DRAFT FINAL REPORT

A Complex Systems Approach to the Extinction of Firms

Volterra Consulting Ltd.

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

UK

August 2002

tel: +44 (0) 20 8878 6333

e-mail: pormerod@volterra.co.uk

Contents

			Page
Summary			3
1.	Introduction		5
2.	Empirical findings		6
3.	The theoretical models		
	3.1	Overview	9
	3.2	The exogenous shock model	10
	3.3	The endogenous extinction model	12
4.	Future work		
	4.1	Towards a deeper understanding of agent-based evolution and extinction: extensions of the present modelling approach	14
	4.2	A particular model of evolution and extinction of agents connected on a scale-free network	16
	4.3	Deliverables	19
	4.4	Estimate of costs	20
References			21
Appendix 1			22
Appendix 2			32
Appendix 3			40
Appendix 4			58
Appendix 5			76
Appendix 6			90

Summary

- There is a growing number of examples of power law distributions in both the social and biological sciences. In particular, a power law provides a reasonable description of the relationship between the frequency and size of the extinctions of biological species in the fossil record.
- We have obtained empirical evidence that the frequency of extinction rates of companies – of companies ceasing to exist as independent entities - can also be approximated by a power law.
- The finding is obtained both with a database of the world's largest 100 companies in 1912 and their subsequent experience through to 1995, and with a database of some 6 million US companies of all sizes subdivided by industry and state in the 1980s and 1990s.
- The slope of this power law is almost identical in value to that estimated for the frequency of extinction rates of biological species.
- It is not obvious that systems containing either firms or biological species have any similarities other than that they are complex systems made up of interacting agents.
- Our findings raise the possibility that a more general mechanism of evolution and extinction is responsible for the extinction rates of both firms and of biological species.
- Two quite distinct types of theoretical model are capable of producing results which are compatible with the empirical evidence on biological extinction. An important distinction in the biological context is made between models in which extinction arises because of external environmental stress (eg: asteroid impact) and those in which it arises internally because of competitive interactions between species.
- Economics offers no satisfactory account of why firms cease to exist, even though the 'death' rate amongst modern US firms amounts to more than 10 per cent of the total population each year.
- But a distinction between external and internal causes is made more generally in economic theory, although the particular jargon which describes the two situations is 'exogenous', or outside the system, and 'endogenous', or within the system.

- The two types of theoretical models in the biological literature can be given interpretations which are entirely plausible in terms of economic theory (with relatively minor modifications).
- However, we find that the 'external shock' model of extinction cannot *both* be given a sensible economic interpretation and at the same time generate behaviour which is compatible with the empirical evidence.
- In contrast, the 'endogenous' model has both a plausible economic interpretation and can generate behaviour which is consistent with the evidence on extinctions.
- This finding is robust with respect to a range of modifications to the model.
- External shocks can be applied to the endogenous model, but provided the magnitude is not too great relative to the strength of the interactions between agents, the results are not affected.
- The results therefore suggest that the extinction patterns of firms are intrinsic to the system, arising from the interconnections between them. Shocks to the external environment may well play a role, but they are not the primary cause.
- The endogenous model performs well and has a plausible economic interpretation. But it is highly simplified. We therefore propose to investigate a number of extensions of the model in order to make it more realistic.
- In particular, the assumption that all agents are directly connected to each other is not realistic. In practice, the direct connections between economic agents involve a degree of sparsity.
- A scale-free network which connects agents both introduces sparsity and appears plausible in a model of connected firms.
- Introducing scale-free networks into the model extends the potential range of applications to a wide variety of social networks which connect individuals – an increasing body of evidence points to the relevance of scale-free networks in describing connections between individuals.
- We propose to analyse a particular example of the evolution and extinction of agents in a scale-free network which connects them, namely, that of a criminal gang. The overall severity of crime which emerges from the model, and the impact of intervention by the authorities can be examined.
- Details of the proposed extensions are set out in Section 4. The cost is $89,350.

1. Introduction

Economic theory contains a very large amount of material on firm behaviour, but little on why they might fail. Most of the literature available looks at why the number of firms in an industry might change in the face of a downturn. The Salter [1] approach is still widely used, which derives a model in which failure depends on the efficiency of different vintages of capital. There is a widespread recognition that one reason for firm extinction, namely mergers, does not occur at random but in waves. A certain amount of statistical time-series analysis exists on this (for example, [2]), and there is a large literature which discusses individual examples of merger. But no satisfactory theory explaining this phenomenon exists.

Corporate demographers have in recent decades carried out a substantial amount of empirical work on the mortality rate of firms. The recent book by Carroll and Hannan [3] contains an exposition of much of this material (including the distinguished contributions made by the two authors), and a large number of references.

This work, however, is concerned almost exclusively with the death of firms within individual industries, a large number of which have been considered. In this context, firms are all in competition with each other to some degree. This is reflected in the empirical curve fitting exercises to data which have been undertaken. The probability of firm mortality is related strongly to the number of firms within the main industry in which the firm operates. But the potential symbiotic influences between firms across industries is not really contained in this approach.

An interesting perspective on the evolution and extinction of agents is provided in the biological literature on the frequency with which species become extinct in the fossil record. Empirical analysis suggests that a power law with an exponent of some –2 provides a reasonable empirical description of the relationship between the frequency and size of species extinctions. Several theoretical models of the evolution and extinction of species ('agents') have been constructed which are capable of generating a similar power law.

There are two main aims of this project:

- to carry out empirical investigation of firm extinctions to investigate the extent to which the relationship between the frequency and size of firm extinctions is similar to that observed in the fossil record of the extinction of biological species

- to investigate the extent to which a variety of models proposed to explain the extinctions of biological species can be used to give a satisfactory explanation of extinctions of firms. An important aspect of this is to consider the extent to which models developed in the biological context can be given sensible economic interpretations

In addition to the specific outputs, the project therefore makes contributions to:

- theoretical and applied economics
- complex systems theory

The findings of the project are discussed in detail in six papers attached as Appendices (1) – (6). The text of this report discusses the main results. Section 2 considers the applied analysis, and section 3 the theoretical models. Section 4 makes recommendations for future work, along with provisional estimates of costs.

2. Empirical findings

A number of studies have examined the relationship between the frequency and size of biological extinctions in the fossil record, and a recent survey is given by Drossel in [4]. The evidence is consistent with the existence of a power law describing biological extinction events, with an exponent estimated to be around -2.

A detailed database exists which considers the experiences of the world's largest 100 industrial companies in 1912 [5]. It charts their progress through to 1995, noting the years in which individual firms ceased to exist as independent entities. These companies, by definition, were extremely large even by the standards of today. US Steel, for example, employed 221,000 workers, and most of the others employed more than 10,000. Yet the overall attrition rate has been rather high, with only 52 of the 100 firms surviving in any independent form in 1995, and only 19 of them still being in the top 100 industrial companies.

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[1] fit of (1) to the data (for N > 0) gives estimated values α = 18.0 and β = -1.76, the latter with a standard error of 0.18. The standard error 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 estimated exponent in the power law relationship is not statistically significantly different from –2. *In other words, the hypothesis that the power law describing the frequency of extinctions of biological species is the same as that describing the frequency of extinctions of large firms is not rejected by the data*.

A detailed discussion of the data, analysis and results is given in Appendix 1. Appendix 2 provides the evidence from a quite different data set. This considers the statistical distribution which appears to characterise the demises of US firms, across the entire universe of such firms using a publicly available data base. The evidence is consistent with the hypothesis that the data follow a power law distribution.

[1] using a non-linear least squares algorithm 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

The data we examine comes from the American Office of Advocacy. Recorded in the data are the frequency of firm 'deaths' on an annual basis from 1989 through 1997 split into nine different industrial sectors for each of the 51 states[2]. This gives rise to 459 series of 9 annual observations in 51 states, a total of 4131 observations in total. We describe each of these observations as a 'group'. In other words, each group specifies a particular industry in a particular year in a particular state. Individual experiences vary widely, but it is interesting to note that on average just over 10 per cent of all active US companies (i.e. companies which trade) become extinct in any given year. Further, the rate of extinction is not correlated with the state of the business cycle. The number of demises was virtually identical in the recession year 1991, when real GDP fell by 0.5 per cent, and in the boom year 1997 when it grew by 4.3 per cent.

The technical details of considering how well the data fits a power law are more intricate than those involved in the analysis of the very large firm data set described above, and they are set out in Appendix 2.

The evidence is consistent with the hypothesis that the data follow a power law distribution with exponent –2.1. Again, this is virtually identical to that in the fossil record.

The analysis of two completely different data sets gives the same result. Namely, that the hypothesis that the power law describing the frequency of extinctions of biological species is the same as that describing the frequency of extinctions of large firms is not rejected by the data.

This obviously raises the possibility of the existence of a more general law of the frequency of extinctions in systems which evolve.

[2] District of Columbia is included as a state in this data set

3. The theoretical models

3.1 Overview

An important distinction is made between biological models in which extinction arises because of environmental stress (e.g.: asteroid impact) and those in which it arises because of competitive interaction between species. An identical distinction is made in economic theory, although the particular jargon which describes the two situations is 'exogenous', or outside the system, and 'endogenous', or within the system. Most economic models of the business cycle, for example, assume that it occurs due to exogenous shocks such as the speed of technological innovation, rather than arising from the interactions of the agents in the economy.

