

An Agent-Based Model of the Evolution of Market Structure and Competition

Paul Ormerod (pormerod@volterra.co.uk)

Bridget Rosewell (browsewell@volterra.co.uk)

Laurence Smith (lsmith@volterra.co.uk)

Volterra Consulting Ltd

**The Old Power Station
121 Mortlake High Street
London SW14 8SN**

Paper presented at Computing in Economics and Finance - 2002 in Aix-en-Provence

Abstract

We consider in this paper the evolution of a market in which a single product is produced, which can be differentiated both on price and quality. The specific focus is upon the consequences of new entrants into a market in which, initially, there is a single monopoly supplier. We set up an agent-based model of firms and consumers, each following particular rules of behaviour for pricing and purchasing.

The model extends beyond the comparative statics of conventional theory. In our model, the market evolves over time, and solutions to the model describe the market structure which evolves and emerges from the process of competition.

The model contains a strong element of contingency, so each individual solution of the model is unique.

On average, across a large number of solutions, the price falls (and the quality improves) from that set by the monopolist to a level very similar to that implied by the Cournot model, given the average number of firms which survive in the model.

However, there is little or no connection across the individual solutions between the price which eventually obtains and the number of firms which survive in the market.

Judged on the conventional criterion of the distribution of market shares, at any point in time the market structure is, in general, anti-competitive. But, as the outcome on market price shows, the model is competitive in any meaningful sense of the word.

The potential range of outcomes for the model is wide, reflecting the importance of contingency in the process of the evolution of market structure. A certain amount of the variation can be accounted for deterministically, but overall the results contain a substantial stochastic component.

1. Introduction

Conventional approaches developed to examine competition and monopoly generally examine the market outcomes in individual situations and compare them with alternative potential outcomes. Often, the comparison is made between the outcome of a perfectly competitive – desirable – market and the actual situation that is observed. Policy is then directed towards moving the industry under examination closer to the description of the competitive outcome. This approach does not always produce the intended consequences, however, or leaves a suspicion that market power is not being properly characterised. For example, the changes imposed in the UK on brewers' rights to tie their outlets were intended to support the existence of smaller brewers. This is not what occurred. Equally, repeated investigations of food supermarkets in the UK have failed to identify the exercise of monopoly power – yet public discontent remains.

One possible reason for these apparent failures is that too little attention is paid to the *process* by which the market might change and develop – or has already changed - in the face of structural change. In addition, there is little understanding of the time over which competitive forces operate and over which stable situations emerge.

This paper focuses on the competitive process in one particular example, where it has already been recognised that conventional analysis may be weak. Dynamic, high tech industries are characterised by factors such as falling average costs over a broad range of output, quick and frequent entry and exit with modest capital requirements, economies of scale in consumption and high rates of innovation. Rapid change means that an analysis of process becomes essential. (see, for example, Posner, 2000).

The American authorities have already recognised in principle the challenges to regulatory thinking which such industries pose. For example, Robert Pitofsky, chairman of the US Federal Trade Commission, stated in giving evidence to a Senate committee on 'Mergers and Corporate Consolidation in the New Economy': 'Merger analysis has moved from strict reliance on structure-based presumptions that focused largely on market share

data to a sophisticated analysis that takes account of the dynamic nature of competition in the real world' (Pitofsky, 1998).

We consider in this paper the evolution of a market in an industry in which a single product is produced. The specific focus is upon the consequences of new entrants into a market in which, initially, there is a single monopoly supplier. Companies can compete on both the price and the quality of their offer. In other words, in common with almost all goods and services which are produced in reality, the product is not completely homogenous.

The model we develop can be applied in its current form to any industry where technical or regulatory change is making entry of new firms possible. Application to a specific industry requires calibration of the model to the particular characteristics of the industry. The paper therefore outlines a template which illustrates a general model of market change.

Our model formalises the principle of contestability of markets. We specify a process which determines the timing and number of firms which enter the market, and the prices and quality levels which they offer. The reaction of existing firms to new entrants is described, as is the behaviour of consumers when confronted with a new set of prices and quality levels.

The monopoly may have arisen for a variety of reasons. For example, in the twentieth century, in many Western European countries, the physical distribution network of utilities such as telecommunications, gas and electricity was captured by the state, and a single monopoly producer was given control of the physical network. A more market driven example is that of the establishment of VHS as the industry standard in videotape. As VHS sales increased, the production costs of recorders fell as the scale of production rose. Further, as the product became established as the industry standard, economies of scale in consumption arose through more related products becoming available, such as

the wider availability of tapes. The computer industry provides another example, with software becoming written increasingly in a particular operating system.

In each of these cases, the monopolist can be thought of as having captured the sales network to consumers. In the case of utilities, this can mean quite literally the physical network over which the product is delivered. More generally, however, this particular concept of a network need not have a physical presence, and refers simply to the connections from a supplier to the consumers. Consumers can only buy from those companies of whose product they are aware. The phrase 'sales network' in this paper means the set of connections from a firm to consumers. By definition, a monopolist has a sales network which connects it to all consumers in the particular market, whether or not the product is supplied across a physical network.

In this paper, we provide an example which starts by assuming that a monopoly has captured the entire sales network – either as a result of state control or for other reasons. We then consider the impact of new entrants to this market. Such new entrants may be made possible by de-regulation and privatisation, but there is also a more general interpretation where such entrants become possible because of new technology. When new entrants come into a market, their ability to sell depends on the customers that they are able to reach, either directly or indirectly. Unless potential customers are able to consider the new offer, they will be unable to take it up. Each new entrant has access to a sales network, which is an important feature of our model. However, the sizes of the networks across new entrants is fixed as a set of parameters in the model, and we do not in this paper introduce potential new economies of scale in consumption.

