in the environment of the current population, then someone *is sure to find this better strategy sooner or later...*" (p. 57, emphasis added).

However, in the games examined above, it is not really reasonable to assume (as Axelrod seems to do) that the equilibrium outcomes have

[55] David Harbord (pers. comm.) has pointed out that the result presented here would not generally hold if bidders could offer different units of capacity at different prices. For example, the low price, high quantity bidder, when faced with a price bid of $p_m$ and a quantity bid of 1, will prefer to bid a last unit of capacity at $p_m$ as along as $C < 1/2\, p_m$. Thus, the actual level of capacity or availability costs, of the "price ceiling", and of any indivisibilities will be important in the more general setting.

arisen from firms actually trying out, in the market, all the strategies that precede equilibrium. This would require an extraordinary lack of rational deliberation[56]. It is therefore important to consider the types of interpretation that can be applied to the evolving strategies. What exactly is it evolves[57]?

To answer this fully would require a detailed model of decision-making within the firm[58]. But even without a fully spelt out model, one candidate for the rȯle is that the GA represents the evolution of strategic scenarios within the firm. Scenario analysis is pervasive in the corporate world. Some scenarios are explicitly worked out on firms' own quantitative models of their industry, whilst others are more qualitatively generated and analysed. In either case, a firm's scenarios embody its deliberations about its environments. Under this interpretation, the GA simulations can be considered to model the decision process within one firm as a series of *hypothetical* calculations. The chromosomes of the GA represent scenarios, and the gradual spread of strategies within each firm's population is intended to mimic the way firms weed out unviable scenarios and keep those that "have a chance"[59]. Convergence of a population to a single strategy (or to stable

[56] Mirrored nicely by the often incredible rational sophistication required by analytic treatments of Nash equilibrium.

[57] "What Evolves?" is a question rightly highlighted by Anderson (1994) as central to the development of a fully evolutionary theory of the firm.

[58] Chattoe (1995) criticises some recent GA applications in economics on the grounds that they often confuse an "instrumentalist" and a "descriptive" rȯle for the GA. In the first, the GA is used as it is in engineering applications, to efficiently search a solution space; in the second, an analogy is implied between the operations of the GA and decision making within the firm, so that GA results become interesting in *themselves*. Chattoe argues that when the first parades as the second, no progress in understanding is made, because no explanation is given of why real firms embody precisely those searching strategies that are implicit in the operations and parameters of the GA. It is true that there is often confusing silence on the questions of "What evolves? And why does it evolve *like this*?", and true also that progress can undoubtedly be made on making the GA more "representational". These need not, however, be reasons to concentrate effort solely on these questions. Not only can be GA have an "instrumental" rȯle in economics (see section 4.2), but there also exists research suggesting that human agents in experimental settings are well mimicked by EPs (e.g. Arifovic (1996), Chen *et al.* (1996)) – experimental validation may allow us to postpone Chattoe's challenge.

[59] Penrose (1990) offers a strikingly similar model of how mathematicians arrive at mathematical truth. He writes (p.546): "There must be a powerful impressive selection process that allows the conscious mind to be disturbed only be ideas that 'have a chance' [...] It seems to me there are two factors involved, namely a 'putting-up' and a 'shooting-down' process [...] Without an effective putting up process, one would have no new ideas at all [...] But one also needs an effective procedure for forming judgements, so that only those ideas with a reasonable chance of success will survive." Similarily, the firm would only actually try out strategies in the market if they, too, "have a chance," and internally, the firm needs a putting-up process which we model by mutation and recombination.

distribution of strategies) can be thought of as the firm achieving some type of "reflexive equilibrium"[60].

As noted in the description of the Pool game, the appearance of Nash equilibria is sensitive to the crossover rate. At a crossover rate below approximately 12% of the population, the strategy variable distributions seem to cycle in suggestive, but out-of-equilibrium ways. Above about 20% of the population, the GA often locks-in to non-Nash equilibria, and population diversity falls very rapidly.

Such sub-optimal results of the GA simulation are both a challenge for the approach and a potential strength. Arthur (1993) describes a learning algorithm in which lock-in is the result of difficulty of discrimination between good outcomes. He argues that "What is crucial to the emergence of optimal action is [...] that learning has time to explore and discover the action with the largest expected value." Moreover, he finds that human subjects in experimental settings typically do not allow exploration enough time.