We consider two models from biology, each of which is a good representative of one of the two alternative approaches. In terms of the extinction of firms, it seems plausible that a model which incorporates both exogenous and endogenous mechanisms of extinction will be required. Firms are obviously endogenously connected to each other in terms of their commercial dealings, and there clearly have been some important shocks leading to extinctions which have been completely exogenous, such as the Russian revolution in 1917 or the nationalisation of many mining companies in Europe in the late 1940s. What is not clear is the respective weight which needs to be given to the exogenous and endogenous mechanisms, which is why we select a model from each of these approaches.

In each of the models, at any point in time each agent is assigned a particular level of fitness. 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. The models evolve on a step-by-step basis and contain rules which specify:

- how the fitness levels of individual agents evolve
- the level of fitness at which agents are deemed to become extinct
- how agents which become extinct are replaced

3.2 The exogenous shock model

In this model, originally proposed by Newman [6], the agents are not connected in any way. Extinction arises solely because of changes in the external environment. This general principle of the model is certainly compatible with economic thinking, given that, as noted above, external shocks play an important role in conventional economic theory.

The detailed results of our analysis of this model are contained in Appendix 3. Appendix 4 reports the effects of extending the model to allow for two potential types of connection between the agents. First, the extinction of any given agent leads to those agents which are directly connected to it also becoming extinct with a given probability, which is a parameter of the model. Second, the extinction of any given agents leads to those agents which are directly connected to it each having their fitness levels re-drawn at random.

In the basic Newman model, each agent has a stress tolerance level, which evolves at random. There is a level of external stress which also evolves at random. When the stress tolerance of an agent falls below the level of external stress in any period, the agent becomes extinct. A rule specifies how extinct agents are replaced.

The random evolution of the stress tolerance levels of agents can be reconciled with economic theory in a natural way. In an economic context, this can be thought of as a firm updating its strategy by a process of trial and error. 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 [7]).

In a biological context, when the model is populated by a very large number of agents it is able to generate a pattern of extinctions in which the relationship between the frequency and size of extinctions follows a power law with exponent of around –2. This result appears to hold across a variety of adaptations of the basic model.

We consider populating the model in an economic context by small numbers of agents (less than 1,000), which we can think of as representing the largest firms in an economy. Very large firms are of great importance not just in terms of the domestic economy, but also in terms of the ability of an economy to compete successfully in international markets.

In these conditions, a power law of extinctions with an exponent of around –2 can be generated, but it is not a particularly good fit to the data. In other words, the model is not necessarily applicable in all economic contexts.

More importantly, there is a parameter in the model which determines the proportion of agents whose stress tolerance – whose strategies in other words – are updated in any given step of a solution of the model over time. In a biological context, a time-scale of agent evolution which is slow relative to the frequency with which changes to the external environment occur makes sense, and this assumption is used in the results reported by Newman and his colleagues.

However, this assumption is much less plausible in an economic context. Agents - firms - can react very quickly to external events. The impact of their reactions may be highly uncertain, and there may be unintended and unforeseen consequences of actions which leave an agent worse off than if it had not reacted. But the point here is that agents *are* able to react quickly. The reactions of firms to the attack on 11 September illustrates this point clearly. No-one knew with any degree of certainty what the economic consequences would be. The one thing which companies did know for certain is that the external environment had altered, and they revised their plans, taking what seemed to them the best actions in the circumstances.

Once agents are allowed to react quickly to changes in the external environment, the resulting relationship between the frequency and size of extinctions can no longer really be said to follow a power law.

This result continues to hold under the various extensions to the model which we report in Appendix 4.

In other words, the rules within the exogenous shock model of agent extinction can in general be given plausible economic interpretations. However, when the realistic assumption is made that agents can react quickly to external events, we conclude that the model of agent extinction in which extinction arises solely from external shocks does not provide a particularly good description of reality in an economic context.

3.3 The endogenous extinction model

This model is based upon the model built to account for biological extinctions by Sole and Manrubia [8]. This model is essentially based upon the interactions of agents. The model contains N agents, and a matrix of couplings J_{ij} that indicates how each agent affects every other agent. A negative J_{ij} indicates that agents i and j are in competition, or that agent j feeds on agent i. A positive J_{ij} shows either that agent i feeds on j, or that they are in symbiosis. If both J_{ij} and J_{ji} are positive, for example, the two agents are in symbiosis.

The interpretation of competition or symbiosis has an inherent plausibility in terms of companies. Certain companies do compete directly, and the growth of one is at the expense of the other. Ford and General Motors is an obvious example. However, in many cases the growth of one company enables others to prosper. If General Motors grows, there will be more opportunities for its suppliers, from raw materials to accountants and consultants. An analogous agent-based model for the overall growth of market economies is described in non-technical terms in Ormerod [9].

The model specifies rules for how the J matrix is updated, how species are deemed to become extinct, and how these latter are replaced.

Again, as with the Newman model, the updating of the fitness of agents – expressed in this model by the J_{ij} matrix, is random. And, again, this is perfectly compatible with the economic concept of agents maximising under uncertainty, as discussed in section 3.2 above.

Overall, it is possible to place a plausible economic interpretation in the rules of this model, though the rule in the original model which specifies how extinct agents are replaced needs to be modified in order to give it the required plausibility.

With this model, even when populated by a small number of agents, a power law with exponent of around –2 gives a good description of the results. This conclusion is robust with respect to a variety of reasonable assumptions regarding the rules of the model, and in particular the rule which updates the fitness of individual agents. The results are discussed in detail in Appendix 5.

Further, the same result also obtains when we subject the model to external, exogenous shocks, which can be either specific to individual agents or common across all agents. Of course, the stronger the sizes of the shocks are compared to the strength of the connections between agents, the more the model becomes like the purely external shock model described in section 3.2. But even for fairly large shocks, the extinction patterns in the model remain well approximated by a power law with exponent of –2. The detailed results on this are set out in Appendix 6.

A way of thinking about this key result is as follows. External shocks themselves can give rise to power law behaviour of a system compatible with the empirical evidence. But a necessary condition for this to hold is that the internal structure which connects agents, and across which shocks are propagated, itself generates appropriate power law behaviour.

The results therefore suggest that the extinction patterns of firms are intrinsic to the system, arising from the interconnections between them. Shocks to the external environment may well play a role, but they are not the primary cause.

4. Future work

4.1 Towards a deeper understanding of agent-based evolution and extinction: extensions of the present modelling approach

The conclusion of the above is that the extinction patterns amongst economic agents are intrinsic to the system, arising from the interconnections between them. The extension of our analysis will therefore be based upon the endogenous extinction model.

There are essentially two components to this work to extend the model:

- we propose to extend the model by making it even more realistic in an economic context, and to see if the desired properties still hold. A particular way of doing this will enable the model to be used to illuminate evolution and extinction more generally in a wide variety of social networks which connect individuals
- we propose to understand more about what might be termed the 'policy properties' of this kind of model

As noted above, the model has a plausible and realistic interpretation. All models are approximations to reality, and substantial simplifications and abstractions need to be made.

In discussions with economists and complex systems theorists, however, the single point which is made most frequently about the model is that in reality firms are not on a completely connected network. A degree of sparsity needs to be introduced to the connection matrix, to allow for the fact that the strategy of agent k does not necessarily have a direct impact on the fitness of agent j. Of course, an indirect connection may very well exist. The actions of agent k may impact directly agent i, which in turn affects agent j.

We therefore propose to introduce various degrees of sparsity into the connection matrix, and to investigate to what extent the results are robust to this.

In this context, extending the model by connecting agents by a scale-free connection matrix can be thought of as an important special case of the more general principle of sparsity in the connection matrix. The agents in a scale-free network are not completely connected to each other, but are connected in a particular way.

Given the increasing evidence that social networks between individuals of various kinds[3] can be approximated by a scale-free network, developing the model with a scale-free network of connections is important in its own right.

This may enable us to extend the application of the model more generally to illuminate evolution and extinction in a wide variety of networks between individuals.

There is a second extension to the model which both increases the realism of its assumptions in an economic context, and may also make it more generally applicable in the context of networks of connections between individuals. In the current version of the model, the number of connections for any given agent remains fixed throughout the lifetime of the agent. We will allow the number of connections to grow with the lifetime of the agent. The economic interpretation in this context is that the longer a firm survives, the greater the probability that its size increases, and therefore the more agents with which it becomes connected.

The final extension is to investigate the impact of intent on the part of agents. Agents may not, under uncertainty, necessarily understand the consequences of their actions. But they may be able to estimate which of the existing agents at any point in time have the highest levels of fitness. New entrants in particular may therefore model themselves on this particular set of agents, rather than on a surviving agent chosen at random. We will investigate the effect on the model of allowing agents a degree of intent.

Some properties of the model related to potential policy considerations which we propose to examine are:

[3] evidence is published regarding, for example, the Internet and connections between sexual partners

- the influence of diversity on the overall fitness of the system
- to investigate quite generally whether any characteristics of the system can be identified which signal, however imperfectly, that a large extinction is imminent

The degree of diversity of the system can be thought of as a measure of the extent to which agents differ from each other. Diversity in evolving systems may be an important contributory factor to the success or otherwise of the system as a whole.