The model was first developed to illustrate the development of the UK telecommunications industry, but is intended to be a more general approach to the questions of market entry and consumer and producer behaviour in such circumstances. We set up an agent-based model of firms and consumers, each following particular rules of behaviour for pricing, quality setting and purchasing. At the outset there is a single monopoly supplier. By construction, the market eventually contains an indeterminate

number of producers, each of which supplies the product close to the minimum price which is both technologically feasible and enables a normal rate of profit to be made. In other words, we begin with a monopoly and end with outcomes on price which on average are close to the perfectly competitive market of standard economic theory.

However, the solutions extend beyond the comparative statics of conventional theory. In our model, the market evolves over time, and solutions to the model describe the market structure which evolves and *emerges* from the process of competition.

Section 2 of the paper sets out the theoretical model. The specification is deliberately kept as parsimonious as possible, whilst at the same time maintaining a reasonable degree of realism. The model, though containing only a small number of behavioural rules and parameters, nevertheless generates a potentially complex range of results. Section 3 presents the properties of the model. Some implications are discussed in section 4.

2. The theoretical model

2.1 Overview

Initially, there is a single monopoly supplier selling a product, which can be differentiated in terms of its overall quality, to a large number of consumers. The model evolves on a step by step basis, with each step of the model representing a period of calendar time. In the case of the general specification of the model, a market structure is eventually reached in which there is an indeterminate number of suppliers. Each of these firms supplies the product at a price below that which obtains in the initial conditions, and which is close to the minimum which is technologically feasible. The original incumbent need not be (though almost always is) one of these suppliers.

2.2 Description of the model

The market is populated by n consumers. We assume for simplicity that they each consume an identical amount of the product in any given period. The amount spent per period by each consumer, and hence total sales of the product, may change over time, but our interest in this paper is on, amongst other things, the market shares of the producers rather than on the total size of the market. So the amount spent by each consumer is the same in any *given* period.

Initially, the market contains a single monopoly supplier, selling the product at a price of p_{mon} (using the subscript 'mon' to indicate the incumbent firm's monopoly price), and with a quality q_{mon} . The model evolves on a step-by-step basis, in which each step is a period of time. We specify a process by which other firms enter the market, both in terms of frequency and in terms of the total numbers entering each period. An obvious way to do this is by a draw from a binomial distribution $B(m, p)$, where we specify the maximum possible number of new entrants in any period m and the mean number of new entrants $m \cdot p$. We specify $(k - 1)$ new entrants in total, where $k < n$.

Each new entrant comes into the market with a (p, q) drawn at random from a uniform distribution on $[p_{\text{min}}, p_{\text{mon}}]$ and $[q_{\text{min}}, q_{\text{mon}}]$. The price p_{min} is the lowest possible price at which, after the process of technological innovation is complete, the product can be offered and a normal rate of profit obtained by the most efficient supplier. The quality q_{min} is the best quality at which the product can be offered, again subject to a normal rate of profit being obtained. Note that, for simplicity, we measure quality inversely, so that the *lower* the quality measure in our model, the *better* the quality is. Note also that we express quality on a single dimension. This does not mean that the product necessarily has only one feature which measures its quality. There could well be several features, which are concentrated into a single measure of overall quality.

All n consumers are connected on a network to the initial monopoly supplier. This could be in the case of telecommunications quite literally a physical network, but the use of the word 'network' in this physical sense is too limiting. 'Network' in this context means, as

noted in section 1, that consumers on the network of firm f_i are both aware of the offer from firm f_i and are willing to consider buying from it.

Each new entrant obtains potential access to a network of consumers. We can specify this to be all n consumers for each firm, or for each firm to have access to a fixed proportion of all consumers, with the precise consumers available to each firm chosen at random. More generally, however, and in the results reported below, each new entrant obtains potential access to a proportion of the total number of consumers drawn at random from a uniform distribution on $[v_{\min}, v_{\max}]$, where $v_{\min, \max} \in [0, 1]$. Once the group of customers to which a firm has potential access has been chosen, it is set up immediately. A further simplification is that the group then remains fixed during all subsequent steps in that particular solution of the model.

There are three obvious reasons why new firms in the market do not have potential access (in general) to all consumers, which can obtain either singly or in combination. First, the regulator could impose restrictions so that, for example and purely by way of illustration from the telephone market, a new entrant could be permitted to offer international calls but not domestic ones. Second, the marketing strategy of the firm may be such that not all consumers are aware that the firm is making an offer in the market. In reality, marketing strategies vary widely in effectiveness, and this is reflected in our model. Third, the firm itself may deliberately target only a small percentage of consumers. In the context of British land line phone calls, for example, several firms now specialise in offering cheap calls to India, say, or to the United States.

In each period, each consumer reviews the price and quality of each of the firms to which he or she is connected. The consumer at any point in time is only permitted to buy from a single supplier. This is not always completely realistic, but is a reasonable assumption to make in this initial specification of the model. The consumer is allocated from the outset a weight w_i , which expresses his or her preference between price and quality (w_i is chosen from a uniform distribution on $[0,1]$). The consumer calculates for each of the

firms on his or her network $w_i * p + (1-w_i) * q$. For the k th firm, we describe $w_i * p_k + (1-w_i) * q_k$ as being the overall value to the i th consumer of this offer – v_{ik} for short.

