How Much Can Firms Know?

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

Bridget Rosewell (brosewell@volterra.co.uk)

Volterra Consulting Ltd

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

February 2004

* 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

Two key stylised facts have been established about the extinction patterns of firms. First, the probability of extinction is highest at the start of the firm's existence, but soon becomes more or less invariant to the age of the firm. Second, a very recent finding, that the relationship between the size and frequency of firm extinctions is closely approximated by a power law.

Economic theory has a great deal to say about many aspects of firm behaviour, but offers relatively little on the deaths of firms. An agent based model of firm evolution and extinction has been developed which has properties which conform closely to the stylised facts. However, in this model, firms are unable to acquire knowledge about either the true impact of other firms' strategies on its own fitness, or the true impact of changes to its own strategies on its fitness.

We examine the effects of allowing firms different amounts of knowledge about the effects of strategy in the context of the agent-based evolutionary model. We investigate the emergent properties of the system in terms of the relationship between the size and frequency of firm extinction, and the survival patterns with respect to the age of agents. In other words, we compare the properties of the model with the two key stylised facts.

There are very considerable returns in the model to acquiring knowledge. There is a sharp increase in the mean agent age at extinction for agents with even a small amount of knowledge compared to those without. Indeed, we find that as both the amount of knowledge available to firms increases and as the number of firms capable of acquiring such knowledge rises, the lifespan of agents begins to approach the limiting, full information paradigm of neo-classical theory in which agents live for ever.

However, even with relatively low levels of knowledge and numbers of agents capable of acquiring it, the model ceases to have properties which are compatible with the two key stylised facts on firm extinctions. The clear implication is that firms have very limited capacities to acquire knowledge about the true impact of their strategies.

1. Introduction

Economic theory has a great deal to say about many aspects of firm behaviour, but offers relatively little on the deaths of firms. For example, in the context of general equilibrium theory, in an n-period world agents act as if they have perfect foresight about the best strategy to follow in any future state of the world [1]. The proof of existence of equilibrium requires that each agent has an infinite amount of knowledge. In such a world, an agent such as the firm never disappears.

Slightly more realistically, a firm could be assumed to have perfect information up to and including the current period, and to act as a rational maximiser in determining its strategy. However, a firm could now become extinct because of its failure to anticipate future random shocks. Certainly, empirical studies of firm growth and size have emphasised the importance of stochastic influences, from the initial work of Gibrat [2], through classic papers of the 1950s and 1960s (for example, [3,4]), to much more recent contributions [5,6]. But the literature of the deaths of firms is surprisingly sparse. Some empirical studies relate firm deaths within individual industries to factors such as the number of firms in the industry at the time the firm enters ([7], for example, provides examples and cites others). But no general analysis of deaths appears to be available.

There do, however, appear to be two important stylised facts about firm extinctions. First, it has been known for some time (for example, [7,8]) that the probability of extinction is highest in the early life of a firm, but declines rapidly and is thereafter more or less invariant with respect to the lifespan of the firm.

Second, it has been shown recently, that the empirical relationship between the frequency and size of firm extinctions is described well by a power law. Analyses of a database of 6 million US firms by state and industry in the 1980s and 1990s [9] and of a database of eight OECD countries over the 1977-99 period [10] both give estimates of the exponent of the power law of –2. Over a longer time-scale, analysis of the experiences of the largest 100 industrial firms in the world in 1912 over the 1912-95 period gives an estimate of –1.75 [11]. The empirical relationship is well-grounded.

An agent-based model of the evolution and extinction of firms has been developed [10] which has properties which conform closely to the empirical evidence exhibited in these two stylised facts, namely:

- the relationship between the size and frequency of firm extinctions is closely approximated by a power law
- the probability of extinction is highest at the start of the firm's existence [illegible]ut soon becomes more or less invariant to the age of the firm

In this paper, we examine the effects of allowing firms different amounts of knowledge about the effects of strategy in the context of the agent-based evolutionary model. We investigate the emergent properties of the system in terms of the relationship between the size and frequency of firm extinction, and the survival patterns with respect to the age of agents. In other words, we compare the properties of the model with two key stylised facts.