Each agent in the system at any point in time is characterised by two vectors – one which describes the impact of the agent on all other agents, and one which describes the impact of all other agents on itself. Distance metrics for the system as a whole will be calculated, and compared to the fitness of the system.

We will also reduce the level of diversity by modifying update/entrant rules and observing the impact on the fitness of the system.

With the second area of investigation, we propose to investigate whether any characteristics of the system can be identified which signal, however imperfectly, that a large extinction is imminent. It is possible that this may not generate any positive results, but it is important to examine this issue.

4.2 A particular extension of the model to the evolution and extinction of individual agents connected on a scale-free network

We propose to develop, building on the model extensions described above, a model of the evolution over time of the overall level of activity carried out by a network of individuals rather than firms. We have in mind networks in which the influence of an individual agent on another is often a key influence on the behaviour of that individual. A criminal gang is an obvious example, and we illustrate the proposed model with this in mind.

The main focus will be to get a better understanding of the impact of intervention in the system by the authorities. With a criminal gang, obvious examples are both the efficiency of and the sentencing policy of the criminal justice system, and the impact of prison on agents.

The Volterra team has experience of innovative modelling of crime. For example, in 2001 we carried out a substantial study for the UK government on 'Non-linear modelling of burglary and violent crime'. We set up systems of non-linear differential equations to describe the evolution and spread of crime[4]. Both non-linear differential equation systems and graph theoretic techniques (networks) are used more generally in understanding the spread of epidemics, in how ideas (such as cultural norms) or fashions diffuse through a system of connected agents.

Recently, a number of papers have examined the diffusion and control of viruses on scale-free networks. (e.g. Dezso and Barabasi [10]). However, the approach has not been applied to crime. Further, a number of important modifications are required in order to capture realistic features of such networks.

The approach may require modification as it is implemented. For example, analysis of scale-free networks described in section 4.1 above may inspire improvements. But the outline of our thinking on the model at present is as follows.

An important factor in the decision of an individual agent to become involved in crime is peer influence and pressure. A small number of individuals exercise influence over a large number of others, but most influence only a few. We can therefore think of a criminal network, or gang, as being approximated by a scale-free network. The network is a directed one. In other words, the fact that agent k influences the behaviour of agent j does not necessarily imply that j influences k.

We postulate that the number of crimes committed by the kth agent varies positively in proportion to the *total* number of connections of the agents which influence k. As a

[4] The study is being published by the British Home Office in the autumn of 2002.

practical illustration: a small group of youths who hang around with each other but who have no real connections to the major criminals of their area will commit relatively few crimes. But once such a connection is established, the probability that they will commit considerably more rises sharply. A big-time criminal exercises more influence over a given agent than a small-time one.

In fact, it is more useful to think not of the absolute number of crimes, but of an index of severity. A gang leader may influence a relatively new recruit to commit murder, for example. This would constitute just one crime, but its severity is obviously high. It is not necessary in the model to specify the mapping between the type of crime and the degree of severity – in any event, there is an inevitable degree of personal judgement involved in any such mapping.

But the key point remains: the position of an individual on the scale of severity of crimes he commits varies positively in proportion to the *total* number of connections of the agents which influence him. Any particular position on the scale could correspond to a small number of serious crimes, or a much larger number of less serious ones.

Agents in this network evolve in two ways:

- agents become extinct i.e. they give up crime. Empirically, the probability of this rises with age
- the longer they have been part of the criminal network, the more influence agents in general have over other agents

When an agent becomes extinct, it is replaced by one which has, initially, a low number of connections.

The exogenous policy variable in the model is prison, which removes the agent from the network. The impact can be introduced in several ways:

- an agent is chosen at random to be sent to prison

- the probability of being chosen varies positively with the number of crimes committed
- the probability varies inversely to the number of crimes committed (criminals become more skilled at covering themselves, can exercise much more intimidation etc.)

An agent can be removed either permanently (e.g. death penalty or very long sentence) or temporarily. If the removal is temporary, after a fixed period the agent can either decide to become extinct (give up crime) or return to the system.

An important feature of the model is that when an agent returns, its influence in the network is increased. This is a well known effect of prison: some are deterred by the experience, but others simply gain more contacts and knowledge. An important part of the exercise is to examine the sensitivity of the overall crime rate to assumptions on how much a re-entering agent increases its influence by the experience of prison.

The total amount of crime (measured on the scale of severity) which emerges from the system will be monitored. The impact of different types of intervention by the authorities, under a range of assumptions, will be examined. For example, it may be very difficult to remove agents who influence large numbers of agents, for they themselves may commit few crimes. It will usually be much easier to remove less influential agents. But if a proportion of these return to the system with a higher level of influence than before, the overall severity of crime committed may rise. It is such issues which will be examined.

4.3 Deliverables

In the existing project, we have delivered:

- a series of working papers describing aspects of the project in detail
- interim reports on progress
- a final report summarising the key findings, and containing detailed appendices on the work carried out

We propose to adopt the same approach in the extensions to the project.

References

1. W.E.G.Salter (1966), *Productivity and Technical Change*, Cambridge University Press

2. R.J.Town (1992), 'Merger waves and the structure of merger and acquisition time-series', *Journal of Applied Econometrics*, **7**, S83-100

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

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

5. 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

6. M.E.J.Newman (1997), 'A model of mass extinction', *J. Theor. Biol.*, **189**, 235-252

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

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

9. P.Ormerod (1998), *Butterfly Economics*, chapter 12, Faber and Faber

10. Z.Dezso and A-L Barabasi (2002), ‘Halting viruses in scale-free networks’, arXiv:cond-mat/0107420

Appendix 1

Marshall's 'Trees' and the 'Global Forest': the Extinction Patterns of Capitalism's Largest Firms

Paul Ormerod (mailto: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

***** 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

We analyse in this paper information on the evolution of the largest firms which existed in the opening decades of the 20th century.

We show that the relationship between the frequency and size of the extinctions is approximated well by a power law relationship. Further, the intervals of time between extinction events can also be described by such a relationship.

Analysis of the most informative data set available suggests that 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.

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) notes how this altered the judgement of Marshall, the leading economist of his day, on the life spans of corporations. In the first edition of his *Principles of Economics*, published in 1890, Marshall suggested that, like trees in the forest, there would be large and small firms, but 'sooner or later age tells on them all'.

As Hannah notes, Marshall was an acute observer of the contemporary economy in the UK, Germany and the US, and by the sixth edition of his book published in 1910, his views had changed. Marshall then held the opinion that 'vast joint stock companies ... often stagnate but do not readily die'. Much more recent work by business historians (e.g. Chandler (1990)) has generally been supportive of Marshall's latter view. By the start of the twentieth century, an entirely new phenomenon of the firm with global reach had been created, and these firms in general still dominate markets.

Hannah constructed 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. Hannah argues that this evidence suggests that, on balance, Marshall's earlier view is more consistent with the evidence than his later one. Very large firms do in fact die. Further, only a distinct minority retained their position in the top 100 companies.

The purpose of this paper is to examine the evidence on firm extinction for the existence of any particular patterns. There is a growing literature on extinction patterns in the fossil record of biological species (see Drossel (2001) for a detailed survey). An important aspect of this is the relationship between the number of species which become extinct in an given year, and the frequency with which the different numbers are observed. Section 2 carries out this analysis.

A further point of interest is the distribution of the timing of extinction events. In particular, we examine the distribution of the gaps between extinction events in the firm data set. In other words, the distribution of the successive number of years in which no extinction is observed. This is examined in section 3. Section 4 provides a brief conclusion.

2. Frequency and size of firm extinctions

The standard approach in the analysis of the extinction patterns of biological species (Drossel op.cit.) is to fit a relationship between the number of agents 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, which varies substantially. 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.

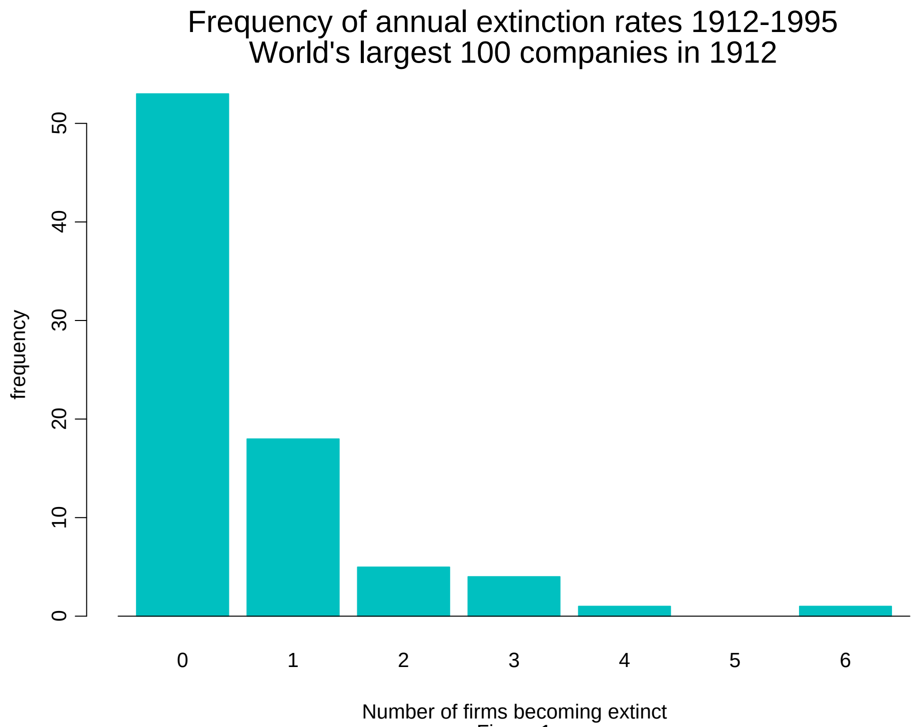

Figure 1

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

[5] 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

The actual values and those fitted by (1) are set out in Table 1. For comparison, the fitted values of an exponential relationship of the form F = α exp(β(N-1)) are also shown. This alternative form is frequently used for comparison with a power law in the biological literature.