The "time to explore" explanation of lock-in can be combined with the scenario interpretation of the GA to shed light on the interpretation to give of the GA's sensitivity to reproduction and crossover rates. High values for these (above about 20% reproduction rate) can be thought of as corporate "single-mindedness", whilst low values (below about 12%) would represent "indecisiveness". A high reproduction and crossover rate, under this interpretation, sees the firm selecting initially promising scenarios on very little evidence of their profitability: good performance in just a few successive generations can lock a firm into these scenarios. At low reproduction and crossover rates, the vagaries of tournament selection may lead no scenario to ever dominate the population[61]. It is somewhere between single-mindedness and indecisiveness that the firm can hope to remain sufficiently open-minded to attain optimal outcomes[62].

### *4.2 Using the GA in applied simulation*

Markets occasionally have to be modelled in greater detail than offered by the stylised abstractions of typical industrial organisation models. This can

[60] "Reflexive equilibrium" need not be a Nash equilibrium. Moreover, this interpretation leaves open the question of whether the firm need delay action in the market until reflexive equilibrium is reached. If the GA is modelling the way firms develop internal models of their environments, then there will be cases in which a market trial yields sufficiently valuable information about the modelled market to be worth risking.

[61] What determines a "high" and a "low" reproduction and crossover rate is almost certainly game-specific. This seems to be a case where the suggestion in Chattoe (1995) that these parameters should be endogenised in the EP simulation could yield interesting insights.

[62] This interpretation sits nicely with management theorists' current emphasis of the firm as a "learning machine."

happen in prescriptive applications, where a firm or regulator is being counselled. For example, von der Fehr and Harbord (1996) consider whether a number of market structures in the Australian electricity industry are likely to lead to abusive dominant positions. The level of detail required here goes beyond the stylised representation of their analytical model, because the regulator is concerned with the actual extent of any likely effect, and not just its direction. Moreover, the choice variables include no longer just price and quantity, but timing, discrete and multiple capacity tranches, entry, transmission capacity et cetera[63].

The method employed by von der Fehr and Harbord[64] is to use a simulation model to calculate pay-offs under different combinations of choice variables, and examine results off-model to find the patterns predicted by the analytic model. The combinatorial explosion that soon occurs when trying to richly simulate markets imposes this restricted scenario analysis: only so many combinations of price, entry and investment options can be run. The solution taken by London Economics and Harbord Associates (1995) is to fix exogenously entry decisions and transmission investments, and then search for Nash equilibria in prices under each scenario. Price choices are themselves constrained to being 1, 2 or 3 times marginal cost.

The approach can be quite successful. In this case, for example, it was shown that one player maintained a dominant position across most scenarios. However, the approach also has its drawbacks. Decisions to enter a market will depend on the type of competition expected in that market, so fixing entry exogenously weakens conclusions. Markets with large sunk costs will often exhibit significant path dependency, but the scenario analysis cannot demonstrate the extent or importance of such effects. Finally, constructing pay-off matrices for all these scenarios is inefficient: the attempt is to "cover" solution space as widely as possible, where an efficient simulation will concentrate effort on promising areas of the solution space.

The EP approach promises to resolve such difficulties. Firstly, the solution method is not constrained by the complexity of the objective function. There is no technical difficulty for EP involved in endogenising

---

[63] A number of other difficulties might arise in trying to apply the von der Fehr and Harbord model to real power pools, which could be explored in an EP framework. For example, Patrick and Wolak (1996), while recognising the value of the von der Fehr and Harbord framework, hypothesise that the two dominant players in the England and Wales market use quantity commitments rather than prices as strategic variables in part because the market rules stipulate that each price bid must cover 48 price-setting periods, whilst quantity bids may be tailored to every price-setting period. Quantity bids are thus attractive strategic variables because they can be more flexibly used. Work in progress is using the GA to test this hypothesis and its likely magnitude (see also note 44).