We obtain a permissible set, as it were, of the different amounts of knowledge which agents are assumed to possess, and which generate properties which are compatible with the stylised facts on firm extinction. Assumptions about the amount of knowledge available to agents which lie outside this set lead to the model being no longer able to replicate the stylised facts.

Section 2 describes the theoretical agent-based model, and section 3 discusses how knowledge and intent are introduced into agent behaviour. Section 4 sets out the results, and section 5 gives a conclusion.

2. The theoretical model of the evolution and extinction of firms

2.1 Overview

The model contains N agents, and every agent is connected to every other. 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. The overall properties of the model emerge from the interactions between agents.

The connections between agents 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. Each firm can be thought of 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 with respect to other firms. 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.

2.2 The theoretical model: a more formal statement

The model consists of a population of N agents which are considered to influence each other via a matrix of uniformly distributed interconnections $J_{ij} \in [-1, 1]$, where J_{ij} is the effect of agent i upon agent j.

The model can be summarised as follows.

1. J_{ij} is initialised by drawing at random from a uniform distribution on [-1, 1]

2. Each agent i has one of its J_{ji} updated, i.e. assigned with a new value chosen at random in the interval [-1, 1]

3. The fitness $F_i(t)$ of each agent is calculated, where $F_i(t) = \sum_{j=1}^{N} J_{ji}(t)$

4. If $F_i(t) < 0$ then the agent is deemed extinct and its interconnections are obsolete, i.e. $J_{ij} = 0$ and $J_{ji} = 0 \ \forall j$.

5. Extinct agents are then replaced. A random "parent" k is chosen from the surviving agents. Each replacement agent i is assigned new connections such that :

   $J_{ij} = J_{kj} + \varepsilon_{ij}$ and $J_{ji} = J_{jk} + \varepsilon_{ji}$, where ε_{ij} is a small random number drawn from a uniform distribution $[-\varepsilon_{\max}, \varepsilon_{\max}]$.

6. Steps 2 to 5 are repeated for n iterations.

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 [see, for example, 13]. 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 in [14], 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

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) $J_{ij} ,J_{ji} < 0$

---

[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.

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, the signs will be the same when two firms carry out activities which are complimentary to each other. From a theoretical perspective, Chamberlin [15] argued, in the context of his theory of monopolistic competition, that 'As for the conventional categories of industries, it seems increasingly evident to me that they have their origin, not primarily in substitution at all but in similarity...'.

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.

The properties of the model are in general invariant to the introduction of a considerable degree of sparseness into the connection matrix, J_{ij} [16]. Further, the introduction of external shocks which are either common to the fitness of all agents or specific to individual agents does not alter the properties of the model, provided that the scale of the shocks is not such as to dominate the model [17].

2.3 Properties of the model

Figure 1 shows the typical relationship between the frequency and the size of extinctions which emerges from the model.

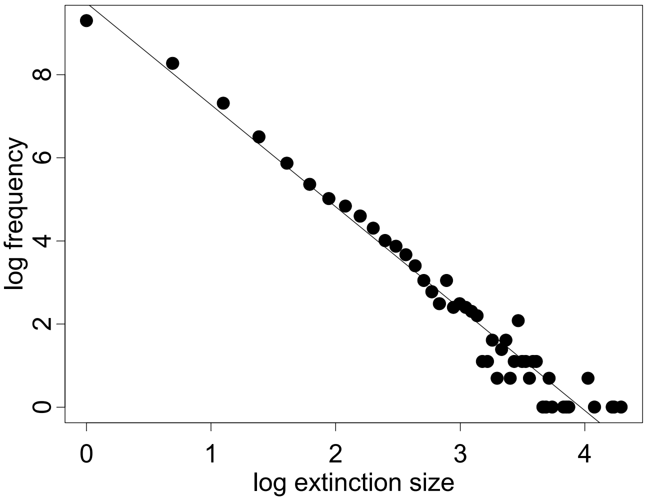

Figure 1 *Typical power law extinction size/frequency relationship from a singe 50,000 iteration run of the standard model.*

The average estimated slope of the power law over a large number of individual solutions of the model is –2.43.

Figure 2 shows the average relationship over 20 typical solutions of the basic model between the probability of extinction and the age of the agent.