Table 1. Frequency of annual extinction rates: Actual and fitted[6]

	Number of extinctions					
	1	2	3	4	5	6
Frequency						
Actual	18	5	4	1	0	1
Power law fit	18.0	5.3	2.6	1.6	1.1	0.8
Exponential fit	17.8	6.3	2.3	0.8	0.3	0.1

Just as in the description of biological extinctions (Drossel, op.cit.), the shortage of data points means that we cannot be too dogmatic about the precise nature of the functional form. However, a power law relationship in this case provides a better description of the data than does an exponential distribution.

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 large firm data set.

An important implication of this kind of relationship between frequency and size of extinctions is that extinctions on any scale can happen at any time. The probability of a very large extinction is very small, but it is greater than zero.

[6] the fitted values are rounded to 1 decimal place

A further data set is given by Fligstein (1990). This reports the membership of the 100 largest firms in the US, selected on the basis of asset size, at the end of each decade from 1919 through 1979. Altogether, 216 firms appear in the data set. Fligstein reports whether a firm was in the top 100, and does not give information on whether a firm ceased to exist as an independent entity. Exact years of entry and exit from the top 100 are not reported, but the status of each firm at the end of each decade.

There is one direct comparison which can, however, be made between the Fligstein and Hannah data sets. In the Fligstein US data, only 34 out of the top 100 in 1919 survived in the top 100 in 1979. Hannah's data set is the world top 100 in 1912, and by 1995 19 were still in the world's top 100 industrial companies. So the 'extinction' rate defined as exit from the top 100 is very similar in the two data sets. In Fligstein, 66 of the original top 100 in 1919 dropped out by 1979, and 81 of Hannah's 1912 firms dropped out by 1995. The annual average 'loss rate' from the top 100 is therefore very similar: 1.10 in Fligstein and 0.98 in Hannah. Of course, there is overlap between the data sets, with 42 of Hannah's 1912 companies being in the 1919 Fligstein data, but it is far from being complete.

Given that 'extinction' in the Fligstein data is defined as being exit from the top 100 and that this data is aggregated over time into periods of a decade, a direct comparison with the results from Hannah is not possible. However, a power law between frequency and size of 'extinction events' also fits the Fligstein data extremely well. Comparing 1919 and 1929, 31 firms exited from the US top 100, the largest number in any decade 1919-79. The lowest observed number of exits was 14 in the 1939-49 period. Fitting a least squares regression between the number of firms exiting in any given decade, and the overall rank of that number[7], the estimated value of the exponent of the power law is

-0.397 with a standard error of 0.054. The fitted values, rounded to the nearest whole number, are very similar to those of the actual data.

[7] In other words, the highest number of exits (31 during 1919-29) is given rank of 1, through to the lowest (1939-49) which is ranked 6.

To emphasise again, the Fligstein data set cannot be compared directly with the Hannah one for a number of reasons, but power laws give good descriptions of both data sets.

3. The distribution of the periods of years between extinction events

Figure 2 plots the information in the Hannah data set relating to 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, as for example in 1931 and 1932 and 1968 and 1969.

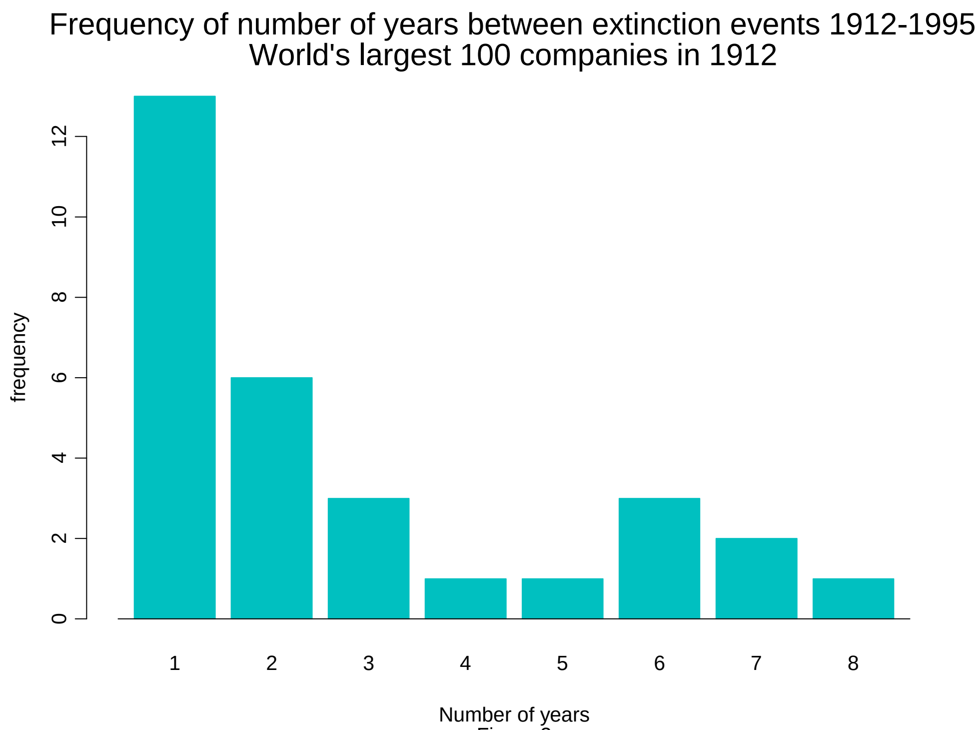

Figure 2

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 reported in section 2. Table 2 shows the actual and fitted values.

Table 2. Number of years between extinction events: Actual and fitted

	Number of years							
	1	2	3	4	5	6	7	8
Frequency								
Actual	13	6	3	1	1	3	2	1
Power law fit	12.0	5.3	3.3	2.3	1.8	1.5	1.2	1.0
Exponential fit	12.7	6.5	3.3	1.7	0.9	0.5	0.2	0.1

4. Conclusion

In this paper, we have examined patterns of extinction amongst capitalism's largest companies. The main focus of the analysis is a data set constructed by Hannah (1999) of the top 100 industrial companies in the world in 1912, which provides information on the year in which individual members of this group ceased to exist as independent entities over the 1912-95 period. A related data set is provided by Fligstein (1990) of the top 100 US companies 1919-79. For a variety of reasons the analyses of the two data sets are not in general comparable in a quantitative sense, but qualitatively the results using the two sets are very similar.

We find that power law relationships offer good descriptions of the extinction patterns of very large companies. In particular, the frequency with which a given number of extinctions in any time period is observed depends upon the inverse of the number raised to a power of itself. This is identical to the empirical relationships which have been used to describe the relationship between the frequency and size of extinctions of biological species. In addition, the distribution of the number of years between years in which extinctions occur also follows a power law.

Power law relationships imply that events on any scale can occur at any time. The probability of a large extinction happening in any particular period is much less than that of a small one, but such events are an intrinsic part of the survival/extinction patterns of very large firms.

References

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

N.Fligstein (1990), *The Transformation of Corporate Control*, Harvard University Press

Appendix 2

Power Law Distribution of the Frequency of Demises of U.S Firms[8]

William Cook and Paul Ormerod

Volterra Consulting, 121 Mortlake High Street, London SW14 8SN, UK

e-mail: wcook@volterra.co.uk, pormerod@volterra.co.uk

Abstract

Both theoretical and applied economics have a great deal to say about many aspects of the firm, but the literature on the extinctions, or demises, of firms is very sparse. We use a publicly available data base covering some 6 million firms in the US and show that the underlying statistical distribution which characterises the frequency of firm demises - the disappearances of firms as autonomous entities - is closely approximated by a power law. The exponent of the power law is, intriguingly, close to that reported in the literature on the extinction of biological species.

The purpose of this paper is to provide empirical evidence on the statistical distribution which appears to characterise the demises of US firms, across the entire universe of such firms using a publicly available data base. The evidence is consistent with the hypothesis that the data follow a power law distribution. The observation of power-law distributions (fractal behaviour) in a system's macroscopically observable quantities is a characteristic property of many-body systems representing the effects of complex interactions amongst the constituents of the system. Recent evidence that related aspects of economic activity are consistent with power distributions at the aggregate level is given, for example, by [1] on the growth of firms, [2] on the sizes of firms and [3] on the duration of recessions in the Western market economies.

[8] 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

Economic theory has a great deal to say about many aspects of the firm. Empirical studies of firm growth and size have emphasised the importance of stochastic influences, from the initial work of Gibrat [4] through classic papers in the 1950s and 1960s [for example 5,6], to more recent contributions such as [1, 2]. However, the literature on demises of firms is surprisingly sparse.