The consumer switches all of his or her business to the firm offering the lowest v_{ik} of the k firms on his/her network (which may not correspond the lowest $w_i * p + (1-w_i) * q$ then on offer to other consumers, because the particular consumer concerned may not be aware of such offers), subject to the following condition. At the outset, each consumer is allocated a 'switching propensity', s_i , which is drawn at random within $[0, 1]$. The exact functional form which is used in this paper is set out below in section 3.2. If the customer identifies a v_{ik} , which is lower than that of his or her existing supplier, v_{im} , he or she will then switch to firm m from firm k , with probability s_i .

There are several reasons for introducing this probabilistic element into the choice. Although the product offers of the firms are very similar, they are not perfect substitutes, for two reasons. First, the lowest (p,q) supplier may specialise in an offer which is not very important to a given consumer. Someone who makes only local phone calls will not be interested in a firm which provides only cheap international calls. Second, even within the same segment of the market, such as local calls, the product is not completely homogenous in that consumers may have doubts about the reliability of a previously unknown supplier.

There are two other possible reasons why consumers will not in general switch to the lowest price producer. First, there may be costs involved in switching. To take an obvious example, if changing suppliers involved having to change telephone number – staying with the telecomm example - for most people the savings on price would have to be considerable to offset the inconvenience involved. Second, consumers may simply exhibit inertia and stay with their existing supplier, perhaps because the savings involved are small.

At the start of the next period, each firm already in the market is given the opportunity to reduce its (p,q) offer. Firms are not certain about the distribution of preference across

consumers regarding price and quality, and so assign equal weight to price and quality in each of the (p,q) offers which they observe¹. They aspire to move to the (p,q) of the firm for which $\omega * p_i + (1 - \omega) * q_i$ is minimised, where ω is the average of the w_i across all consumers.

However, firms differ in their ability to adapt their organisation in order to deliver the desired (p,q) offer. We can think, for example, we can think of firms as differing in their level of X-efficiency. The ability of the firm to do achieve the desired (p,q) depends on the firm's flexibility level ϕ_i . At the outset, each firm is allocated a flexibility level, ϕ_i , which is drawn at random from a uniform distribution on $[\phi_{\min}, \phi_{\max}]$, where $\phi_{\min, \max} \in [0, 1]$. In each period, each firm switches to the lowest $\omega * p_i + (1 - \omega) * q_i$ with probability ϕ_i .

Consumers then review their choice of suppliers given the revised set of (p,q) from existing suppliers, and given the (p,q) offered by new entrants (if any) in that period.

The model proceeds on a step by step basis, with new firms entering until the pre-specified total number of firms, k , has been reached. The model can then continue to be solved in exactly the same way over subsequent steps, except that there are no more new entrants.

In the limit, of course, if the model is allowed to run for a sufficiently long time with a very large number of potential entrants, both p and q will in general converge towards zero. However, a key feature of the model is that it is not intended to describe the outcomes of two equilibria – the first with the monopolist setting (p,q) at $(1,1)$, and the second a perfectly competitive outcome in which both p and q approach zero in the limit, the second of which obtains after a totally unspecified period of time. Rather, we

¹ technically, they assign a weight equal to the average across all consumers. Of course, when the model is populated by a large number of consumers, the average weights on price and quality will be, with very high probability, very close to equal

recognise that economic processes are rooted in time. It is not at all useful to say that, once a monopoly has been opened up to competition, *eventually*, at some unspecified point in the future, a different equilibrium will prevail in which (p,q) is no longer at $(1,1)$ but is arbitrarily close to $(0,0)$. Indeed, Atkinson (1969.) provides evidence to suggest that the timescale of transition between equilibria in many theoretical economic models is extremely long. We are interested in how the market evolves over a specified and realistic time-scale – up to 10 years, say. And we are interested in the relationship between (p,q) and the market structure which evolves over such a time scale.

We do not specify an explicit process for exit from the market by unsuccessful firms. Once a firm has entered the market, it is able in principle to acquire customers both in the period in which it enters, and in each subsequent period. A firm may enter in a given period, only to find its (p,q) to be undercut by competitors, so that it makes zero sales in that period. By reducing its (p,q) in the next period, it is possible for the firm to then gain some customers, but this number is likely to be small. Clearly, the more periods in which a firm fails to make any sales, the less likely it becomes that it will do so in future periods. More firms may enter the market, and others already in may reduce their (p,q) . Firms which fail to make any sales in the period in which they enter the market will, on average, remain in this condition, so their 'exit' takes place at the same time as their entry.

More generally, we do not specify any minimum number of sales which are required for a firm to remain in the market. Once a firm enters the market, it remains there, regardless of its level of sales.

Finally there is the issue, frequently neglected in economics, of how to translate the concept of time in theoretical models into realistic time scales. Often, the question is avoided altogether by recourse to the method of comparative statics. In other words, the equilibrium solution of a model after a posited change has taken place is compared to the equilibrium situation beforehand, and the issue of time is not addressed. Atkinson (op.cit.) demonstrated some of the potential pitfalls of this neglect, showing, for example,

that the implied time scale over which adjustment between equilibrium paths took place in neo-classical models of economic growth could be of the order of 100 years.

Our model evolves on a step by step basis, where each step corresponds to a unit of time. There is no unequivocal, natural definition of the unit of time in the model. But switching by consumers on the basis of price is a central feature of the model. Given that in many utility markets, for example, consumers receive bills quarterly, it seems sensible to define a unit of time in the model as being one quarter in real time.

3. The results

3.1 Preliminary remarks

As we noted in the introduction, the model has been kept as parsimonious as is consistent with achieving a reasonable degree of realism in describing how markets with new entrants might evolve. It is, however, complex, and the exploration of its properties is not straightforward. In terms of parameters, for example, we have

- The numbers of new entrants
- the prices at which new firms enter the market
- the quality of the product with which new firms enter the market
- the flexibility of each firm, measuring its ability to react to lower prices and better quality levels offered by competitors
- the switching propensity of each consumer

The time-scale over which the model converges to a stable distribution of market share at a price close to the minimum possible price and a quality close to the best can be affected substantially by the choice of the ranges from which these parameters are drawn. Further, the outcome as far as the incumbent monopolist is concerned is affected by its position in the range of flexibility.

