

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

Introduction

Evolutionary programming (EP)¹ techniques have been recognised by a number of economists² as being of potential use to the discipline. These researchers have tended to apply EP methods to relatively complex problems that require advanced tools even for analytical solution³. In this paper, I describe the results of applying a simple genetic algorithm to a number of much 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 more 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.

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 ways⁴. This potential was recognised by the precursors of EP in economics⁵, but more recently it has been overlooked in favour of a more theoretical view of EP's rôle⁶. 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, or the data available to them is incredible).

Description of the model

GAs are a special sort of optimisation algorithm. All optimisation 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 optimal⁷. 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

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

² An overview is presented in Lane (1994)

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 very successful application of GA simulation is in the study of repeated prisoner's dilemmas. See Marks (1992).

⁴ A good example of this sort of analysis is described in Harbord and von der Fehr (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).

⁵ Nelson and Winter (1982) argue that the way we get computers to solve complex problems is interesting because they can teach us about the way humans solve these problems. See also Anderson (1994).

⁶ 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.

⁷ A thorough description of GAs is given in Holland (1975). A brief overviews can be found in Koza (1992) and Marks (1992).

specified number of generations, or once some other externally defined criterion is satisfied. Thus, GA's are a method of optimisation 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.

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 model⁸? The 6 steps involved are:

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 representation⁹ is used here. We limit the monopolist's search domain to prices between 0 and 31, which allows us to represent all possible strategies as a five digit binary number. So, for example, a price of 31 is represented by 11111, and a price of 1 by 00001. 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). 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 tournaments¹⁰ 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" strategies¹¹ (i.e. prices yielding most profit) to breed and allow the least fit to disappear from the population. The breeding method used

⁸ 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).

⁹ 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 codings are given in Davis (1991).

¹⁰ 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 permutations of strategies becomes large.

¹¹ 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.

here resembles genetic recombination: two parent strategies are chosen, a crossover position 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 (00001) and 31 (11111) were chosen as parents, with a crossover position of 3, the two offspring prices would be 3 (00011) and 29 (11101). It is also possible in this step to include a mutation operation which randomly switches the value of a digit with some probability.

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.

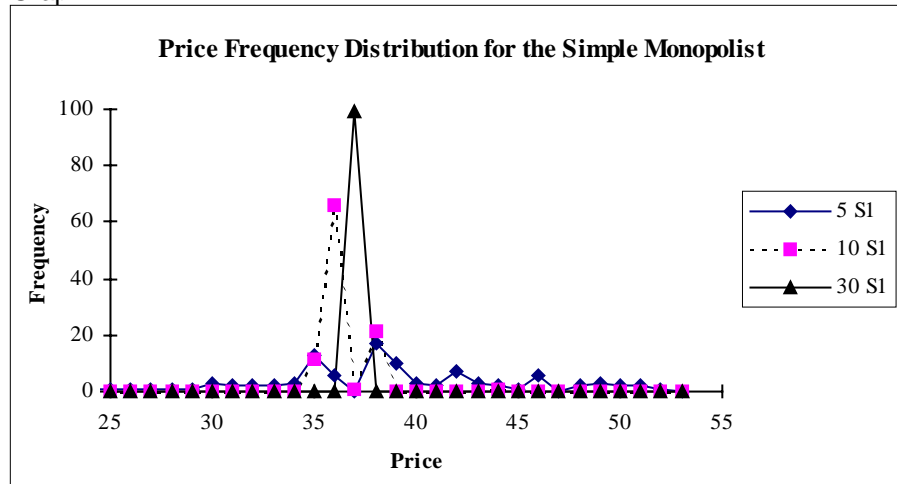
The search method described in these steps probably does not at first sight seem to have any obvious economic interpretation¹² (what, for example, is the analogue of strategy recombination within the firm's decision making process?), and the reader could justly wonder at this stage why the GA should 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 naïve (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.

Graph 1 below is a representation of the state of the population of strategies after 5, 10 and 15 generations¹³. The population is comprised of 100 possible prices 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.