The disappearance of a firm as an autonomous entity - its demise - can take place for a variety of reasons, such as merger, take over and bankruptcy. The proximate reasons for demises over the 1912-95 period amongst the 100 largest industrial companies in the world in 1912 are given in [7], and similar evidence over the 1919-79 period of the 100 largest companies in the US on a decade by decade basis is provided by [8]. Very small firms constitute the overwhelming majority of incorporated businesses in terms of numbers, many of which are controlled by a single shareholder. Here, the reasons for a firm ceasing to exist can be even more diverse. In addition to bankruptcy, the owner may, for example, close a firm down simply because he or she has decided to focus attention on a different area of business. A small number of empirical studies which relate firm demises within individual industries to factors such as the number of firms in an industry at the time the firm enters are cited in [9], but no general analysis of demises appears to be available.

The data we examine comes from the American Office of Advocacy. Recorded in the data are the frequency of firm 'deaths' on an annual basis from 1989 through 1997 split into nine different industrial sectors for each of the 51 states[9]. This gives rise to 459 series of 9 annual observations in 51 states, a total of 4131 observations in total. We describe each of these observations as a 'group'. In other words, each group specifies a particular industry in a particular year in a particular state.

Data is also provided on the total number of firms for each of the 4131 groups. The number varies enormously, with the smallest being mining in Hawaii in 1997, with just six firms. The largest is service sector firms in California, with 265,710, also in 1997. To avoid any potential problems which might arise from the very small size of

[9] District of Columbia is included as a state in this data set

some groups, we exclude from the analysis any group which has in total less than 200 firms. This leaves out 221 observations, just over 5 per cent of the total, leaving a sample of 3910 observations to be analysed.

The average total number of firms across the United States over this period was 5.73 million, of which on average 611,000 died each year. The total number of firm demises varied very little from year to year, the maximum being 664,000 in 1996 and the minimum 577,000 in 1994. The number of demises bears little connection with the overall state of the economy. In the recession year 1991, for example, real GDP fell by 0.5 per cent and in the boom year 1997 it grew by 4.3 per cent. Yet the number of demises was very similar in both years, being 630,000 and 648,000 respectively.

The natural focus of analysis across the different groups is not, however, the total number of demises in each group but demises as a proportion of the total number of firms in each group. This eliminates the expected linearity between number of demises and group size, the correlation between the proportion of firms disappearing and the size of the group being only -0.02. Figure 1 plots the histogram of 'death' proportions.

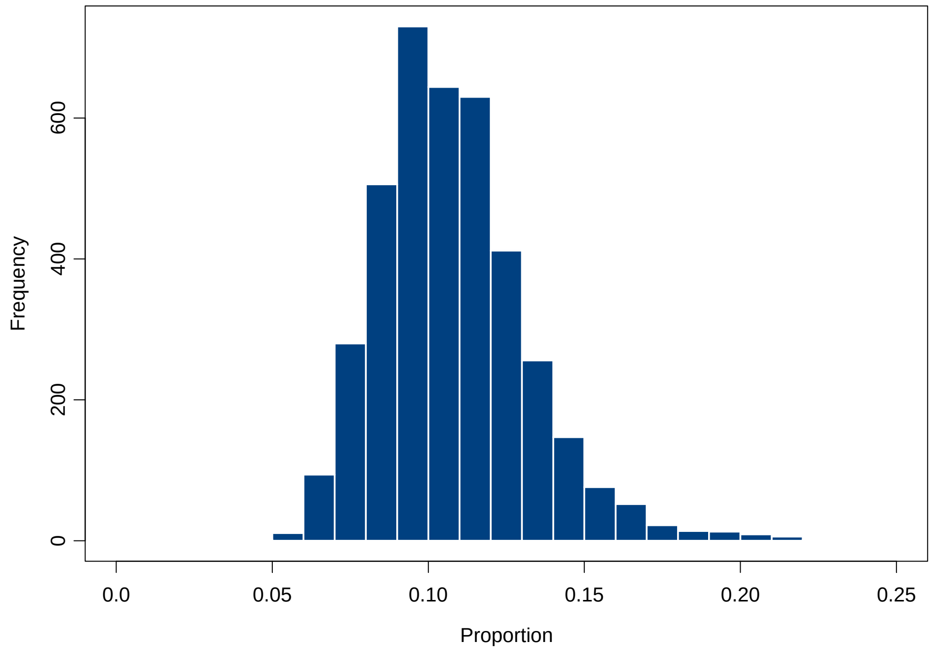

Figure 1: *Histogram of firm demises as a proportion of the total number of firms in the group. The group is defined by year, industry and state. Groups with less than 200 firms in total are omitted.*

At first sight, the data appears to be approximated closely by the lognormal distribution. The null hypothesis of lognormality is only rejected at $p = 0.0009$ using the Kolmogorov-Smirnov test.

However, this could be distinctly misleading. Each observation on the number of demises relates to the total number over the course of a year in any particular group. In other words, it is a *temporal* aggregation across a year of the demises which take place in any given category on each working day, approximately 250 per year.

The distribution we observe from the annual data is not necessarily that which gives rise to the data as it is actually generated on a daily basis. If day-to-day disappearances of firms are independent and identically distributed, we can treat each annual observation, A_i, as being the sum of N independent random variables. By the Central Limit Theorem, sums of iid variables of finite variance will converge towards the Gaussian distribution as N increases. But the actually observed annual data is very definitely non-Gaussian (the null hypothesis of normality is rejected on a Kolmogorov-Smirnov test at $p = 0.0000$).

Indirect evidence on the independence of the daily events whose sum makes up the annual observations is available from the time-dependency of the annual data. The correlation between the current year's proportion of demises and the previous years is 0.530 for the aggregate data - for the total number of demises divided by the total number of firms. However, this falls to an average of 0.265 across the 9 industries across time, and to an average of 0.161 across the 451 separate time-series for each industry in each state. In other words, the correlation over time between observations falls as the level of disaggregation of the data increases. This implies that the assumption that the daily observations - when the data is temporally disaggregated - of firm demises are independent is not an unreasonable one to make.

The assumption that the daily observations are identically distributed need not hold exactly for the Central Limit Theorem to be applicable. Rather, the variance of any separate distributions which exist must have variances which are not too dissimilar, so that no single variance dominates over all the others [10]. Given that the proportions of demises is in [0,1], again this does not seem an unreasonable assumption.

Only when the data is restricted to a sub-set of only 221 observations immediately around the mean, just over 5 per cent of the total, is normality not rejected at the conventional level of $p = 0.05$. In other words, the actual distribution of firm 'deaths' which we observe with data aggregated over a year converges only very slowly to a Gaussian. Each observation consists of the sum of some 250 daily events, yet despite this only a small fraction of the data is well described by the Gaussian distribution. The sum of independently, identically distributed random variables is known to converge only very slowly to a Gaussian when the underlying distribution follows a power law[10] [10]. We therefore postulate that the underlying distribution follows a truncated power law of the form:

[10] when the variance is finite. With a power law with infinite variance, the convergence is to the more general class of Levy distributions, of which the Gaussian is a special case

$$P(X = i) = \frac{(i+1)^{-r}}{\sum_{j=0}^{m} (j+1)^{-r}} \qquad i = 0, \ldots, m \qquad (1)$$

In this form, the distribution is dependent on two parameters, r and m. The rate at which the distribution falls is described by the parameter r, tailing off faster as r increases. The maximum value the distribution can take is given by m. We introduce m because allowing limitless demises per time period does not model a possible real life scenario. The generated data was scaled to the range of proportions observed.

We experimented with a range of values for m and r (for $r > 2$). The maximum annual number of demises observed in any group in the data was 30,366, a daily average rate of some 121, though for 99 per cent of the groups the implied daily average is less than 50. We examined the outcome with m chosen over the range 10 to 500. In each case, we drew 250 values, and summed them, and repeated the procedure 3,910 times, to obtain a series the same length as the number of observations in the actual data set examined. The best fit to the actual data was obtained with $m = 230$ and $r = 2.1$. The correlation between the (sorted) artificially-generated data set and the actual data is 0.995, and the null hypothesis that the two distributions are the same is only rejected on a Kolmogorov-Smirnov test at $p = 0.115$.

In terms of robustness of the results, the null hypothesis that the two distributions are the same is not rejected on a Kolmogorov-Smirnov test at the standard level of significance of $p = 0.05$ for values of m between 120 and 310. Interestingly, the literature on the extinction rates of biological species reports that the frequency/size relationship can also be approximated by a power law with an exponent close to 2 in absolute value [for example, 11].

The two sorted sets of data are plotted in Figure 2.