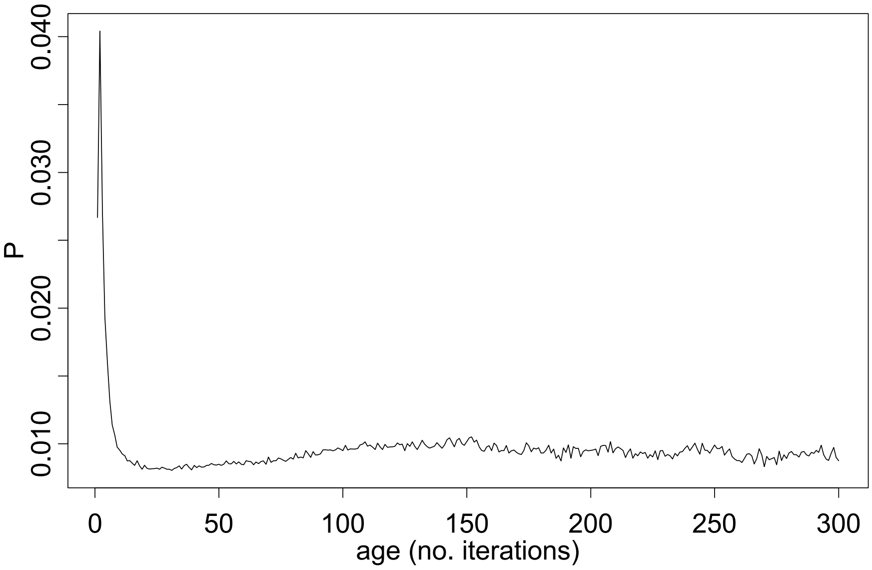

Figure 2 *the number of agents which become extinct at iteration* t *as a proportion of the number of agents who reach age* t, *a proxy for extinction probability, vs. no. of iterations.*

3. Knowledge and intent in agent behaviour

There are two separate aspects in the model which relate to the extent of the firm's knowledge about the external environment in which it operates:

- the amount of knowledge which the firm has about the impact of the strategies of other firms on its own fitness at any point in time
- the amount of knowledge which the firm has about the impact of changing its own strategy in order to alter the effect of the strategies of other firms on it.

In the model described above, no matter how carefully an issue is examined, no matter how experienced a firm, the impact of any change which it makes to its strategy cannot be known in advance.

The existing update rule for the effect of a firm's strategy is that each agent has one of its J_{ji} updated, i.e. assigned with a new value chosen at random in the interval [-1, 1]. In other words, the effects of a firm's strategy are updated by a random draw.

We introduce two separate ways in which a firm is permitted to have knowledge and to act with intent:

- a parameter ε , $\varepsilon > 0$, is added to the value chosen at random in the uniform interval [-1, 1] (with the condition that $\left|J_{ji}\right| \leq 1$)
- the updated value is chosen at random from [-1, 1], but the agent is allowed to select the particular J_{ji} which is updated, rather than it being chosen at random

In the general version of the model described in section 2, the expected value of each updated J_{ji} is of course zero. The first variant of the model permits the firm to exercise some control over the impact of changing its strategy by setting the expected value equal to ε, $\varepsilon > 0$[2].

---

[2] subject to the constraint that the maximum value for any J_{ij} is 1

In other words, the firm does not have knowledge about which of the J_{ji} to update in terms of its own best interests. However, it is able to use its knowledge of the market to choose a change in strategy which on average will improve its overall fitness. The higher the value of ε, the greater the degree of knowledge and control which an agent is able to exercise over the outcome of its strategy.

Of course, the firm may 'choose' to update a particular J_{ji} which is close to 1, and so only have a very small probability of improving its fitness. Equally, the basic draw from [-1, 1] may allocate a value close to –1, so that even with the addition of ε there is still a negative impact on fitness. But on average, its fitness will be improved.

The second variant allows the firm to select the particular J_{ji} which is updated. In fact, we allow firms to select the one which is closest to –1. In other words, the one which has the largest adverse impact on the fitness of the firm at that point in time. The updated value is chosen from [-1, 1], so the firm cannot act with knowledge as it is able to do in the first variant of the model. However, it has more knowledge in this version in the sense that it knows the current effect of the impact of its existing strategies on its fitness.