¹² Indeed, it is not even obvious that the steps constitute a good optimisation procedure. Holland (1975) shows that the GA is a good all-purpose optimiser. 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 intuitive overview of the proof is given in Koza (1992)

¹³ 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 (see below) require an analysis of the population as a whole. Schematically, we view the population as the reservoir of ideas from which the firm (eventually) comes to randomly pick its actual market strategy.

Graph 1



There are several things to note about the graph. Firstly, the GA this was run on allowed only integer solutions, so that the 30th 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 an optimal solution very quickly. In fact, even by the 5th generation, we see a remarkable amount of “organisation” in the population of strategies¹⁴. The 10th 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 below) 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.

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)¹⁵.

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”.

¹⁴ Remember that the starting population is randomly generated, and therefore approximately uniform.

¹⁵ 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 Harbord and von der Fehr (1993).

However, there is a considerable increase in the complexity of the mechanisms at work. Each player has a separate population of strategies. These populations now co-evolve, since 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 ...).

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 that 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 intuitive¹⁶: each producer wants to price just 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 five binary digits.
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.
5. Birth, death and mutation was carried out for each population separately, as described above.
6. The simulation was stopped when an equilibrium seemed to be reached.

Graph 2 below shows the evolution of market price over the 300 generations the GA ran for. By the end of the run (in fact, by the 150th generation), the market price in tournaments had fallen to 1. This is almost equivalent to the analytical result, in which price falls to 0. The discrepancy arises from the fact 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.

Graph 2

¹⁶ This is described more formally in Tirole (1989), pp210-211.




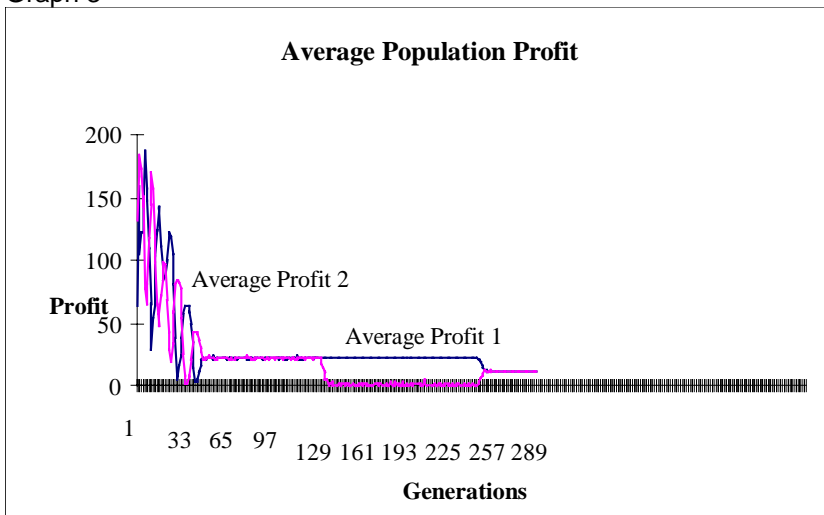
Table 1 below shows the frequency with which prices are encountered in each population by the end of the run. Almost all surviving price are at $p=1$, and the one that is not, $p=0$, we can safely assume has arisen out of mutation and will not survive (since if both set price at 1, they share the market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

Table 1

Price	Frequency in Population 1, %	Frequency in Population 2, %
0	1	0
1	99	100

Graph 3 below 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).

Graph 3



Cournot Competition

The Cournot Model

Cournot competition arises in the following sort of setting¹⁷:

¹⁷ This is taken from Kreps (1990) p443.

“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 reduce¹⁸.

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 π_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 interpreted as being quantities, and the payoffs in the market simulation are as described above.


Graphs 4 and 5 show the frequency distribution of prices in the two populations at the end of the run (289 generations). The analytic solution, $q_1 = q_2 = 9$, accounts for 92.5% of all the strategies. Thus, although the system is clearly attracted to the Nash equilibrium, a small number of other strategies also survive (unlike the Bertrand outcome). In fact, there is very little change in the distributions from the 80th to the 289th generation, with the Nash equilibrium accounting for 87.5% of strategies in the former.