An additional complication is the choice of process by which both the timing and numbers of new entrants are determined. This affects not just the time-scale but the actual outcomes of the model, as does the choice of the networks of consumers to which the new entrants are connected.

In this paper, we therefore report on a realistic scenario for the timing and the number of new entrants. A relatively low number of entrants is introduced over a reasonably long time-scale, so that in any given period the most likely number of new entrants is zero. Each of these entrants has a potentially high level of connectivity. In other words, the value chosen for v_{\max} , the maximum proportion of total consumers which can be connected to any new entrant, is high

In constructing this scenario, we have in mind a type of economic structure which is intended to describe a process in which entry into the market is mainly by other large firms. The firms may be carrying out important brand extensions, or may even be investing in a new and more modern type of physical infrastructure.

3.2 Parameterisation of the model

We populate the model with 1,000 consumers. This is a sufficient number such that every eventuality can occur. The use of a greater number of consumers is not required in order to understand the properties of the model.

We introduce 19 new firms over a 10 year period (giving a total potential of 20 firms in the market). The entry pattern is drawn from a binomial distribution $B(10, 1/20)$. So the maximum number of entrants in any quarter is 10 but the mean number of entrants is $1/2$. The draw is constrained so that there are exactly 19 entrants. The proportion of the total number of consumers to which these firms are attached on a network is drawn at random from a uniform distribution on $[0, 1]$. In other words, new entrants get *potential* access to

between 0 and 100 per cent of consumers. The expected mean value is, of course, 50 per cent across the entrants.

The price of the incumbent monopolist is set initially equal to 1, and the prices of new entrants are drawn at random from a uniform distribution on $[0,1]$. This does not, of course, mean that the price might literally be zero in reality. The range $[0, 1]$ simply re-scales the range across which prices in any actual situation might vary.

It is worth expanding this point. As Vickers (1994) notes 'in the textbook competitive model all sellers of the good are assumed already to know the minimum cost of production, ... it is only through the process of competition that the facts will be discovered'. There are essentially two reasons why prices might fall in this model. First, through competitive pressure as new entrants come into a market previously dominated by a monopolist. Second, as technological innovation over time enables the product to be supplied at prices which were not previously feasible, in the sense that profit could not be made at such prices.

In practice, these two reasons will often be mixed together, but they are conceptually distinct. We might think, for example, of a utility operating with mature technology under a state monopoly. The legal process of de-regulation could permit entrants into this market for the first time, each with access to very similar levels of technology, but operating with varying degrees of efficiency. In these circumstances, prices are reduced purely by competitive pressures. In contrast, we can imagine a situation in which a company has acquired over time an effective monopoly. A technological innovation is made which enables competitors to enter the market, and continuing investment in research and development enables prices to fall further and further.

Our model does not in its present form distinguish between these two sets of circumstances. In both cases, prices are drawn uniformly from $[0, 1]$. Yet in practice, prices are likely to fall further in the second of these cases than in the first. The range $[0, 1]$ simply re-scales both these ranges in the model.

The flexibility parameter, ϕ_i , for each firm is also drawn at random from a uniform distribution on $[0,1]$. This means that, on average, the incumbent will have a flexibility level of 0.5, exactly in the middle of the distribution. In practice, whether through years of inertia brought about by the previous monopoly position, or because of actions of the regulator, the incumbent may well have a flexibility level which is initially lower than the average. However, we explore here the implications of the incumbent having, on average across the set of solutions, the average level of flexibility.

The switching parameter, s_i , for each consumer is set up as follows. First, a random draw from a uniform distribution on $[0, 1]$ is made, giving a σ_i for each consumer. The switching parameter s_i is then given by $s_i = 1 - \sqrt{\sigma_i}$. This gives a distribution of the s_i such that most consumers have a fairly low propensity to switch in any given period, and a small number have a high propensity to switch. The probability density function is plotted in Figure 1 below. The mean value of this distribution is $1/3$. Supposing, for example, that a consumer with this value of s is confronted on four successive occasions with the choice of switching, the probability that he or she will not do so is just under 0.20. If the s_i were chosen instead simply drawing from a uniform distribution on $[0, 1]$ giving a mean of $1/2$, this probability would of course be only 0.0625, a figure which is reached after approximately seven periods with the choice of s_i which we actually use.

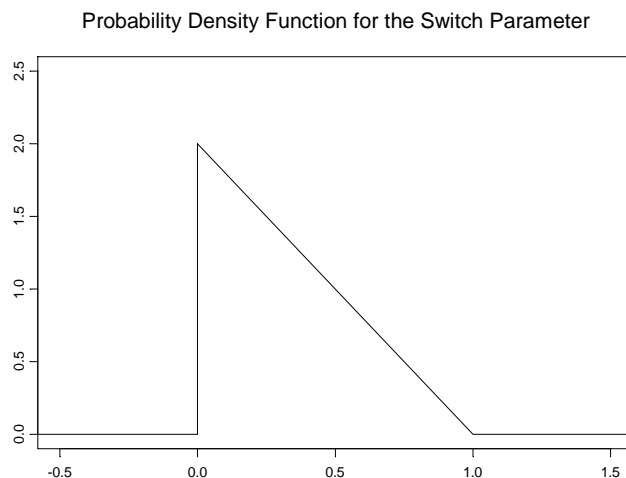


Figure 3

3.3 The results

The probabilistic nature of the model means that its properties must be explored by simulation. Under each scenario, we therefore carry out 1000 separate solutions of the model. Each time, we start from the initial conditions that the monopolist has 100 per cent of the market, selling at $p_{\text{mon}} = 1$ at a quality $q_{\text{mon}} = 1$, where 1 is the worst level of quality and 0 the best. The timing and numbers of firms which enter, and their networks to consumers, also varies across each of the simulations. The scenarios are different from each other in five respects:

