An Agent-Based Model of the Extinction Patterns of Capitalism's Largest Firms

Paul Ormerod (pormerod@volterra.co.uk)*

Helen Johns (hjohns@volterra.co.uk)

Laurence Smith (lsmith@volterra.co.uk)

Volterra Consulting Ltd

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

November 2001

***** Corresponding author

We are grateful to the Institute for Complex Additive Systems Analysis of the New Mexico Institute of Mining and Technology for financial support towards this research

Abstract

Power-law distributions (fractal behaviour) in a system's macroscopically observable quantities are a characteristic property of many-body systems representing the effects of complex interactions amongst the constituents of the system. Power law distributions are both self-similar and scale free, demonstrating that events may occur on all lengths and time scales.

The empirical relationship between the frequency and size of extinctions of capitalism's largest firms is described well by a power law. This power law is very similar to that which describes the extinctions of biological species in the fossil record.

We develop an agent-based model of the evolution and extinction of firms based on simple principles of economics. The properties of the model conform closely to the empirical evidence.

The model contains N agents, and all pairs of agents are connected to each other. Individual agents interact with each other, sometimes in co-operative symbiosis and sometimes in direct competition.

The model evolves in a series of steps. The rules of the model specify a) how the connections are updated b) how the fitness of each agent is measured c) how an agent becomes extinct and d) how extinct agents are replaced.

1. Introduction

At the turn of the nineteenth century, large corporations were being built on an unprecedented scale, mainly due to a massive wave of mergers and acquisitions. Hannah (1999) provides a data set of the 100 largest industrial companies in the world in 1912. These are firms which had survived the merger boom at the turn of the century, and were large even by the standards of today. US Steel employed 221,000 workers, and most of the others employed more than 10,000.

By 1995, only 52 of these firms survived in any independent form. Nineteen of the survivors remained in the top 100 industrial companies in 1995, but 24 of them were smaller than they were in 1912. This data set contains information on the exact years in which individual firms ceased to exist as independent entities. Fligstein (1990) supplies a data set on the top 100 US firms 1919-79, but this evidence relates to periods of a decade rather than to individual years.

Ormerod et.al. (2001) show that, for both these data sets, the empirical relationship between the frequency and size of extinctions is described well by a power law. The information in the Hannah data set also shows that the intervals between extinction events follows power law behaviour.

The purpose of this paper is to set out a theoretical model of agent extinction which is compatible with the empirical evidence. Section 2 describes the model. Section 3 discusses the nature of the empirical evidence and how the success of the model is judged. Section 4 presents the results of the model. Section 5 gives a brief conclusion.

2. A theoretical model of agent extinction

The model contains N agents, and all pairs of agents are connected to each other. These connections can be thought of as representing the way in which the net impacts of the overall strategies of firms impact on each other. Both the strength and the signs of the connections vary.

In formal terms, the connections are embodied in a matrix of couplings, J_{ij}, which indicates how each agent i affects every other agent j, with $J_{ij} \in [-1, 1]$. It is important to emphasise that the J_{ij} are not simply the cross-price elasticities which might be estimated between products in, say, a Nearly Ideal Demand System. They represent the net effect of a firm i's overall strategy on firm j, and not just the impact of relative price. Competition between agents, for example, is the broad concept noted by Vickers (1994) in the *New Palgrave Dictionary of Economics*, where it is defined as 'a rivalry between individuals (or groups or nations), which arises whenever two or more parties strive for something that all cannot obtain.' Price may certainly be an element in defining the value of the connection from agent i to agent j, but so is, for example, advertising, R and D and effort levels.

The overall fitness of an agent is measured by the sum of its connections to all other agents[1]. More exactly, it is the sum of influences on each agent of all other agents. Fitness in this context is fitness for survival, and is a wider concept than, for example, just volume of sales or profits. There are many examples in business history of very large firms with high levels of profits which have collapsed very rapidly due to drastic mistakes of strategy by the management[2].

Three combinations of pair-wise connections are possible in terms of the signs of the J_{ij} : i) J_{ij}, $J_{ji} > 0$; ii) $J_{ij} > 0$, $J_{ji} < 0$, or vice versa; and iii) J_{ij}, $J_{ji} < 0$

Case (i) represents a situation in which firms benefit from each other's presence in a market. The situation could arise through co-operation or tacit collusion. More generally,

[1] these include the connection of the product/firm to itself, as it were, the J_{ii}. A firm may possess qualities which lead to positive or negative effects on its own fitness. For example, a firm may attempt to occupy a niche for which, in any given period, the demand is very weak, and is therefore handicapped in its attempts to survive. The properties of the model are in any event not affected in any significant way if the J_{ii} are set equal to zero.