Graph 4




¹⁸ Quantity is said to be a strategic substitute. See Tirole (1989) pp 218-220.

Graph 5



This persistence of non-Nash strategies is reflected in the market outcome prices, shown in graph 7. Unlike the development of prices in the Bertrand simulation, which soon becomes very stable (perturbed only marginally by random mutations), the Cournot price path shows considerable variation about the Nash equilibrium price. The variations occur because there is always some chance that an unusually large number of non-Nash strategies will be pitted in the market tournaments, yielding higher or lower quantities offered to market.

Graph 7



The Chain of Monopolies

The Chain of Monopolies model

In the chain of monopolies model¹⁹, 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 because it provides the simplest demonstration of the slightly paradoxical fact that a chain of monopolists supply a smaller price at a higher quantity than an integrated monopolist would.

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.

¹⁹ The vertical chain of monopolists is described in Tirole (1989), p174.


- 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 $dp_1/dp_1 = 0$ and $dp_2/dp_2 = 0$, which occurs at $p_1 = 19 \frac{1}{3}$ and $p_2 = 26 \frac{2}{3}$.

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.

Graph 8 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 distributions. Thus, the system seems to behave similarly to the Cournot case in this game, with the Nash equilibrium a strong attractor, but persistence of a few non-Nash strategies. 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.

Graph 8



A Simple Electricity Pool

The model used here is a simple version of that developed by Harbord and von der Fehr (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 being that price required to pay 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.

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 enough²⁰. 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.

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:

²⁰ This is only one of the several cases considered by Harbord and von der Fehr (1993).

- 10 p2 if p2 > p1
- 8 p1 if p1 > p2
- 9p1 if p1 = p2
- 10 p1 if p2 > 45 > p1
- 0 if p1 > 45
- $\Pi_2(p_1, p_2) = \Pi_1(p_2, p_1)$


This pay-off function simply says 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 last three conditions should be evident.

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.

The mixed strategy equilibrium has no very intuitive description. The equilibrium mixing probabilities (allowing discrete integer price bids between 0 and 45) are shown in Graph 8.

Graph 8
P(Price)

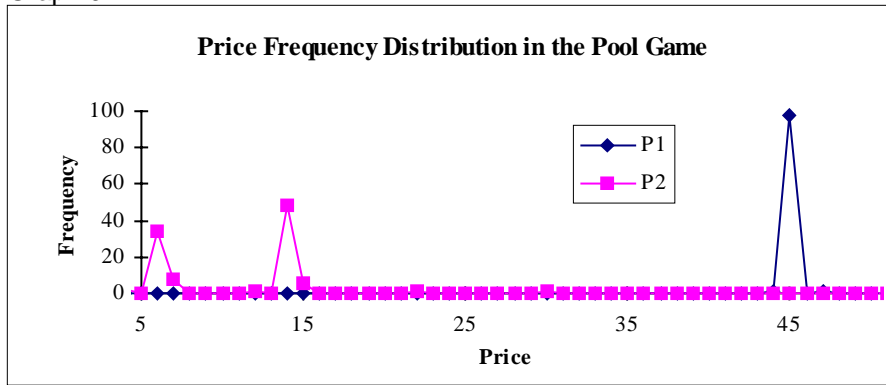


The GA simulation

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


Graph 9 shows the distribution of prices in the two populations after 85 generations. This corresponds exactly to one of the pure strategy Nash equilibria, with producer 1's strategy population consisting entirely of the maximum price bid (45), and producer 2 bidding indifferently anywhere below 36 (the grouping at prices of 14 and 6 is accidental, although groupings at arbitrary values below 36 occurs whenever this model is simulated).