- the flexibility of each firm, φ_i , is drawn separately for each scenario
- the switching propensity of individuals, s_i , is also drawn separately
- the prices and qualities at which firms enter the market differ between the scenarios
- the timing at which firms try to enter varies between the scenarios
- the proportion of consumers to which each firm gains potential access also differs between scenarios

Figure 2 plots the histogram of the average price which emerges after 40 periods of the solutions of the model.

Histogram of average price outcomes after 40 periods




Figure 2 Histogram of average price outcomes in 1,000 solutions of the model after 40 periods

The single most frequently observed outcomes for the market price is in the range 0.05-0.10. In other words, price does fall to a level close to the minimum which is feasible. Occasionally, however, the price remains relatively high. In terms of the distribution of outcomes across a number of solutions of the model, the quality of the product evolves similarly to price, although the outcome of the two may obviously differ in any given scenario.

The mean level of market price after 40 periods is 0.145, with a minimum of 0.00007 and a maximum of 0.650. The inter-quartile range is between 0.057 and 0.206.

Figure 3 plots the histogram of the outcomes, again after 40 periods, across the 1000 solutions of the model of the market share of the initial monopolist.

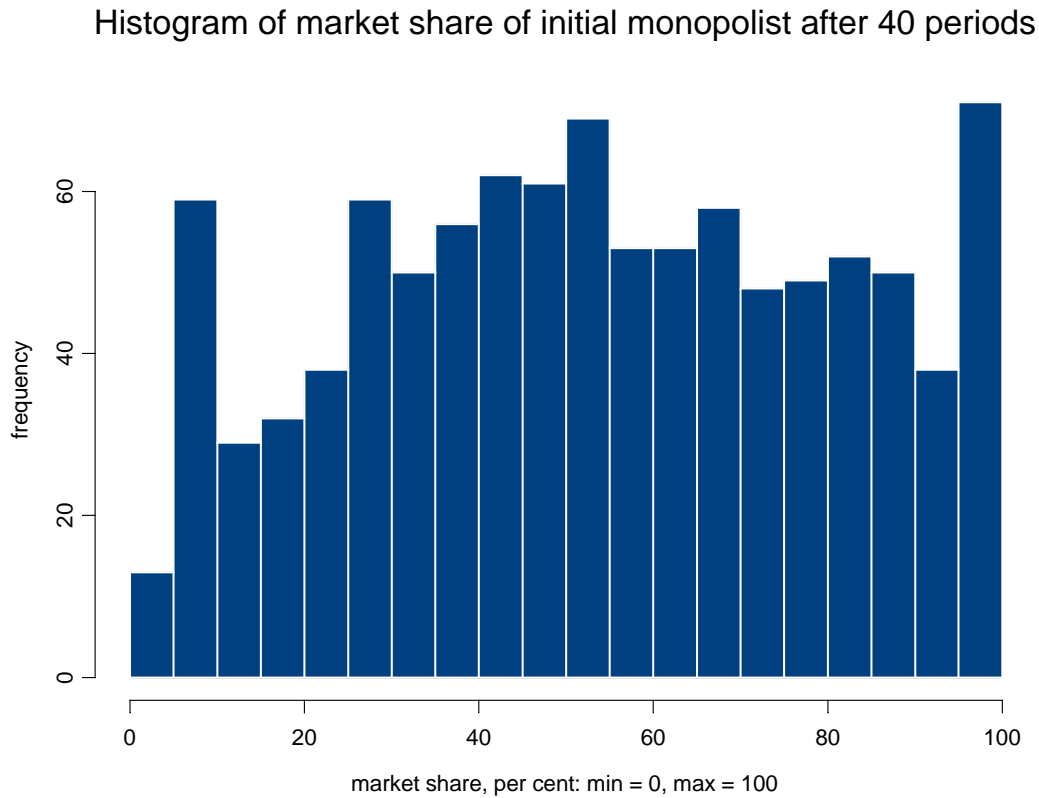


Figure 3 *Histogram of market share of the initial monopolist in 1,000 solutions of the model after 40 periods*

Quite frequently, the incumbent monopolist retains a very high market share. Despite this, as Figure 2 shows, the market price usually falls very sharply. The average market share of the monopolist after 40 periods is 53.5 per cent, with a minimum of 3.4 and a maximum of 100 per cent. The inter-quartile range is wide, between 32.1 and 75.9 per cent.

An important determinant of the eventual market share of the monopolist is its flexibility, as Figure 4 below shows.

Flexibility of initial monopolist and market share after 40 periods




Figure 4 *Scatter plot of 1,000 solutions of the model between the flexibility of the initial monopolist and its market share after 40 periods*

A high level of flexibility is by no means a guarantee of a high eventual market share, but the simple correlation between the two variables is 0.712.

A wide range of outcomes is possible for the number of firms with non-zero sales after 40 periods, as Figure 5 shows.

Histogram of number of selling firms after 40 periods




Figure 5 *Histogram of number of selling firms in 1,000 solutions of the model after 40 periods*

The mean number of firms is 8.2, so that on average almost 12 out of the 20 firms fail completely i.e. have no sales at all. This seems compatible with the outcomes which are observed in practice (see, for example, Carroll and Hannan (2000)).