[2] Marconi in the UK is a recent example. Marconi is the old GEC company which existed for many years, and is far from being a flash-in-the-pan dot.com firm. Yet its stock market value fell from £35 billion to under £1 billion during the course of a single year because of appalling strategic errors by its management

the signs will be the same when two firms carry out activities which are complimentary to each other.

Case (ii) arises when two products are in competition, and the overall strategy of one is such that it gains fitness at the expense of its rival. Case (iii) is a more intense example of the competitive case (ii). In this instance, the degree of competition is such that the firms carry out actions which reduce both their fitness levels. An example is when two firms become engaged in a price war which ultimately reduces both their profit levels.

The connections between agents evolve over time. In other words, firms alter their strategies. We can think of each firm as attempting to maximise its overall fitness level. In the model, the firm proceeds by a process of trial-and-error in altering its strategy for any given product. The model is solved over a sequence of iterated steps, and at each step, for each agent one of its connections is chosen at random, and a new value is assigned to it.

This process is completely compatible with the conventional rationalisation of the maximisation hypothesis in orthodox economic theory. Agents are assumed on the one hand to maximise their individual utilities, yet on the other it is recognised that under conditions of uncertainty it is impossible for individual agents to follow maximising behaviour, because no one knows with certainty the outcome of a decision. The two views are reconciled, and maximisation is nevertheless deemed to occur, because it is argued that competition dictates that the more efficient firm will survive and the inefficient ones perish (the classic statement of this is Alchian (1950)).

An agent is deemed to become unable to survive if its overall fitness falls below zero. At any step in the solution of the model, more than one agent can become extinct. If m agents become extinct in any given step, an extinction of size m is defined to have taken place.

Agents which become extinct are replaced by new agents. In a version of this model developed to account for the patterns of extinction of biological species, Sole and

Manrubia (1996) postulate that the new entrants copy very closely surviving species. At each step in the solution of the model when an extinction has taken place, a surviving agent is chosen at random as the template for the new entrants. The connections of each new agent are the same as those of the template agent, except for a small random change in the value of each connection.

In this particular economic context, this replacement rule is not completely unreasonable. Firms sometimes do become very large by copying closely firms which are already very large. For example, firms occasionally acquire companies which are bigger than themselves. However, more usually, firms which grow sufficiently to enter the group of the world's largest companies often have distinctive qualities of their own. New sectors of the economy become important, such as financial services or computing.

The replacement rule we use reflects this factor. The connections of a new entrant are in the first instance chosen at random from the interval [-1, 1]. However, we distinguish the net impact of surviving agents on the new entrant, the J_{ji}, from the impact of the new entrant on the survivors, the J_{ij}. In the former case, a firm which has grown sufficiently to enter the set of the world's largest companies can be assumed to have at the time of entry a fitness level which is greater than zero. If the J_{ji} were simply chosen at random, the mean fitness of new entrants would be zero. We therefore add to each of these elements the mean fitness level of the surviving firms at the time of entry, divided by the total number of agents. In other words, the average fitness level of a new entrant will be equal to the average fitness level of surviving agents.

In the case of the impact of the new entrant on surviving agents, the J_{ij}, we simply assume that the overall impact has a mean value of zero.

In terms of a formal statement of the model, we have:

The model contains N agents and a matrix of couplings, J_{ij}, which indicates how each agent i affects every other agent j, with $J_{ij} \in [-1, 1]$. The model is solved over a sequence of iterated steps, and at each iteration the following occurs:

i) for each agent i, one of its J_{ij} is replaced with a new value chosen at random from a uniform distribution on [-1, 1].

ii) the overall fitness of any given agent is measured by $f_i = \sum_j J_{ji}$, and any agent for which $f_i < 0$ is deemed to be extinct. If m agent become extinct, an extinction of size m is deemed to have taken place.

iii) an extinct agent is replaced by a new entrant into the system. The connections of a new entrant are chosen at random from [-1, 1]. The mean fitness level of the surviving agents (divided by the number of agents) at the time of entry is then added to the interactions of other agents with the new entrant, the J_{ji}. The interactions of the new entrant with other agents, the J_{ij}, are simply chosen at random from [-1, 1].