Graph 9



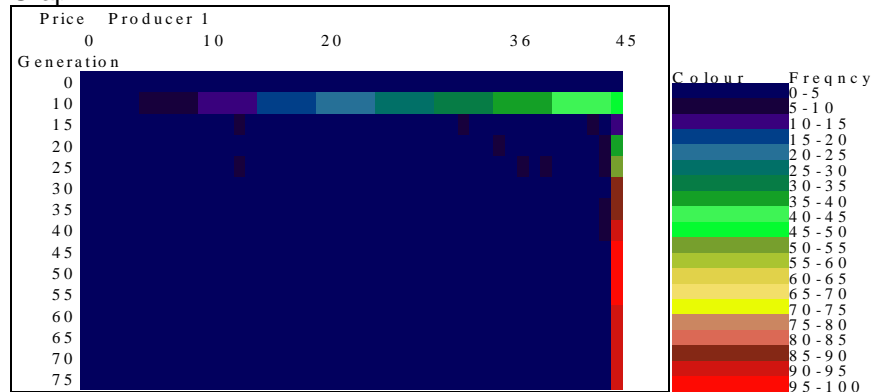
Which producer ends up with the higher profits (lower bids) is entirely a matter of chance. To see in more detail how this occurs, Graph 10 shows the evolution of profit streams for the two producers over the length of the run. The high/low profit outcome appears soon after the tenth generation, and stabilises soon after the fiftieth.

Graph 10

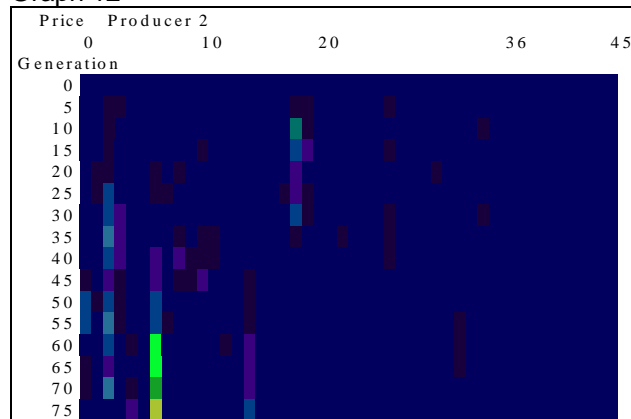


To see how the populations develop to produce the outcome, we need to look at the evolution of the frequency distributions over the course of the run. This is shown in Graphs 11 and 12, which map the frequency of price occurrences in the populations every five generations, from the start to generation 75. The zero'th generation is an entirely homogenous shade, showing the near uniform distribution of prices. We see that already by the 10th generation, producer 1's strategies are bunching around high values, whilst producer two's are still relatively undifferentiated. From the evolution of profit figure (Graph 10), we see that neither producer is obviously more profitable than the other at this point. In fact, producer 1 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 benefit accrues to the other.

Graph 11.



Graph 12



The GA was run with a number of different parameter settings (mutation rates, size of reproduction pool), to examine whether any combination of settings would yield the mixed strategy equilibria as stable fixed points. Although some parameter settings destroyed the pure strategy equilibrium described above (see discussion for more detail), none produced an outcome that was recognisably the mixed strategy equilibrium²¹. This suggests that the mixed strategy equilibrium is simply too knife-edged to be the outcome of a heuristic process²².

The original Harbord / von der Fehr 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. The simulation described above was modified to state that if both producers 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.

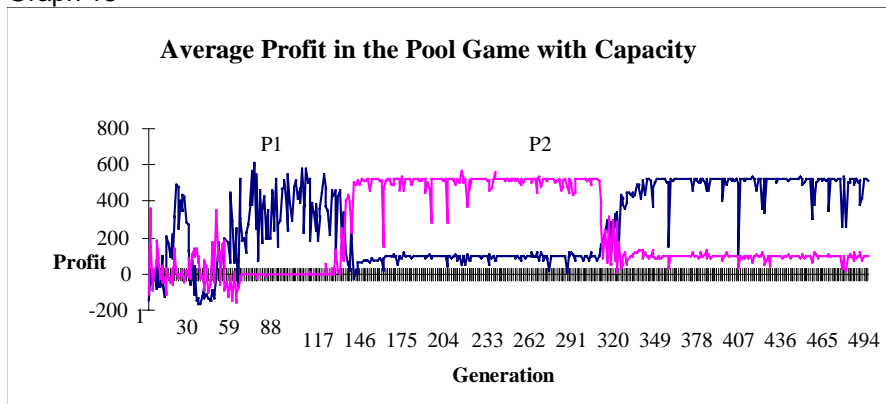
Graph 13 shows the average profits earned by the two producers over the course of the run. We find the high profit/low profit equilibrium again (from the 60th generation). However, what is

²¹ These runs produced almost cycling distributions of prices, and it is difficult to test whether the mixed strategy distribution is any sort of attractor of these distributions.