Figure 6 sets out the distribution of the average market share of the 8 largest selling firms after 40 periods. (Of course, in a number of the solutions of the model, when fewer than 8 have non-zero sales, the value for some of the ‘largest’ 8 in these particular solutions is zero).

Distribution of average market share after 40 periods

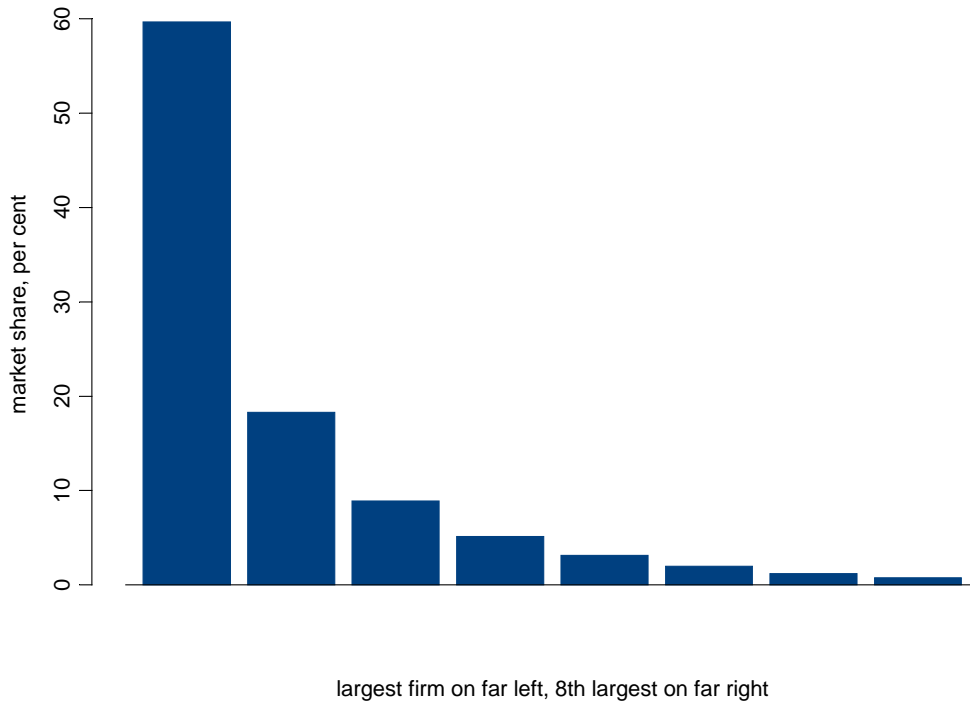


Figure 6 *Distribution of market share of the 8 largest firms in 1,000 solutions of the model after 40 periods*

The figure shows the distribution regardless of the identity of the firm. In the majority of solutions of the model, the largest firm is the initial monopolist, but this is not always the case.

Interestingly, a good approximation to the size distribution of the largest 8 firms after 40 periods is provided by a power law. Axtel (2001) shows that this a general characteristic of the distribution of firm sizes in the United States. A log-log least squares fit of average market share in Figure 6 on the rank of the firm by market share (largest has rank equal to 1, etc.) gives an R^2 of 0.983 and an estimated exponent of -2.09 with a standard error of 0.203. An exponential also gives a good approximation to the distribution of firm size, but the power law is better.

3.4 Does the most flexible firm survive?

The concept of flexibility in this model is, of course, that of the ability to respond quickly to competitive pressures. A firm with high flexibility is much more likely to be able to cut its price immediately to match the lowest price on offer than one with a low flexibility.

In this analysis, we exclude the incumbent monopolist because of its special circumstances of access to the network of consumers. We monitor in each solution of the model the firm with the lowest level of flexibility, and observe its eventual market share. The same is done for the firm with the highest level of flexibility. Table 1 sets out the summary statistics for these factors

Table 1 Summary statistics : firms with lowest and highest flexibility and the eventual market shares

	min.flex	min.share	max.flex	max.share
Min.:	0.000	Min.:0.000	Min.:0.479	Min.:0.000
1st Qu.	0.014	1st Qu.:0.000	1st Qu.:0.931	1st Qu.:0.000
Mean:	0.050	Mean:0.020	Mean:0.950	Mean:0.027
3rd Qu.	0.069	3rd Qu.:0.011	3rd Qu.:0.985	3rd Qu.:0.002
Max.:	0.424	Max.:0.424	Max.:0.999	Max.:0.694

Min flex is the vector of the least flexibility in each simulation; max flex is the vector of the highest flexibility in each simulation; min share is the vector of the end share of the firm with the least flexibility in each simulation; max share is the vector of the end share of the firm with the highest flexibility in each simulation

Note: 1) the 3rd quartile of the distributions for min.share and max.share are less than the mean – in each case a relatively small number of large values increases the mean

2) both the min.flex and the min.share happen to be 0.424 – it is not a typing error

The difference between 'min share' and 'max share' is clear. For example, the maximum eventual market share of the firm with least flexibility in 1,000 solutions of the model is 42.4 per cent, and for the firm with highest flexibility it is 69.4 per cent. A formal Kolmogorov-Smirnov test of the null hypothesis that the vectors of 'min share' and 'max share' have the same distributions is rejected at $p = 0.000$. So flexibility does play some part in the commercial success of new entrants.

However, a high level of flexibility is no guarantee of success. No fewer than 728 out of the 1,000 firms with maximum flexibility have an eventual market share of zero, compared to 619 of the firms with minimum flexibility. A great deal depends upon the timing of entry. As the model stands, if the most flexible firm enters late in the process, the market price has a high probability of already being very low, so the firm cannot capture any customers. The reaction of the incumbent monopolist is also important, given its unique network which connects it to all consumers. An early price reduction by the monopolist makes it much more difficult for subsequent entrants to succeed.