3. Empirical evidence on the extinction patterns of the world's largest companies in 1912

The standard approach in the analysis of the extinction patterns of biological species (see Drossel (2001) for a detailed survey) is to fit a relationship between the number of firms which become extinct in any given year, and the frequency with which these are observed.. In other words, no attempt is made to replicate the actual time-series observed for extinctions. Instead, the focus is on the properties of the underlying distribution which could give rise to the historical realisation which is actually observed.

Figure 1 plots the frequency of annual rates of extinction from the Hannah data set. This relates to the experiences of the world's top 100 industrial countries in 1912 over the 1912-95 period. In most years, no single giant firm became extinct, but four firms became extinct in the year 1919, and no fewer than six in 1968.

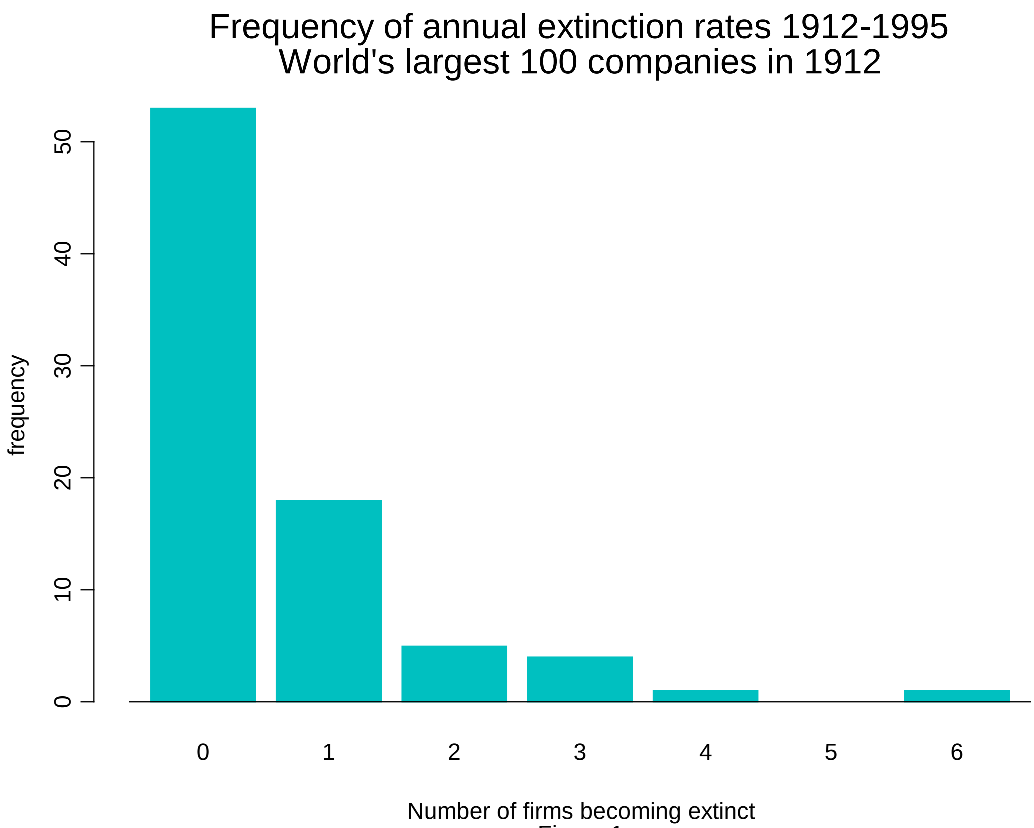

Figure 1

A power law of the form

$$F = \alpha . N^{\beta} \qquad (1)$$

describes the data well, where F is the frequency with which the annual number of extinctions is observed over the 1912-1995 period, and N is the annual number of extinctions.

A least squares[3] fit of (1) to the data (for N > 0) gives estimated values of α of 18.0 and of β of -1.76, the latter with a standard error of 0.18. The standard error

[3] using a non-linear least squares algorithm in S-Plus rather than the conventional log-log least squares fit, because there are no examples in the data of 5 firms becoming extinct in any single year and hence the dependent variable takes the value zero for this observation

of the equation is 0.94. Comparing this latter to the standard error of the data, 6.75, the equation fits the data well.

The evidence from the Fligstein data set, using the largest 100 companies in the US in 1919 over the 1919-79 period, also suggests very clearly that the frequency/size extinction relationship can be described well be a power law. But, mainly because the information is aggregated into periods of a decade, this data set is not nearly as informative in this context as the one provided by Hannah

A more detailed discussion of the empirical evidence is given in Ormerod et.al. (op.cit.).