4. The results

We populate the model with 100 agents. The size/frequency extinction relationship which emerges with this number of agents is very similar to that with large numbers of agents. Each version of the model is solved 500 times over 50,000 steps (with the first 10,000 iterations disregarded), and the average of the results across the 500 solutions is reported.

In each case, we investigate the effect of allowing different numbers of agents to use the rules which give them more knowledge.

We present results on:

- the exponent of a fitted extinction size/frequency power law
- the relationship between the probability of extinction and the age of the agent and the mean agent age at extinction
- the mean number of extinctions per 10,000 iterations of agents incapable of intent, i.e. using the purely random strategy update rule
- the mean number of extinctions per 10,000 iterations of agents capable of intent

The variant in which agents are allowed to identify and update the individual connection which impacts most negatively on their fitness can be described very briefly. Essentially, in this variant of the model, extinctions soon become extremely rare, and the average fitness of the agents approaches the maximum possible value of 100.

Intuitively, the reasons for this are as follows. First, suppose that no extinctions take place. The minimum value of the J_{ji} for each agent is always the one which is updated. If the new value is still the lowest for any given agent, it will be redrawn again on the next iteration, until it is greater in value than the J_{ji} which was originally the second lowest value for that agent. Eventually, all the J_{ji} will take values greater than the one which was originally the highest, and then it too will be updated. It is easy to see that in the absence of extinctions, all the J_{ji} will in the limit approach the value of 1.

Second, consider the mean nt and the initial conditions of its J_{ji}. These will be distributed uniformly in the interval [-1, 1], with an average value of zero. In other words, the lowest value of the J_{ji} will be $-\omega$, a value close to -1. So, in the process of updating, it is very unlikely that a value will be chosen which is lower than $-\omega$, thereby bringing about the extinction of the agent. In fact, the mean value of the updated J_{ji} will be zero, so the overall fitness of the agent will on average rise from zero to $+\omega$. In the next update, the lowest of the J_{ji} for the agent will still be $-\omega$, so even if a value is drawn which is lower, the agent will not become extinct. In fact, the expected value of this update is also zero, so that the overall fitness of the agent will be close to $+2\omega$, and so on.

The main results are therefore concerned with the rule which enables agents to update strategy in a way in which any updated connection has an expected value greater than zero.

We examine the model under a variety of choices of ε and of the number of agents for which $\varepsilon > 0$. For purposes of description, we designate the latter as N_C, signifying the number of agents which is capable of exercising intent.

As a benchmark, with the version of the model described in section 2, the exponent is 2.43 ([illegible] absolute value), the mean agent age at extinction is 99 (i.e 99 iterations of the model), the mean agent fitness is 11.7, and the mean number of extinctions per 10,000 iterations is 100.

Table 1 shows the absolute value of the estimated exponent, β, of the log-log least squares regression of the size and frequency of extinctions.

Table 1

Exponent of extinction size/frequency power law

ε / N_C	**1**	**10**	**25**	**50**	**75**	**100**
0.5	2.45	2.88	3.89	5.02	NEO	NA
0.35	2.44	2.76	3.66	4.94	NEO	NA
0.2	2.42	2.54	2.98	4.57	5.44	NEO
0.15	2.42	2.49	2.70	3.77	5.14	NEO
0.1	2.41	2.45	2.51	2.76	3.60	5.03
0.075	2.42	2.44	2.48	2.58	2.82	3.44
0.05	2.42	2.43	2.45	2.49	2.54	2.64
0.01	2.41	2.42	2.42	2.42	2.43	2.43

Note: The left hand column shows the value of ε and the top row shows the number of agents capable of intent ($\varepsilon > 0$). "NA" means that no extinctions at all are observed, and "NEO" means that not enough extinction observations occur to attempt to fit a power law.

For $\varepsilon = 0.01$, the power law relationship between size and frequency remains virtually the same as in the model in which agents have no knowledge at all on the impact of strategy changes. Similarly, if just a single agent has this capability, the power law is not affected regardless of the value of ε.