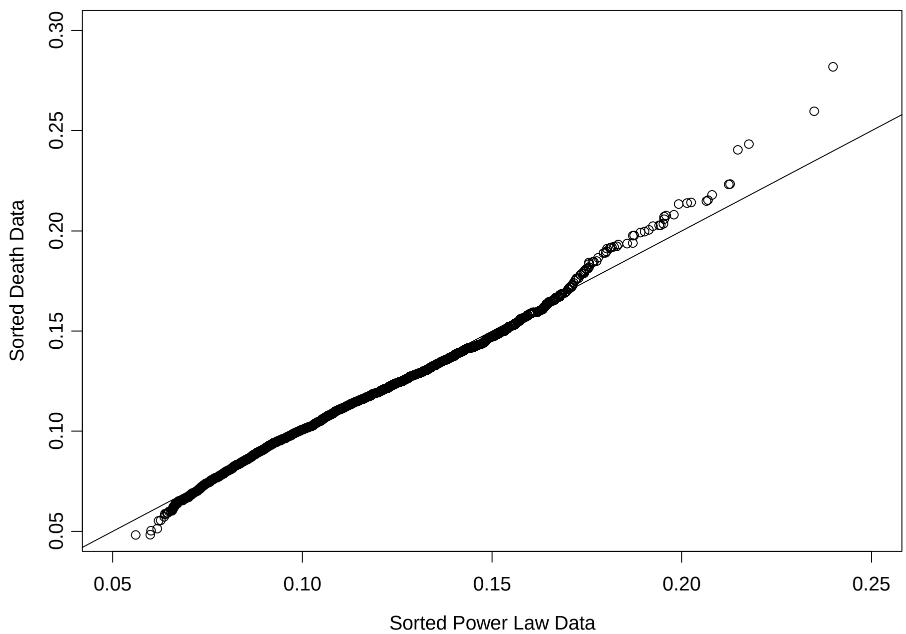

Figure 2 *Data generated from the truncated power law (1) with m = 230 and r = 2.1 plotted against the actual data, both data sets being sorted by size. Each observation of the generated data is the sum of 250 observations drawn from (1).*

We notice small deviations from the true distribution in the lower and upper tails. It is likely that the swing at the lower end is due to the small group size association still remaining.

In summary, the underlying statistical distribution which characterises the frequency of firm demises -the disappearance of a firm as an autonomous entity- in the United States is approximated well by a power law. The exponent of the power law is, intriguingly, close to that reported in the literature on the extinction of biological species

References

1. 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

2. R.L.Axtell (2001), 'Zipf Distribution of US Firm Sizes', *Science*, 293, 1818-1820

3. P.Ormerod and C.Mounfield (2001), 'Power Law Distribution of Duration and Magnitude of Recessions in Capitalist Economies: Breakdown of Scaling', *Physica* **A**, 293, 573-582

4. R.Gibrat (1931), 'Les Inégalitiés Economiques', Librarie du Receuil Sirey, Paris

5. I.G.Adelman (1958), 'A Stochastic Analysis of the Size Distribution of Firms', *American Economic Review*, 53, 893-904

6. P.E.Hart (1962), 'The Size and Growth of Firms', *Economica*, 24, 29-39

7. 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

8. N.Fligstein (1990), *The Transformation of Corporate Control*, Harvard University Press

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

10. J-P.Bouchard and M.Potters (2000), *Theory of Financial Risks: from Statistical Physics to Risk Management*, Cambridge University Press

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

Appendix 3

Patterns of Agent Extinction under External Shocks

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

***** 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

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.

We consider a model of the evolution and extinction of agents from the biological literature in which extinction arises solely because of changes in the external environment. The agents in the model are not connected in any way.

Each agent has a stress tolerance level, which evolves at random. There is a level of external stress which also evolves at random. When the stress tolerance of an agent falls below the level of external stress in any period, the agent becomes extinct. A rule specifies how extinct agents are replaced.

We examine the properties of the model when it is populated by small numbers of agents. The sensitivity of the properties are examined with respect to the rule by which the stress tolerance of agents is updated, and to the distribution and its parameters from which the external stress variable is drawn.

In general, the model does not give a particularly good description of the observed relationship between the frequency and size of extinction of firms. A good description can be obtained if we assume that each agent updates its stress tolerance level very infrequently, but this assumption is not plausible in an economic context. A very close approximation to a power law can also be obtained when agents update tolerance levels by small amounts each period and the level of external stress is very low, but the exponent on the power law differs from that which is observed.

We conclude that for small numbers of agents, a model in which extinction of autonomous, unconnected agents takes place solely through changes in the external environment does not offer a very satisfactory account of firm extinction.

1. Introduction

The observation of power-law distributions (fractal behaviour) in a system's macroscopically observable quantities is a characteristic property of many-body systems representing the effects of complex interactions amongst the constituents of the system. It is a property which emerges at the level of the system as a whole from the interactions of the individual agents which comprise the system. Power law distributions are both self-similar and scale free, demonstrating that events may occur on all length and time scales.

The economy is being analysed increasingly from the perspective of complex systems theory (for example, [1]). Evidence is growing for the existence of power-law distributions in aggregate economic behaviour (for example, [2-4]). An analysis of the extinction rates of the world's 100 largest industrial companies in 1912 over the period from 1912 to 1995 shows that a power law with an estimated exponent of close to -2 gives a good empirical description of the data [5].

The size distribution of extinction events amongst biological species can also be described by a power law with an exponent of around -2 [6-8]. A number of models based upon the principle of self-organised criticality of a system arising from interactions between agents have been developed to account for this empirical power law relationship (for example, [9-11], with a recent survey given by [12]).

In economic theory, an important distinction is made between models in which changes occur because of events which are external to the system - 'exogenous' in the jargon of economics - and those in which changes take place because of interactions within the system itself - 'endogenous'. An identical distinction is made between biological models of extinction, in which extinction arises because of environmental stress (eg: asteroid impact) and those in which it arises because of competitive interaction between species, with examples of these being given by [11] and [9] respectively.

In terms of the extinction of firms, it seems plausible that a model which incorporates both exogenous and endogenous mechanisms of extinction will be required. Firms are obviously endogenously connected to each other in terms of their commercial dealings, and there clearly have been some important shocks leading to extinctions which have been completely exogenous, such as the Russian revolution in 1917 or the nationalisation of many mining companies in Europe in the late 1940s. What is not clear is the respective weight which needs to be given to the exogenous and endogenous mechanisms.

In [13], we analyse and adapt the model set out in [9] to examine the extinction patterns of firms from the perspective of a model in which extinction occurs solely through endogenous interactions between the companies.

The purpose of the present paper is to analyse the properties of the Newman model [11], in which agent extinction takes place purely through exogenous shocks, again in the context of extinctions of firms rather than of biological species.

2. Description of the theoretical model

2.1 The basic model

The Newman model is populated by N agents, each of which is characterised by one number, x_i, which stands for its stress tolerance. Initially, these are chosen at random from a uniform distribution on [0, 1]. The second variable in the model is the level of external stress, η, which is chosen at each time step independently and at random from a distribution $p_{stress}(\ )$ which is either Gaussian or exponential.

The model evolves in a sequence of steps. In each step, a value for the stress η is chosen, and all agents for which $x_i < \eta$ become extinct. A step in which m agents become extinct is referred to as an extinction event of size m. Each extinct agent is replaced immediately with a new one with a value of x_i chosen at random from a uniform distribution on [0, 1][11].

In addition, in each step, a small fraction, f, of all agents obtain a new random value of x_i. In an economic context, this can be thought of as a firm updating its strategy by a process of trial and error. 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 [14]).

2.2 Extensions of the model

Using an algorithm which calculates the properties of the model with an infinite number of agents, the frequency of extinction sizes is shown to follow a power law with an exponent $\cong -2.0$ [9, 15]. This result appears to hold across a variety of adaptations of the basic model. A fuller description of these is given in [15], along with appropriate references, and a summary is as follows.

For example, the choice of other distributions such as the Poisson, power law or stretched exponential for $p_{stress}(\ )$ makes little difference to the results. Further, in the version of the model described above, there is no interaction between agents. If the agents are placed on a simple one-dimensional lattice, and the extinction of any given agent leads in the same step to the extinction of its immediate neighbours regardless of their levels of stress tolerance, x_i, the frequency of extinction events still follows a power law with exponent $\cong -2.0$. Similarly, the result holds if new agents

[11] It can be shown analytically [15] that the results are not sensitive to the choice of this particular distribution

are assumed to inherit their x_i from surviving species rather than having them allocated at random from [0, 1]. A more sophisticated version of the model introduces different types of external stress. Each agent has a stress tolerance level for each of the types of stress, and becomes extinct whenever one or more of the values of the external stress variables exceeds its relevant stress tolerance level. Again, the model reproduces the power law on the sizes of extinction events.

Finally, results for a version of a model in which extinct agents are not replaced immediately but gradually over a period of time are discussed in [15]. This appears to raise some difficulties for the properties of the model when they are compared to those of the fossil record on biological species extinctions. However, this is not really relevant in an economic context, where it is reasonable to assume that new firms immediately replace extinct ones. For example, if we consider extinction patterns amongst the very largest firms [5, 13], whenever one of the top, say, 100 firms becomes extinct, by definition it is replaced in this category by the next largest. More generally, the American Office of Advocacy provides evidence on firm births and deaths on an annual basis from 1989 to 1997, a period which spans both a recession and a strong economic recovery. The average total number of firms across the United States in this data base over this period was 5.73 million, of which on average 623,000 "died" during the course of a year, and 705,000 were "born".

2.3 The focus of the current paper

A substantial proportion of the output of the American economy, and indeed that of the other Western market economies, is accounted for by a small number of very large firms. The first major phase of mergers, acquisitions and expansions which led to this situation took place in the decades immediately around 1900. By the time of the First World War, the wave of corporate restructurings had been consolidated, and firms capable of operating on a global scale then dominated the US economy for the first time (for example, [16, 17]).