3.5 Market price and the market share of the monopolist

We can also consider the outcomes of this approach and compare the results with other ways of thinking about the process of competition. For example, it is often thought that reducing the market share of a monopolist (for example by competition policy) will ensure lower prices. We can examine whether there is any connection here between the eventual market share of the monopolist and the prevailing market price. Figure 7 shows the scatter plot between these two variables.

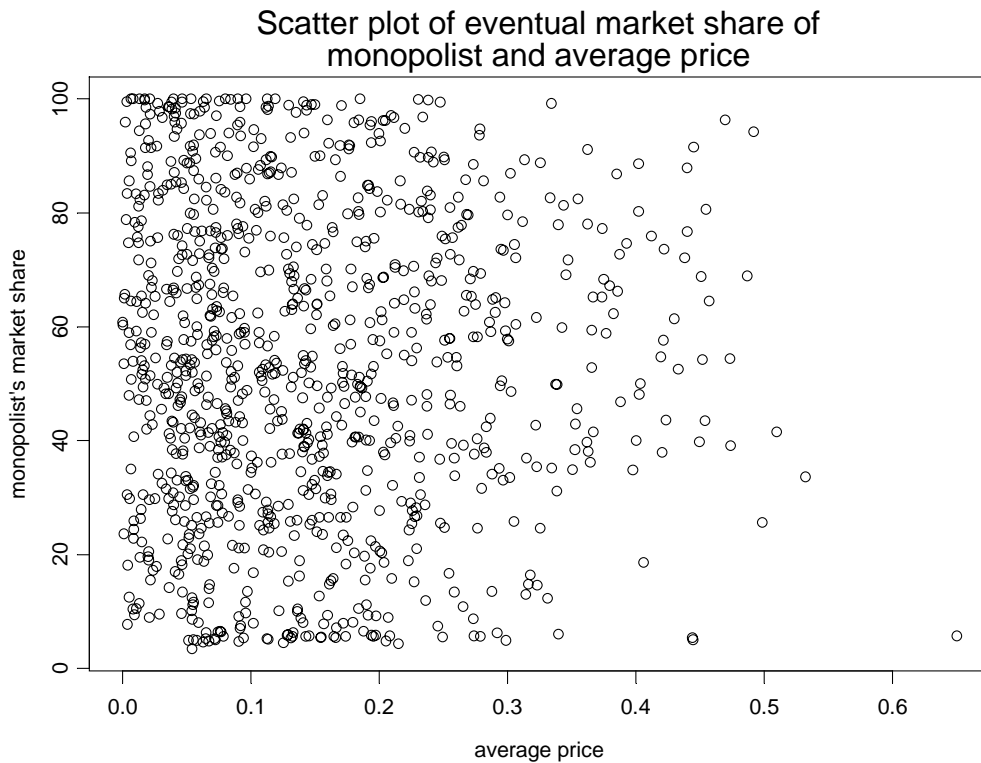


Figure 7. Scatter plot of 1,000 solutions of the model between the eventual market share of the initial monopolist and the average price after 40 periods

The simple correlation between the two is -0.014 .

3.6 Market price and the number of selling firms

We can also look at the relationship between the total number of firms and the final market price. Here, standard economic theory implies a relationship between the equilibrium market price and the number of firms in the market. The fewer the number of firms, the more the price will be above the level which just covers both costs and a normal rate of profit.

The Cournot model of oligopolistic markets is used widely. The exact relationship between the equilibrium market price and the number of firms in the market depends upon the particular form of the market demand function which is assumed to exist. As an approximation, however, it is proportional to $\text{cost} \cdot (1 + 1/N)$, where N is the number of firms in the market and where cost includes a normal level of profit. As N increases, the price converges on cost.

There are two important differences between our model and the standard economic approach. First, as with the incumbent's market share, there is little or no connection between the market price which eventually obtains and the number of firms in the market. Second, for any given number of firms in the market, the market price in our model is below that of the Cournot model, often substantially so.

Figure 8 below plots the relationship between the eventual market price and the number of firms in the market. It is clear that there is little or no connection between the two. The simple correlation is in fact 0.05.

Figure 8. Scatter plot of eventual market price against number of firms in the market