However, as the table shows quite clearly, as ε and N_C are increased, the exponent generated by the model begins to deviate from its actual value. Further, particularly for the combinations of ε and N_C in the upper right hand part of Table 1, a power law ceases to give a reasonable description of the size-frequency relationship. The behaviour of the model when most or all agents have knowledge is particularly erratic, and, indeed, at high values of both as ε and N_C, extinctions effectively disappear from the model entirely.

Figures 3a-e show the relationship between the probability of survival and the age of the agent for a range of values of ε. Again, compared to the standard model, the ability of the model to replicate the stylised facts about this relationship breaks down as the value of ε increases.

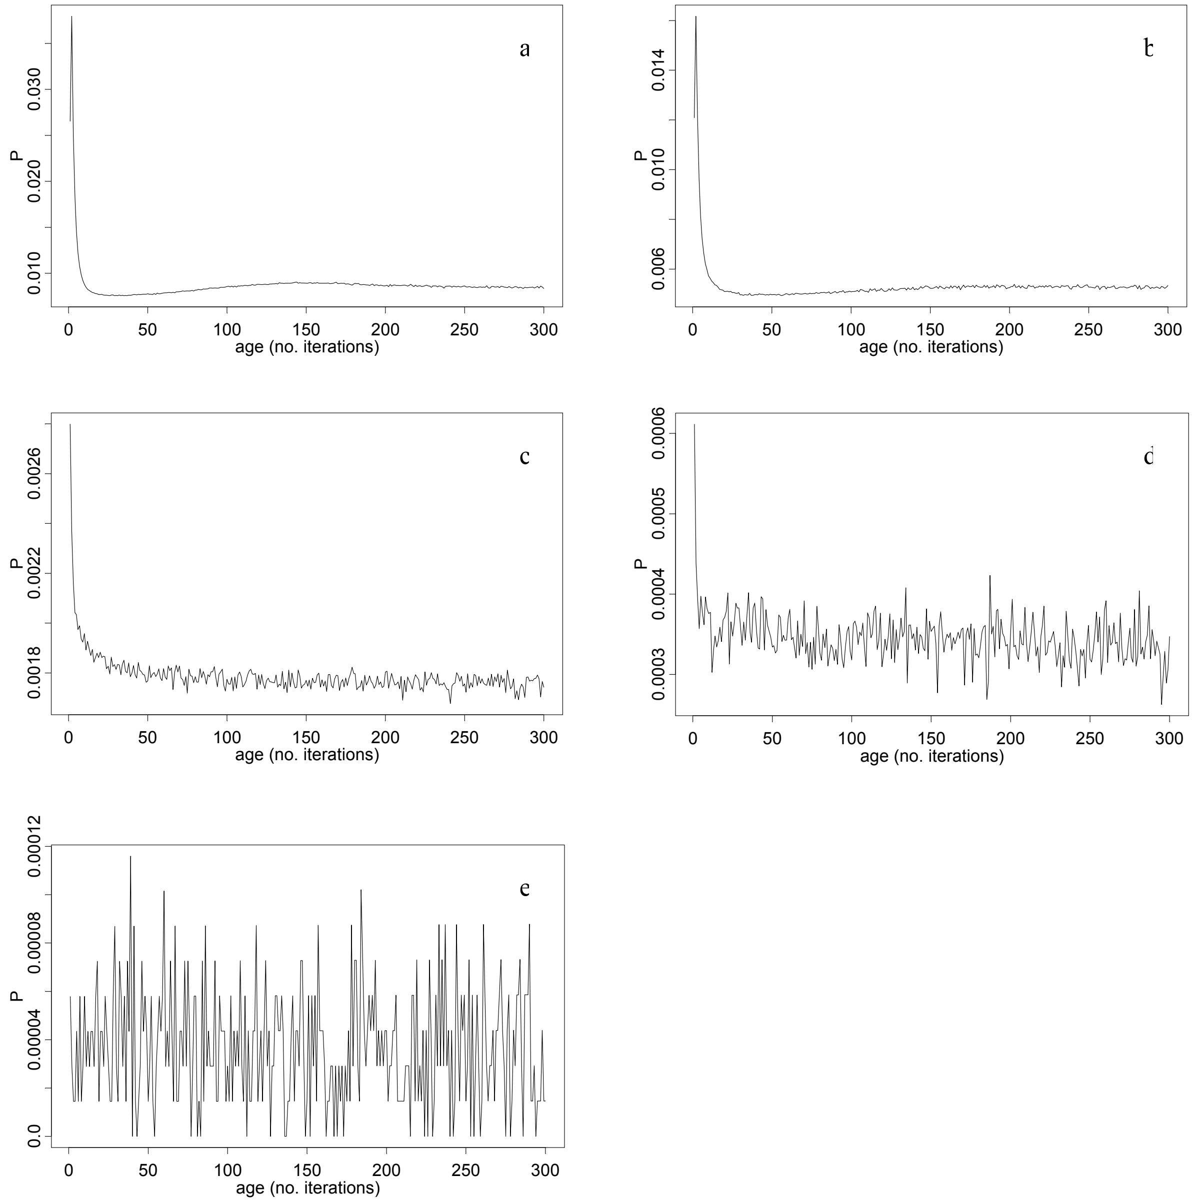

Figure 3 *Plots showing P (the number of agents with intent who become extinct at iteration* t *as a proportion of the number of agents with intent who reach age* t*), a proxy for extinction probability, vs. no. iterations for varying* ε*. (a)* $\varepsilon = 0.01$ *(b)* $\varepsilon = 0.05$ *(c)* $\varepsilon = 0.1$ *(d)* $\varepsilon = 0.15$ *(e)* $\varepsilon = 0.2$. $N_C = 100$ *for all. This becomes less well-defined as fewer and fewer agents with intent actually become extinct.*

Tables 2a and 2b demonstrate the high potential returns available to agents capable of acquiring knowledge about the impact of their strategies. Table 2a shows the mean age (in terms of numbers of iterations of the model) to extinction of agents without this capability, and Table 2b the mean age of agents with this capability.

Table 2a

Mean agent age at extinction for agents incapable of intent

ε / N_C	**1**	**10**	**25**	**50**	**75**	**100**
0.5	102.10	125.09	157.36	200.85	238.39	NA
0.35	101.11	117.18	142.47	179.20	211.65	NA
0.2	100.17	107.35	120.58	144.39	167.71	NA