[64] The details of the actual runs are available in London Economics and Harbord Associates (1995).

variables like investment timing or entry, whereas analytic models soon find the problems intractable, and traditional scenario analysis is ill-suited to searches in highly dimensioned spaces. Secondly, an EP simulation will explore path dependent outcomes and other such "near equilibrium" outcomes. It does this because the mutation and crossover operators can be thought of as "trembles", and significant path dependency occurs when the impact of trembles is large. Finally, the EP approach is computationally more efficient than the "pay-off calculator" approach of scenario analysis: solution space exploration is concentrated on "promising" areas, and is not pre-imposed by the modeller.

## 5 Conclusion

This paper has shown how EP can be used to search for equilibria in simple, standard games from industrial organisation theory. The technique has performed well in this setting. It has also been suggested that one promising area for EP in economics is to supplant the usual scenario analysis used by market analysts, and that EP would help to make applied conclusions seem less arbitrary[65].

However, as a move is made to representing more complex choice problems, the simple GA structure described above will soon become insufficient. For example, it is not capable of representing conditional choice (which is why all the games explored were one shot, simultaneous move games). Several techniques exist which can overcome this limitation. One is Holland's Classifier Systems, another Koza's genetic programming[66]. This latter seems particularly promising. The genetic programming method explores the space of possible programs addressing a problem, and finds analogues of the recombination and mutation operators of the GA. Firms (and their regulators) already use computer programs as aids to decision, so evolving computer programs to represent firm's, behaviour and decision making processes has very clear analogue[67].

---

[65] Lane (1993) argues that "The more richly detailed a model is, the intriguing it is to its designers – but the less likely it is to capture anyone else's imaginations or interest." This is all the more the case when important variables are chosen by the modeller rather than the model.

[66] Holland and Miller (1991), Koza (1992). Lane (1993) argues that the type of GA used here can actually be thought of as the simplest possible Classifier System.

[67] Genetic Programming can be thought of as a way of endogenising Nelson and Winters' (1982) "routines". See Anderson (1994).

## References


Anderson ES (1994) Evolutionary Economics – Post-Schumpeterian contributions Pinter, London

Arifovic J (1996) Strategic Uncertainty and the Genetic Algorithm Adaptation, Simon Fraser University Department of Economics Working Paper

Arifovic J (1996a) The Behavior of the Exchange Rate in the Genetic Algorithm and Experimental Economies. Journal of Political Economy 104 (3): 510–541

Arthur WB (1993) On designing economic agents that behave like human agents. Journal of Evolutionary Economics 3: 1–22

Axelrod R (1984) The Evolution of Cooperation. Basic, New York

Back T, Fogel D, Michalewicz Z (1995) Handbook of Evolutionary Computation. OUP/IOP Publishing, New York

Binmore K (1992) Fun and Games – A text on game theory. Heath, Lexington, MA

Chattoe E (1995) Modelling Economic Interaction Using a Genetic Algorithm. In: Back T Fogel D, Michalewicz Z (eds) Handbook of Evolutionary Computation, pp G7 1:1–1:6. OUP/IOP Publishing, New York

Chen S-H, Duffy J, Yeh C-H (1996) Equilibrium Selection via Adaptation: Using Genetic Programming to Model Learning in a Coordination Game. Evolutionary Programming Conference, San Diego, CA (1996)

Davis L (ed) (1991) Handbook of Genetic Algorithms. Van Norstrand Reinhold, New York

Dosi G, Nelson RR (1994) An introduction to evolutionary theories in economics. Journal of Evolutionary Economics 4: 153–172

Fehr von der NH, Harbord D (1993) Oligopoly, Capacity Competition and Uncertainty Under Different Regulatory Regimes. University of Oslo, Department of Economics, Working Paper

Fehr von der NH, Harbord, D (1993a) Spot Market Competition in the UK Electricity Industry. The Economic Journal 103: 531–546

Fehr von der NH, Harbord D (1996) Market Structure and Competition in Electricity Markets: The Australian Debate. Unpublished

Fudenberg D, Kreps D (1993) Learning Mixed Equilibria. Games and Economic Behavior 5: 320–367

Fudenberg D, Levine DK (1996) Learning and Evolution in Games. *http://levine.sscnet.ucla.edu/Papers/Essay/ESSAY7.htm*