There is, of course, a very large number of small firms whose contribution to total output is substantial, and the average size of business organisations in the closing decades of the twentieth century fell quite markedly [18]. Nevertheless, a key feature of the US economy is a concentration of output amongst a small number of firms. By way of illustration, over 80 per cent of the total war production of America in the Second World War was carried out by just 100 firms [19].

Very large firms are of great importance in terms of the ability of an economy to compete successfully in international markets. The experience of large firms in the American, British and German economies in the twentieth century is given in [16]. And the phenomenal expansion of the Japanese economy in the second half of the twentieth century, when its per capita dollar income rose from around 25 per cent of that of the United States in the mid-1950s to 80 per cent by 2000, was based to a large extent on the external success of its giant manufacturing companies.

The focus of the present paper is to analyse the properties of the Newman model when it is populated by small numbers of agents, which we can think of as representing the largest firms in an economy. An understanding of their extinction patterns is of particular importance.

3. The results

3.1 Background

We consider the properties of the basic model discussed in section 2.1 above for N = 50, 100, 250 and 500. The model is solved in each case over 10,000 iterations[12] a total of 500 separate times ; the results reported are the averages across the 500 solutions.

The results for different values of N are qualitatively similar for any given set of parameters of the distribution from which the level of external stress is drawn. In

[12] This number was reached after experimentation with up to 100,000 iterations showed that the smaller number is sufficient to determine the properties of the model when using small numbers of agents.

general, the lower N is, the higher the absolute value of the exponent of the power law, but the estimated values are all in the neighbourhood of -2. For example, with the external stress variable, η, following an exponential distribution with mean of 0.05 and with f = 0.01, the exponent of the power law when N = 50, 100, 250 and 500 is, respectively, -2.30, -2.04, -1.74 and -1.55. With η drawn from a normal distribution with standard deviation of 0.05, the estimated values of the exponent are, respectively, -2.54 , -2.15, -1.74 and -1.49.

In addition to the different number of agents, we examine the sensitivity of the model with respect to:

- the distribution, $p_{stress}(\ )$, from which the level of external stress, η, is chosen. Specifically, we examine the exponential and the normal distributions
- different parameters of $p_{stress}(\ )$ for each of the two distributions
- the proportion of agents in each period, f, whose levels of stress tolerance - or strategy as we think of it in an economic context - are updated in each period.

The choice of parameters of the distribution of external stress, $p_{stress}(\ )$, needs to be considered in the context of the distribution of the stress tolerance of agents. The latter is drawn initially from a uniform distribution on [0, 1]. The mean tolerance of agents evolves in general from an initial value of 0.5 to some 0.7[13], because agents which survive have a higher than average tolerance to stress. If the parameters of the distribution of external stress, $p_{stress}(\ )$, are set too high, large extinctions will occur frequently simply because of the values of the external stress variable. For example, with an exponential distribution with a mean of 0.25, a value of greater than 0.7 will be drawn approximately once every 16 steps. Even more importantly, a value greater than 1 will be drawn just under 2 per cent of the time, leading to the complete extinction of all agents[14], whereas with a mean of 0.05, the probability of this occurring is effectively zero.

[13] obviously, this value will depend upon the parameters which are chosen, but this value is reasonably typical

[14] it is possible in the version of the model discussed in section 3.3 for a very small proportion of agents to obtain stress tolerance levels greater than 1

3.2 The basic model

The arguments above indicate that it is unlikely that a power law of agent extinctions will be observed if the parameters of the external stress distribution are set at too high a level: there will be frequent large extinctions. This is readily confirmed from simulations of the model. Figure 1 plots (on a log-log scale) the frequency of extinctions and the size of the extinction. These results were obtained with N = 100 and f = 0.01, but they are typical of the properties of the model at relatively high values of the external stress distribution parameters. Figure 1a shows the results with the exponential distribution with the mean set at 0.1, and Figure 1b shows the normal with the standard deviation set at 0.15.

In each case, even at these values of the stress distribution parameters, the results of the model are not described very well by a power law. There are relatively large numbers of very large extinctions including, in the case of the illustration of the exponential distribution, examples of the complete extinction of all 100 agents.

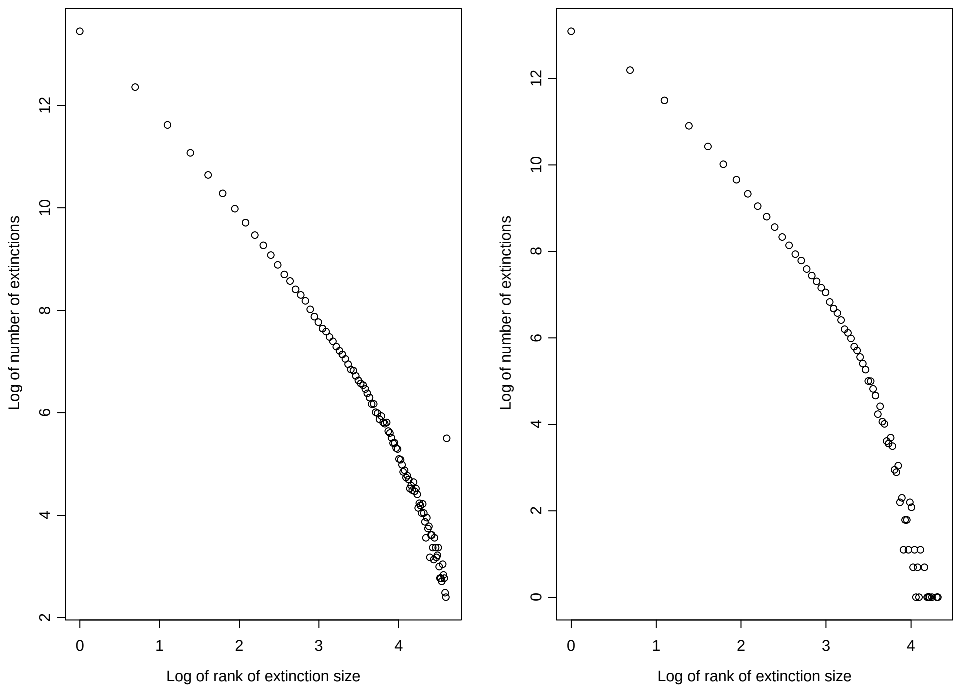

Figures 1 a and b. *Log-log plots of frequency of extinctions and extinction size. In each case, N = 100 and* f = 0.01. *In Figure 1a,* η *is drawn from an exponential distribution with mean = 0.1, and in Figure 1b from a normal distribution with standard deviation = 0.15. In both cases, a power law does not give a particularly good description of the data.*

The results set out in Figures 1a and b are perhaps obvious. Less obvious, but much more important in terms of the interpretation of the model, is the choice of f, the proportion of agents whose stress tolerance is updated in each period. The system receives an external shock - a change in the external stress value - every single period. The parameter f determines the average number of periods between changes to the stress tolerance of the agents of the system. So with f = 0.01, for example, 1 per cent of the agents update their stress tolerances each period, so on average any given agent will update its stress tolerance every 100 periods[15]. In other words, with f = 0.01,

[15] provided of course that it has not become extinct over the course of any relevant period

individual agents evolve only very slowly relative to the frequency with which changes in external stress take place.

In a biological context, a time-scale of agent evolution which is slow relative to the frequency with which changes to the external environment occur may well make sense. However, it is much less plausible in an economic context. Agents - firms - can react very quickly to external events. The impact of their reactions may be highly uncertain, and there may be unintended and unforeseen consequences of actions which leave an agent worse off than if it had not reacted. But the point here is that agents *are* able to react quickly. The reactions of firms to the attack on 11 September illustrates this point clearly. No-one knew with any degree of certainty what the economic consequences would be. The one thing which companies did know for certain is that the external environment had altered, and they revised their plans, taking what seemed to them the best actions in the circumstances.

In this version of the Newman model, however, power law behaviour can only be generated at very low values of f. This is illustrated in Figures 2a and b below, which show the log-log frequency/size extinction plot for the exponential and normal distributions of external shocks respectively, with $f = 0.2$. In other words, an agent reacts on average only every five periods to the external shocks which take place every single period.