0.15	99.99	104.69	113.48	129.82	147.42	NA
0.1	99.56	102.58	107.56	116.90	127.38	NA
0.075	99.69	101.79	105.36	111.89	118.90	NA
0.05	99.54	101.20	103.16	107.36	111.64	NA
0.01	99.21	99.78	100.30	100.96	101.65	NA

Note: when N_C *equals 100, by definition there are no agents incapable of intent*

With no capability of acquiring knowledge, the model gives a mean age at extinction of 99.

The mean age of agents incapable of acquiring knowledge rises as the number of agents which *are* capable, for the following reason. The 'parenting' rule in the model describes how the connections of new agents which replace those which become extinct are determined. Essentially, a new agent copies the connections of a surviving one, with a small variation around them. Now, the average fitness of agents capable of intent rises as ε and N_C increase, and so every new agent, whether capable of exercising intent or not, begins with an expected fitness level which is higher, and so takes longer to become extinct.

Table 2a, however, is in dramatic contrast with Table 2b.

Table 2b

Mean agent age at extinction for agents capable of intent

ε / N_C	**1**	**10**	**25**	**50**	**75**	**100**
0.5	6,368	21,691	23,392	NA	NA	NA
0.35	2,076	7,178	26,477	29,393	23,930	NA
0.2	540	753	1,630	5,710	11,852	17,783
0.15	333	391	556	1,087	1,898	2,973
0.1	212	227	257	333	438	563
0.075	172	179	193	221	257	301
0.05	141	145	150	160	172	185
0.01	106	107	107	108	109	110

Qualitatively, as ε and N_C are increased, the model begins to approach the limiting, full information paradigm of neo-classical theory in which agents live for ever. And the returns to becoming capable of acquiring even a small amount of knowledge are very high. The mean age at extinction when $\varepsilon = 0.01$ varies between 106 and 110, compared to a value of 99 when $\varepsilon = 0$ for all agents. And when $\varepsilon = 0.05$, the mean age varies between 141 and 185.