Fudenberg D, Levine DK (1997) Theory of Learning in Games. Draft, *http://levine.sscnet.ucla.edu/papers/contents.htm*

Goldberg DE (1988) Genetic Algorithms in Search, Optimisation and Machine Learning. Addison Wesley, Reading, MA

Green R, Newbery DM (1992) Competition in the British Electricity Spot Market. Journal of Political Economy 100: 929–953

Holland JH (1975) Adaptation in Natural and Artificial Systems. University of Michigan Press, Michigan

Holland JH, Miller JH (1991) Artificial Adaptive Agents in Economic Theory. American Economic Association – Proceedings of the 103 rd Annual Meeting. pp. 365–370

Holland JH (1992) Adaptation in Natural and Artificial Systems, 2nd edn. MIT Press, Cambridge, MA

Kane D (1996) Hayekian Uncertainty, a Formal Model. Working Paper, Harvard University School of Government

Koza JR (1992) Genetic Programming. On the programming of computers by means of natural selection. MIT Press, Cambridge, MA

Kreps DM (1990) A Course in Microeconomic Theory. Harvester, Wheatsheaf

Kreps DM, Scheinkman J (1983) Quantity Precommitment and Bertrand competition Yield Cournot Outcomes. Bell Journal of Economics 14: 326–337

Lane DA (1993) Artificial Worlds and Economics, Part II. Journal of Evolutionary Economics 3: 177–197

London Economics, Harbord Associates (1995) Market Power in the Australian Electricity Market. A Report to the Industry Commision, London Economics Australia, Melbourne

Lucas N, Taylor P (1993) Characterising Generator Bidding Behaviour: Bidding Strategies in the pool. Utilities Policy, April, pp. 129–135

Marimon R, McGrattan E, Sargent T (1990) Money as a means of exchange in an economy with artificially intelligent agents. Journal of Economic Dynamics and Control 14: 329–73

Marks RE (1992) Breeding Hybrid Strategies: Optimal Behaviour for Oligopolists. Journal of Evolutionary Economics 2: 17–38.

Mitchell M (1996) An Introduction to Genetic Algorithm. MIT Press, Cambridge, MA

Nelson PR, Winter SG (1982) An Evolutionary Theory of the Firm. Belknap, Harvard

Newbery DM (1995) Power Markets and Market Power. The Energy Journal 16(3): 39–66

Noe TH, Pi L (1995) Genetic Algorithms, Learning and the Dynamics of Corporate Takeovers. Georgia State University Department of Finance Working Paper

Patrick HP, Wolak FA (1996) The Impact of Market Rules and Market Structure on the Price Determination Process in the England and Wales Electricity Market. Stanford University Center form Economic Policy Research, Discussion Paper

Penrose R (1990) The Emperor's New Mind. Oxford University Press, New York

Selten R (1983) Evolutionary stability in extensive two person games. Mathematical Social Sciences 5: 269–363

Tirole J (1988) The Theory of Industrial Organisation. MIT Press, Cambridge, MA

## Annexe A

This appendix provides the details of model parameter settings used in the runs. The table numbers refer to the section in the text where the models are described. The first table provides a brief description of the meaning of the entries. The last Table (A.4) gives, for each simulation, the number of "similar"[1] runs repeated. The values shown in the tables are for the runs described in the text.

**Table A.1**

| | |
|---|---|
| Description | Identifies the simulation being described. |
| Section in paper | Where the run is described in the text. |
| Number of populations | The number of co-evolving populations of strategies. |
| Strategies in population | The number of strategies in each population |
| Tournaments per generation | The number of "market trials" that make up each generation. |
| Chromosome coding | The "alphabet" used to define strategies |
| Chromosome mapping | The interpretation placed on the "alphabet". |
| Selection method for tournament | The method used for selecting strategies for tournaments. "Uniform Random" was used throughout, by which every strategy has an equal chance of being picked in every tournament. |
| Selection method for breeding | The way strategies re-combine to create new strategies. The "crossover rate" in the text refers to the proportion of strings in each generation involved in producing new strings for the next generation. |
| Crossover operator | The way re-combination occurs. |
| Mutation operator | The way random changes can influence the strategies in the next generation. |