²² This conclusion sits slightly uneasily with Maynard Smith's (1974) pioneering work in evolutionary game theory. In his original article, he argues that in animal displays of strength, persistence time will be determined by a randomising strategy. The GA simulations suggest that it is on the whole difficult for an evolutionary process to result in a mixed strategy.

striking is that the producers occasionally switch positions (around the 140th and 320th generations).

Graph 13

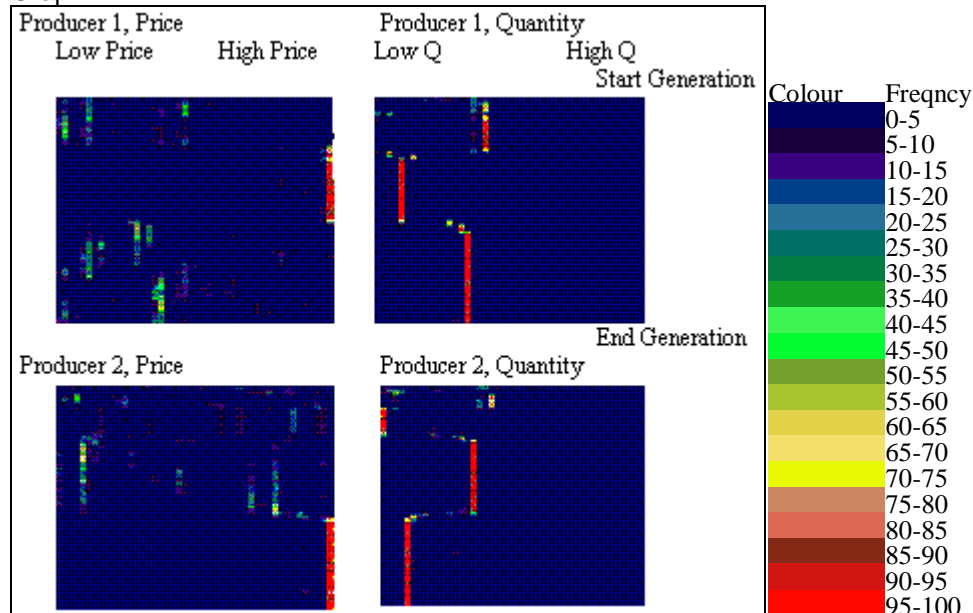


These rapid switches can be examined in more detail in Graph 14, which describe the state of the strategy populations over the course of the run. The top left box has producer 1's price, the top right box his quantity, and the second row contains the same information for producer 2. We see that producer 1 initially settles into a strategy of offering a high quantity at a low price. The quantity is high enough to drive producer 2 out of the market altogether (thus avoiding incurring the sunk capital cost) - in Graph 13, we see this by producer 2's zero average profit between generations 60 and 120, and in graph 14, we see this by the very high frequency of zero quantities in this portion of the run. However, producer 1 soon starts to reduce quantity and increase price, which opens a niche for producer 2 to increase quantity (in the second quarter or so of each row in Graph 14, we see producer 1's high price strategies increasing in frequency, and his low quantity strategies falling, whilst the high quantity strategies of producer 2 become more frequent. This is the first switch seen in Graph 13, around the 120th generation. However, in the South West quadrant of Graph 14, we also see that producer 2 is not (on average) pricing sufficiently low to avoid being undercut by producer 1, which is precisely what eventually happens: producer 1 increases quantity and reduces price. This is the second switch we see around generation 320 in Graph 13. At this point, producer 1 is pricing sufficiently low to avoid another switch of fortunes.

The final equilibrium of the system is a residual monopolist pure strategy Nash equilibrium, just as in the case in which capacity choice is exogenous. The firms thus evolve towards opting for capacity levels that avoid Bertrand outcomes²³.