The evidence for biological extinctions suggests, intriguingly, that a power law with an exponent of -2 provides a good description of the data. This is very close to the -1.76 fitted with the Hannah data set.

We also examined the Hannah data set regarding the number of years between an extinction of at least one of the top 100 industrial firms in 1912. The most frequent observation is one year, which means that in this case extinctions took place in successive years. A power law provides a reasonable fit to the data and, again, this is somewhat better than that given by an exponential distribution. The estimated exponent in the power law least squares fit is -1.18 with a standard error of 0.22. The overall fit, however, is not quite as good as that of the frequency data. Again, more details are given in Ormerod et.al. (op.cit.).

4 Properties of the theoretical model

The initial task in analysing the model is choosing the number of agents with which to populate the model. In the data set provided by Hannah, by 1995 no fewer than 48 of the world's top 100 industrial companies in 1912 had disappeared in any independent form. Ideally, the dates at which they ceased to be members of the top 100 would be available, and extinction could be defined as exit from the top 100. We could then set N = 100, and note the extinctions of the original set of agents. However, this information is not available, and an alternative approach is needed.

We populate the model with 500 agents, which we assume are all very large companies, although size is not an explicit factor in the model. It does not seem unreasonable to assume that the very largest 100 companies derive their fitness levels from their dealings and interactions with a population of these 100 plus 400 other very large companies.. Of course, they will all be involved with many firms of many different sizes. But a local company, say, with the contract to clean the headquarter offices of a major oil company is unlikely to have any discernible effect on the ability of that firm to survive.

With just one or two exceptions, the surviving companies from the 1912 top 100 still operated as substantial companies in 1995, so again it is not unreasonable to assume that even when they dropped out of the top 100 they continued to interact with the survivors plus the other large firms which populate our model. Extinction is therefore defined as dropping out of the set of 500 very large companies. This is not strictly compatible with the empirical data set, but it is a good approximation to it.

Because of the stochastic nature of the model, repeated solutions are required in order to establish its properties. We report results obtained from 500 separate solutions. Each time, the first 10,000 iterations are discarded in order to eliminate any transient behaviour arising from the choice of the initial J_{ij}[4].

In each solution, 100 of the initial 500 agents are chosen at random and designated as the largest 100 firms in the total population. It is their extinction

[4] Experiments both with this model and variants of it suggest that in fact, empirically, a far smaller number need to be eliminated in order to achieve this end

patterns which are monitored, and the solution is halted whenever 50 of them become extinct. We therefore have evidence from 500 separate solutions of the model of the extinction patterns of 50 out of the 100 companies.

Figure 2 plots the relationship between the frequency with which extinctions of different sizes are observed, and the size of the extinction. The slope of the least-squares fit is -1.83 with a standard error of 0.17. The slope is very similar to that which is estimated from the actual data on large firm extinctions.