**Table A.2.1**

| | |
|---|---|
| Description | Simple monopolist |
| Section in paper | 2.1 |
| Number of populations | 1 |
| Strategies in population | 100 |
| Tournaments per generation | 100 |
| Chromosome coding | 6 binary digit |
| Chromosome mapping | Integers from 0 to 63 represent prices |
| Selection method for tournament | Uniform Random |
| Selection method for breeding | 16% crossover rate (i.e. fittest 16 strategies survive, breaed, have one offspring, who replacing the bottom 16%). |
| Crossover operator | 2 Randomly matched parents, with random position on gene |
| Mutation operator | 5% probability that any offspring bit will be "flipped" |

[1] Very few runs were repeated identically – parameters like breeding rate, mutation or tournament per generation were often modified to test sensitivity. Details of the simulation were also occasionally modified (what to do, for example, in the case of a price tie in the Bertrand of Pool games). When sensitivity was large enough to affect the structure of outcomes, runs have not been counted as "similar".

**Table A.3.1**

| | |
|---|---|
| Description | Bertrand duopoly |
| Case | 3.1 |
| Number of populations | 2 |
| Strategies in population | 100 |
| Tournaments per generation | 100 |
| Chromosome coding | 6 binary digit |
| Chromosome mapping | Integers from 0 to 63 represent prices |
| Selection method for tournament | Uniform Random |
| Selection method for breeding | 16% crossover rate |
| Crossover operator | 2 Randomly matched parents, with random position on gene |
| Mutation operator | 5% probability that any offspring bit will be "flipped" |

**Table A.3.2**

| | |
|---|---|
| Description | Cournot duopoly |
| Case | 3.2 |
| Number of populations | 2 |
| Strategies in population | 100 |
| Tournaments per generation | 100 |
| Chromosome coding | 5 binary digit |
| Chromosome mapping | Integers from 0 to 31 represent quantities |
| Selection method for tournament | Uniform random |
| Selection method for breeding | 16% crossover rate |
| Crossover operator | 2 randomly matched parents, with random position on gene |
| Mutation operator | 5% probability that any offspring bit will be "flipped" |

**Table A.3.3**

| | |
|---|---|
| Description | Serial monopolists |
| Case | 3.3 |
| Number of populations | 2 |
| Strategies in population | 100 |
| Tournaments per generation | 100 |
| Chromosome coding | 6 binary digit |
| Chromosome mapping | Integers from 0 to 63 represent prices |
| Selection method for tournament | Uniform random |
| Selection method for breeding | 16% crossover rate |
| Crossover operator | 2 Randomly matched parents, with random position on gene |
| Mutation operator | 5% probability that any offspring bit will be "flipped" |

**Table A.3.4**

| | |
|---|---|
| Description | Electricity pool |
| Case | 3.4 |
| Number of populations | 2 |
| Strategies in population | 100 |
| Tournaments per generation | 100 |
| Chromosome coding | 6 binary digit |
| Chromosome mapping | Integers from 0 to 63 represent prices |
| Selection method for tournament | Uniform random |
| Selection method for breeding | 18% crossover rate |
| Crossover operator | 2 randomly matched parents, with random position on gene |
| Mutation operator | 5% probability that any offspring bit will be “flipped” |

**Table A.3.5**

| | |
|---|---|
| Description | Electricity pool with endogenised capacity |
| Case | 3.5 |
| Number of populations | 2 |
| Strategies in population | 100 |
| Tournaments per generation | 50 |
| Chromosome coding | 6 binary digit + 5 bindary digits |
| Chromosome mapping | First 6 bits represent integers from 0 to 63 (prices), next 5 bits linearly mapped from 0 to 20 (quantities) – i.e. step increments of 20/31. |
| Selection method for tournament | Uniform Random |
| Selection method for breeding | 16% crossover rate |
| Crossover operator | 2 randomly matched parents, with random position on gene (i.e. from 11 positions) |
| Mutation operator | 10% probability that any offspring bit will be “flipped” |

**Table A.4**

| Model | Number of “similar runs” |
|---|---|
| Simple monopolist | 3 |
| Bertrand duopoly | 3 |
| Cournot duopoly | 3 |
| Serial monopolists | 3 |
| Electricity pool | 10 |
| Electricity pool with endogenised capacity | 14 |