Overall, it is clear that the model does not permit a great deal of knowledge acquisition before its properties begin to deviate from the stylised facts generated by the behaviour of actual firms. The model properties continue to conform to the stylised facts if either:

- all agents are capable of acquiring a small amount of knowledge
- *or* a very small proportion of agents can acquire considerable amounts

Otherwise, the properties of the model deviate from those of reality.

5. Conclusion

There are two important stylised facts relating to firm survival and extinction:

- the size/frequency of extinction relationship follows a power law with exponent around –2
- the probability of extinction is highest in the early time period of a firm's life. This falls as the firm becomes older, and eventually becomes generally invariant to the age of a firm

An agent-based model of the evolution and extinction of firms based on basic principles of economics which is similar to, though not identical to, models in the biological literature has properties which conform closely to the empirical evidence.

Firms are connected by a matrix whose elements measure the net impact of a given firm's strategy on the fitness of another particular firm. Firms alter their strategy over time. However, they have no explicit knowledge at any point in time either of the true impact on their fitness of either their own or other firms' current strategies. Further, they have no knowledge of the effect which changing their strategy will have.

Their cognitive ability is in very sharp contrast to the explicit, rational maximising agent of much of economics. Nevertheless, agents acting in this way produce overall outcomes which are compatible with the stylised facts on firm extinctions.

We allow firms to obtain a limited amount of knowledge, and to be able to act with a certain amount of intent. However, their cognitive ability is still very much less than that of an explicit rational maximiser.

There are very considerable returns to acquiring knowledge, for even a small amount leads to a sharp increase in the mean agent age at extinction for agents with knowledge compared to those without. Indeed, we find that as both the amount of knowledge available to firms increases and as the number of firms capable of acquiring such knowledge rises, the lifespan of agents begins to approach the limiting, full information paradigm of neo-classical theory in which agents live for ever.

However, even with relatively low levels of knowledge and numbers of agents capable of acquiring it, the model ceases to have properties which are compatible with the two key stylised facts on firm extinctions. The clear implication is that firms have very limited capacities to acquire knowledge about the likely impact of their strategies.

References

1. R.Radner, 1968, 'Competitive Equilibrium Under Uncertainty', *Econometrica*, 36
2. R.Gibrat ,1931, 'Les Inégalitiés Economiques', Librarie du Receuil Sirey, Paris
3. I.G.Adelman ,1958, 'A Stochastic Analysis of the Size Distribution of Firms', *American Economic Review*, 53, 893-904
4. P.E.Hart ,1962, 'The Size and Growth of Firms', *Economica*, 24, 29-39
5. L.A.N Amaral, S.V.Buldyrev, S.Havlin, H.Leschorn, P.Maass, M.A.Salinger, H.E.Stanley and M.H.R.Stanley 1997, 'Scaling Behaviour in Economics: I. Empirical Results for Company Growth', *J. Phys I France*, **7**, 621-633
6. R.L.Axtell ,2001, 'Zipf Distribution of US Firm Sizes', *Science,* 7 September
7. G.R.Carroll and M.T.Hannan, 2000, *The Demography of Corporations and Industries*, Princeton University Press
8. G.Dosi
9. 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
10. W.Cook and P.Ormerod, 2003, ‘Power law distribution of the frequency of demises of US firms’, *Physica A*, 324, 207-212
11. C.DiGuilmi, M.Gallegati and P.Ormerod, 2003, ‘Scaling invariant distribution of firms’ exit in OECD countries’, Dept of Economics, University of Ancona, Italy
12. H.Johns, P.Ormerod and L.Smith, 2003, ‘An agent-based model of the extinction patterns of capitalism’s largest firms’, paper presented at the World Econophysics conference, Bali, August 2002, and *Physica A*, forthcoming
13. A.Deaton and J.Muellbauer, 1980, ‘Almost ideal demand systems’, *American Economic Review*
14. J.Vickers ,1994, *Concepts of Competition*, inaugural lecture in the University of Oxford, Clarendon Press, Oxford
15. E.H.Chamberlin ,1957, *Towards a More General Theory of Value*, chapter 5, Oxford University Press
16.

17. P.Ormerod, H.Johns and L.Smith, 2002, 'External Shocks in a Model of Endogenous Agent Extinctions', paper prepared for the Institute of Complex Additive Systems Analysis, New Mexico, and available at www.volterra.co.uk