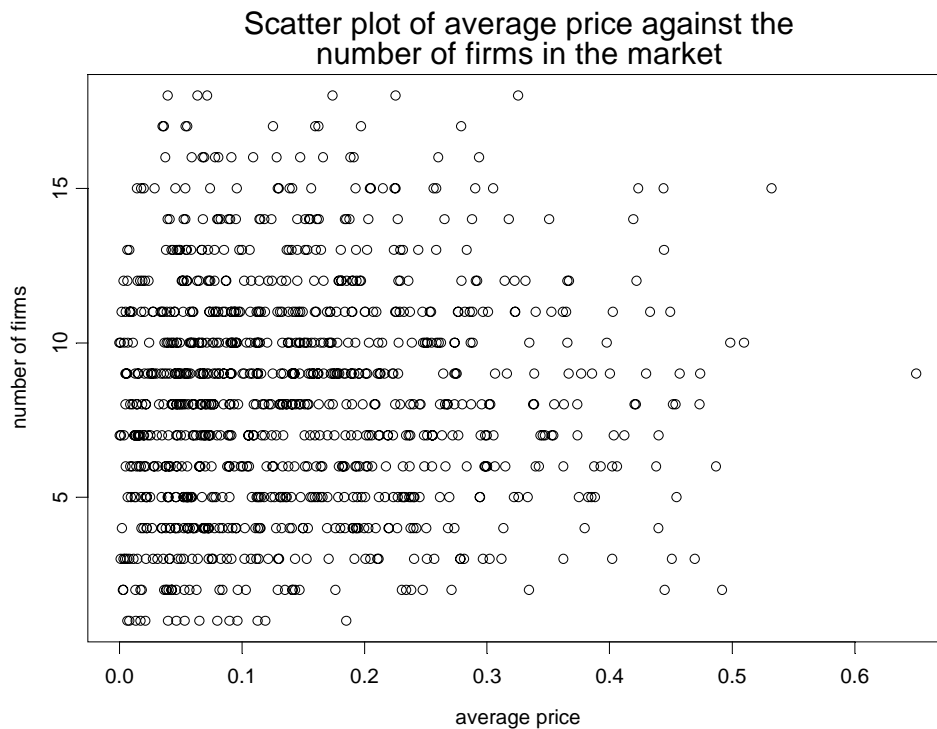


Figure 8 Scatter plot of 1,000 solutions of the model between the number of firms in the market and the average price after 40 periods

The key difference between our model and that of, say, the Cournot model is that with the latter there is a deterministic relationship between the number of firms in the market and the market price which obtains. The more the number of firms, the closer the price becomes to the theoretical level of a perfectly competitive market. In our model, in any particular solution of it there is no necessary connection at all between price and the number of firms. As Figure 8 shows clearly, a very low price can obtain with just one or two firms in the market. Equally, a relative high price may exist with 10 or even 15 firms in the market.

This difference between the Cournot model and our own is much more important than any similarities. However, we can compare the *average* market price which obtains in our model across 1,000 simulations and the *average* number of firms which survive. Purely by coincidence, given the average number of firms which survive in 1,000

solutions of the model, the average price across these solutions is very similar to that of the in the standard Cournot model.

On *average*, after 5 years there are 5.54 firms in total in the market in the simulations of our model, rising to 8.21 after 10 years. Two widely used illustrations of the Cournot model are with a linear and log-linear market demand function, respectively. With a linear demand function, the mark-up on cost is $(1 + 1/(N+1))$, and with a log-linear one it is $(1 + 1/(N-1))$. These imply, respectively, a market price after 5 years which is 15 and 22 per cent above cost. After 10 years the figures are 11 and 14 per cent above cost.

A price of zero in our model is defined to be the lowest technologically feasible price which enables costs to be covered, including a normal profit margin. In these simulations, the *average* market price is 21.5 per cent above this after 5 years, and 14.5 per cent above after 10 years. In other words, our model implies very similar margins *on average* to conventional Cournot solutions. However, to stress again, there is a fundamental difference between the two approaches. Namely, in our model the mere presence of only a small number of firms in the market does not prevent a highly competitive market price from becoming established.

4 Summary and Implications

This paper has established a model of the competitive process based on a set of rules of behaviour which are both basic and plausible. Firms act in order to enter markets profitably and to increase their market share. Consumers behave rationally, purchasing a new product if it is cheaper.

Our model formalises the principle of contestability of markets. We specify a process which determines the timing and number of firms which enter the market, and the prices and quality of the product which they offer. The reaction of existing firms to new entrants is described, as is the behaviour of consumers when confronted with a new set of prices and qualities.

The results of the competitive process, once entry occurs, leads to price in the marketplace which is generally close to the competitive minimum. Individual solutions of the model vary in the speed with which this minimum is approached. In all cases, however, there is a reduction in price and an improvement in quality, and consumers are offered more choice.

In summary, the key results are as follows:

- there is a strong element of contingency in the model, and the evolution of the market over time can differ substantially between individual solutions carried out with the same values of the fixed parameters. However, it is still possible to draw general conclusions about the typical behaviour of the model over large numbers of solutions.
- the market price generally falls from the level set by the initial monopolist to close to the minimum which is both technologically feasible and consistent with a normal margin of profit (quality also improves to become close to the best which is feasible)
- the market price is *on average* very similar to that implied by the Cournot equilibrium given the *average* number of firms with non-zero sales
- however, in any individual solution of the model, the market price which eventually obtains is not really influenced by the number of firms which remain in the market
- the monopolist retains, in general, a substantial share of the market
- judged on the conventional criterion of the distribution of market shares, at any point in time the market structure is, in general, anti-competitive. But as the outcome

on market price shows, the model is highly competitive in any meaningful sense of the word. The original monopolist retains market share by responding to the competition created by new entrants and by lowering its price accordingly.

- the majority of new entrants fail, which seems to fit empirical evidence
- the distribution of market share is approximated closely by a power law, which again conforms with empirical evidence

The results of this approach to the issue should give regulators and policy makers pause for thought when considering contestable markets. For example, it is not the case that a competitive market (in the sense of having a competitive price), will necessarily have lots of firms, or will have driven down the original incumbent's market share. Further, although market share is often used as an indicator – indeed as a primary indicator – of the presence of monopoly power which may lead to anti-competitive behaviour, these results show that this can be seriously misleading. Finally, the existence of an incumbent by itself does not necessarily tell us much about whether the price is low or high and whether the market is competitive or not.

References:

A.B. Atkinson (1969), 'The timescale of economic models: how long is the long run?', *Review of Economic Studies*, XXXVI (2), 137-52

R.L. Axtell (2001), 'Zipf distribution of US firm sizes', *Science*, 293, 1818-1820

G.R. Carroll and M.T. Hannan, (2000), *The Demography of Corporations and Industries*, Princeton University Press

Posner

Pitofsky

J.Vickers (1994), *Concepts of Competition*, inaugural lecture in the University of Oxford,
Clarendon Press, Oxford