²³ The same result is found analytically, in a slightly different game, in Kreps and Scheinkman (1983).

Graph 14



Discussion of the results

In all the simulations described above, the GA comes recognisably close to selecting a Nash equilibrium. However, in all but the Bertrand simulation, the strategy mixes do not precisely correspond to the Nash outcome. Marks (1992) suggests that the GA converges to evolutionarily stable strategies (ESS), which are states in which a population is resistant to small invasions of mutant strategies²⁴.

The concept of Nash equilibrium²⁵, however, has a strong hold on economic modelling, so why should we be interested in a method that generates ESSs? The attraction of Nash equilibrium derives from the rational compulsion it embodies: it seems that no rational agent, if shown a Nash equilibrium, could possibly deviate from it. However, as the settings become more complex, it is similarly difficult to imagine how boundedly rational agents might perform the necessary analysis and computation to determine what the Nash equilibrium of their game is.

The ESS, on the other hand, compels by its simplicity and inexorability. Axelrod (1984), justifying his use of the notion of Collective Stability²⁶, writes:

“The motivation of applying collective stability to the analysis of people’s behaviour 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...**” (p57, emphasis added).

This captures well the attraction of ESSs over Nash equilibria. Whereas the Nash optimiser of most economic models requires fantastic computing and analytic skills, the ESS requires only “trial and error, or just plain luck”.

²⁴ The concept was developed by Maynard Smith (1974). The extension by Riley (1978) to finite population sizes is particularly relevant to the case of the GA’s equilibria. Selten’s concept of trembling hand perfect equilibrium (Selten(1983)) is closely related. Marks (1992) contains a useful discussion of the relationship between GA equilibria and other equilibrium concepts.

²⁵ And the refinement, sub-game perfect equilibrium (Selten (1983)).

²⁶ Collective Stability (CS) is very similar to Evolutionary Stability (ES). ES entails CS, but not vice versa.

However, in the games examined above, it is not really reasonable to assume that the equilibrium outcomes have arisen from firms actually trying out, in the market, all the strategies that precede equilibrium. This would require a lack of rational deliberation almost as incredible as the rational sophistication required by Nash equilibrium. It is therefore important to consider the types of interpretation that can be applied to the evolving strategies. What exactly is it that evolves²⁷?

To answer this fully would require a detailed model of decision-making within the firm. But even without a fully spelt out model, one candidate for the rôle 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 dealt with. Under this interpretation, the GA simulations above actually model the decision process within one firm as a series of "if... then..." calculations, with firms weeding out unviable scenarios²⁸.

As noted in the description of the Pool game, the appearance of Nash equilibria is sensitive to the number of strategies breeding, surviving and dying in each generation. Below approximately 12% of the population, the strategy variable distributions seem to cycle in suggestive, but definitely out-of-equilibrium ways. Above about 20% of the population, the GA often locks-in to non-Nash equilibria, and population diversity falls very rapidly²⁹. 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."³⁰

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 the size of the breeding pool. High values for the breeding pool (above about 20%) can be thought of as corporate "single-mindedness", whilst low values (below about 12%) would represent "indecisiveness." It is somewhere between single-mindedness and indecisiveness that the firm can hope to remain sufficiently open-minded to attain optimal outcomes³¹.

Using the GA in applied simulation

Markets occasionally have to be modelled in greater detail than offered by the stylised abstractions we have looked at so far. This can happen in prescriptive applications, where a firm or regulator is being counselled. For example, Harbord and von der Fehr (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

²⁷ "What Evolves?" is a question rightly highlighted by Anderson (1994) as central to the development of a fully evolutionary theory of the firm.

²⁸ Penrose (1989) offers a strikingly similar model of how mathematicians arrive at mathematical truth. He writes (p546): "There must be a powerful impressive selection process that allows the conscious mind to be disturbed only by 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." Similarly, 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 re-combination.

²⁹ 20% of the population breeding in each generation means, with each couple have two offspring strategies, that 40% of the population is replaced in generation. It is easy to see how just a few generations of such extensive replacement will reduce strategy diversity and lead to lock-in.