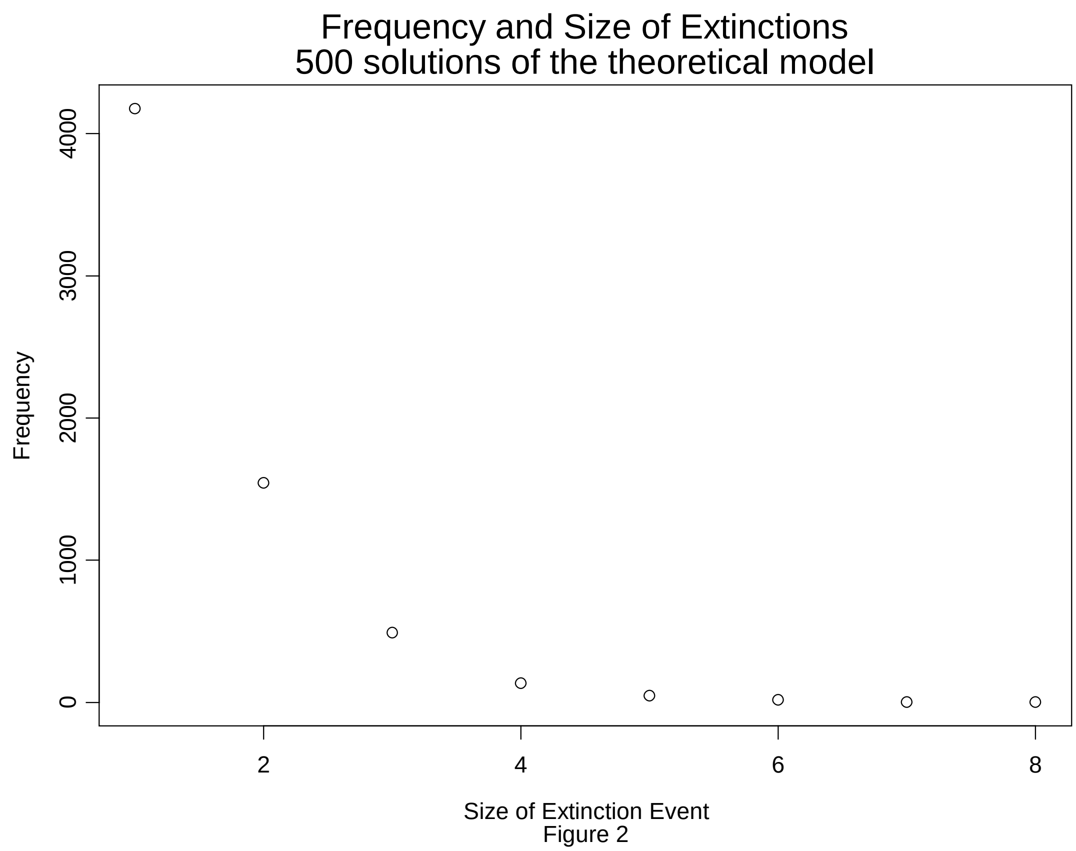

Figure 2: plot showing the power law relationship between extinction size and frequency. The frequency data show the number of times over the 500 solutions of the model in which the particular number of agents become extinct in any given period

We also analysed the number of periods between extinctions of at least one of the designated top 100 firms. Again, this was carried out using the output of 500 solutions of the model as described above for the frequency data. The solutions of the model generate a large number of data points, with for example 9022 occasions when extinctions were observed in successive periods. There is a long, sparse tail in the model-generated data, with a single observation of a gap of 26 periods being recorded.

The exponent of the power law of the actual data is -1.18 ± 0.22. The power law fitted to the model-generated data gives an estimated exponent of -1.78 ± 0.04, which is significantly different from the actual data at the conventional level of $p = 0.05$. In the case of the model-generated data an exponential distribution fits as well as a power law, but the estimated exponent in this regression is again significantly different from that estimated with the actual data at $p = 0.05$. With the model-generated data, the power law tends to over-predict the number of large gaps between extinction events, and the exponential distribution under-predicts them. But both describe the bulk of the data well.

So, in terms of the waiting times between extinction events, a power law relationship provided a good description of both the actual and the model-generated data, although the exact quantitative nature of the differs somewhat.

5. Conclusion

We develop in this paper a theoretical model of agent evolution and extinction based upon straightforward principles of economics. The model is similar though not identical to models of extinction in the biological literature.

We consider evidence from a data set containing information on the world's 100 largest industrial companies in 1912. We also consider a data set of the top 100 US firms over the 1919 - 1979 period.

The relationship between the frequency and size of the extinctions on an annual basis is approximated well by a power law relationship. The exponent of the fitted power law is very similar to that reported in the literature on the extinction of biological species in the fossil record. Further, the gaps between extinction events can also be described well be a power law.

The relationship between the frequency and size of extinctions generated by the model is very similar to that which is observed in the actual data. The model-generated data on gaps between extinction events is also approximated by a power law, though the slope of the relationship is somewhat greater than that of the actual data. The paper raises the possibility that there are general mechanisms at work which account for the extinctions of agents.

References

A.Alchian (1950), 'Uncertainty, Evolution and Economic Theory', *Journal of Political Economy*, **LIX**

A.D.Chandler (1990), *Scale and Scope: the Dynamics of Industrial Capitalism*, Harvard University Press

B.Drossell (2001), 'Biological evolution and statistical physics', cond.mat/ 0101409, forthcoming in *Advances in Physics*

L.Hannah (1999) 'Marshall's "Trees" and the Global "Forest": Were "Giant Redwoods" Different?' in N.R.Lamoreaux, D.M.G.Raff and P.Temin, eds., *Learning by doing in markets, firms and countries*, National Bureau of Economic Research

P.Ormerod, H.Johns and L.Smith (2001), "Marshall's "Trees" and the Global "Forest": the Extinction Patterns of Capitalism's Largest Firms", www.volterra.co.uk

R.V. Solé and S.C.Manrubia (1996), 'Extinction and self-organised criticality in a model of large-scale evolution', *Phys. Rev. E*, **54**, R42-R45

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