³⁰ Arthur (1993) uses data from experimental psychology to show that humans in laboratory settings typically do not allow this to happen.

³¹ This interpretation sits nicely with management theorists' current emphasis of the firm as a "learning machine."

longer just price and quantity, but timing, discreet and multiple capacity tranches, entry, transmission capacity et cetera.

The method employed by Harbord and von der Fehr³² 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 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, as noted above, the mutation and recombination 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.

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 arbitrary³³.

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 programming³⁴. 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

³² The details of the actual runs are available in London Economics and Harbord Associates (1995).

³³ Lane (1993) argues that "The more richly detailed a model is, the more intriguing it is to its designers - but the less likely it is to capture anyone else's imagination or interest." This is all the more the case when important variables are chosen by the modeller rather than the model.

³⁴ See Holland (1986), 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.

decision, so evolving computer programs to represent firm's behaviour and decision making processes has a very clear analogue³⁵.

³⁵ Genetic Programming can be thought of as a way of endogenising Nelson and Winters' (1982) "routines". See Anderson (1994).

References

- Anderson, E.S. (1994), "Evolutionary Economics - Post-Schumpeterian contributions", Pinter, London.
- Arthur, W.B. (1993), "On designing economic agents that behave like human agents", *Journal of Evolutionary Economics*, 3 pp 1-22.
- Axelrod, R. (1984), "The Evolution of Cooperation", Basic, New York.
- Davis, L. ed (1991), "Handbook of Genetic Algorithms", Van Nostrand Reinhold, New York.
- Goldberg, D.E. (1988), "Genetic Algorithms in Search, Optimisation and Machine learning.", Addison Wesley, Reading, Mass.
- Holland, J. H. (1975), "Adaptation in Natural and Artificial Systems", MIT Press, Cambridge, MA.
- Holland, J. H. , Miller, J. H. (1991), "Artificial Adaptive Agents in Economic Theory", *American Economic Association - Proceedings of the 103 rd Annual Meeting* pp 365-370.
- Holland, J.H. , Holyoak, K.J. , Nisbet, R.E. , Thagard, P.R. (1987), "Induction: Processes of Inference, Learning and Discovery", MIT Press, Cambridge, MA.
- Koza, J. R. (1992), "Genetic Programming. On the programming of computers by means of natural selection", MIT Press, Cambridge, MA.
- Kreps, D.M. (1990), "A Course in Microeconomic Theory", Harvester Wheatsheaf.
- Kreps, D.M. , Scheinkman, J. (1983), "Quantity Precommitment and Bertrand competition yield Cournot Outcomes", *Bell Journal of Economics*, 14 pp 326-337.
- Lane, David A. (1993), "Artificial Worlds and Economics, Part II", *Journal Of Evolutionary Economics*, 3 pp 177-197.
- London Economics , Harbord Associates (1995), "Market Power in the Australian Electricity Market. A report to the Industry Commission", 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 pp 329-73.
- Marks, R.E. (1992), "Breeding Hybrid Strategies: optimal behaviour for oligopolists", *Journal of Evolutionary Economics*, 2 pp 17-38.
- Maynard Smith, J. (1974), "The Theory of Games and the Evolution of Animal Conflict", *Journal of Theoretical Biology*, 47.
- Nelson, R.R. , Winter, S.G. (1982), "An Evolutionary Theory of the Firm", Belknap, Harvard.
- Penrose, R. (1990), "The Emperor's New Mind", Oxford University Press, New York.
- Riley, J.G. (1979), "Evolutionary Equilibrium Strategies", *Journal of Theoretical Biology*, 76 pp 109-123.
- Selten, R. (1983), "Evolutionary Stability in extensive two person games.", *Mathematical Social Sciences*, 5 pp 269-363.
- Tirole, J. (1988), "The Theory of Industrial Organisation", MIT Press, Cambridge, Mass..
- von der Fehr, N.H. , Harbord, D.), "Market Structure and Competition in Electricity Markets: The Australian Debate.
- von der Fehr, N.H. , Harbord, D. (1993), "Spot Market Competition in the UK Electricity Industry", *The Economic Journal*, 103 pp 531-546.