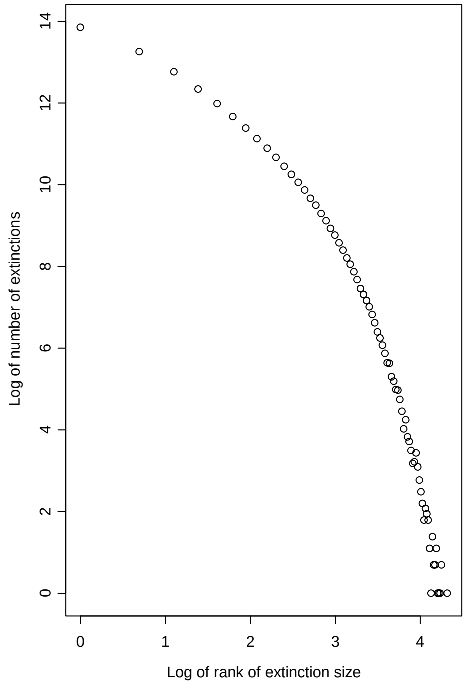

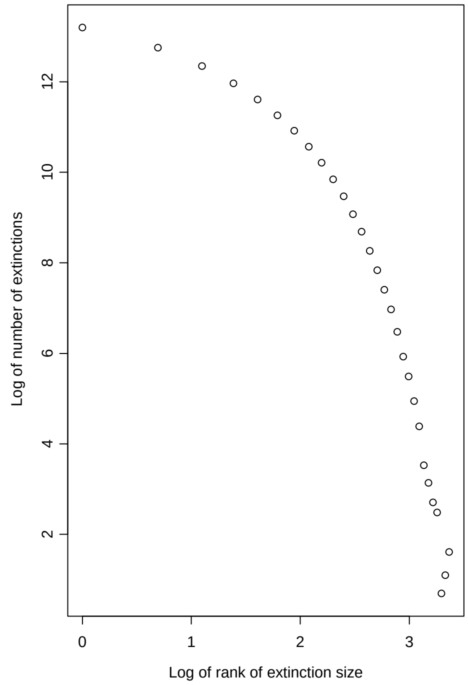

Figures 2 a and b. *Log-log plots of frequency of extinctions and extinction size. In each case, N = 100 and* $f = 0.2$. *In Figure 2a, the external stress variable* η *is drawn from an exponential distribution with mean = 0.05, and in Figure 2b from a normal distribution with standard deviation = 0.05. In both cases, a power law does not give a good description of the data.*

3.3 An alternative rule for the evolution of stress tolerance levels

The above version of the model is described in biological terms as one of 'punctuated equilibrium'. In other words, agents update their individual stress tolerance levels relatively infrequently. Further, when they do update, the change is potentially large, with the new value being drawn at random from the uniform distribution on [0, 1]. As an example, the potential size of the changes can be seen by examining the distribution of the differences between two variables, each of 1 million observations, drawn at random from the uniform distribution on [0, 1]. The mean of the differences is, of course, zero. But the inter-quartile range is between -0.29 and +0.29. In other words, 50 per cent of agents in this example experience a change in stress tolerance which is greater than 0.29, a fairly large jump. Of course, in the model the stress tolerance distribution of surviving agents is not uniform, but is within [0, 1].

We now consider a version in which each agent updates its stress tolerance by small amounts in each single period. In other words, agents react with the same frequency as changes in the external environment. The update rule in this version of the model is that the stress tolerance of agent i at time t, $x_{i,t}$, evolves as follows: $x_{i,t} = x_{i,t-1} + \varepsilon_t$, where ε_t is a normally distributed random variable with mean zero and variance σ^2, and σ^2 is small.

This version of the model is capable of somewhat better approximations than that of the basic model to a power law relationship between the frequency and size of extinctions with an exponent of around -2. Consider, for example, the model with N =100 and the external stress variable, η, drawn from an exponential distribution with mean of 0.05. Setting the standard deviation of ε, σ, at 0.01, 0.05 and 0.10, the estimated exponent on the power law relationship is -2.22, -2.84 and -2.94 respectively.

However, the power law least squares regressions tend to over-predict the number of medium and large extinctions generated by the model. A typical set of results is plotted in Figures 3a and b below.

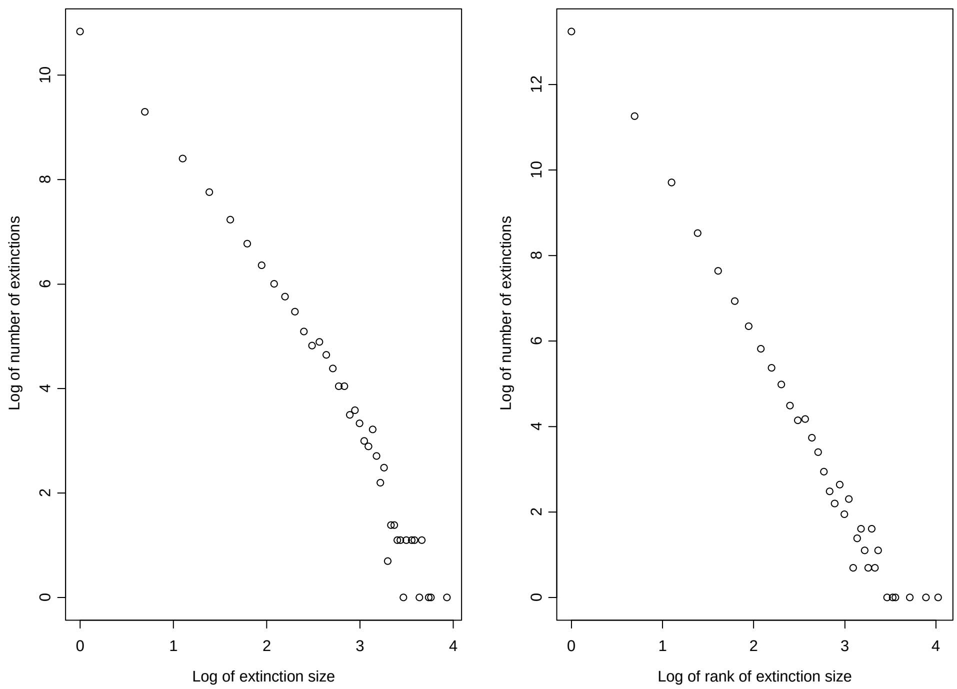

Figures 3a and b. *Log-log plots of frequency of extinctions and extinction size. In each case, N = 100 and f = 0.01 and η is drawn from an exponential distribution with mean = 0.05. The stress tolerance of all agents is updated in each period. The change in the stress tolerance in the previous period is drawn from a normal distribution with standard deviation of 0.01 in Figure 3a and 0.10 in Figure 3b.*

When both the external stress variable takes on only small values and the stress tolerance of agents changes by only small amounts from period to period, a power law describes the data more accurately. For example, with η drawn from an exponential distribution with mean = 0.01, this variable hardly ever takes a value of greater than 0.1. Indeed, over 99 per cent of all the values in such a distribution fall within the range [0, 0.05], so the level of external stress is small.

Importantly, however, the absolute value of the estimated exponent of the relationship in such cases tends to be considerable greater than 2. For example, Figures 4a and b plot the frequency/size of extinctions relationship for a exponential distribution of

$p_{stress}(\ )$ with mean 0.01 and 0.02 respectively, and for ε with a standard error of 0.01 and 0.02 respectively. The fit to a power law is good, but the estimated exponents are -4.36 and -3.63 respectively.

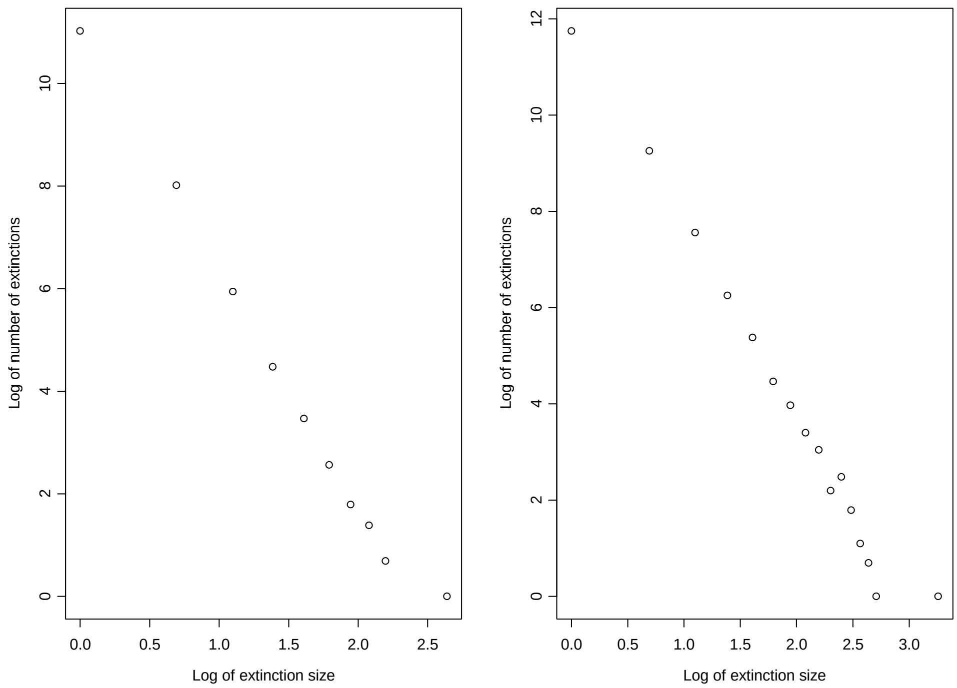

Figures 4a and b *Log-log plots of frequency of extinctions and extinction size. In each case, N = 100 , f = 0.01, the external stress variable, η, is drawn from an exponential distribution and the stress tolerance of all agents is updated in each period. In Figure 4a, the mean of η is 0.01 and the standard deviation of ε is also 0.01. In Figure 4b these are 0.02 and 0.02 respectively.*

4. Conclusion

We consider in this paper an agent-based model of evolution in which the agents act autonomously and are not connected in any way. Agent extinction can only arise through changes in the external environment.

Each agent has its own stress tolerance level, which evolves over time. The model contains a level of external stress, which is updated by a random process in each period. If the stress tolerance of an agent falls below that of the external stress level, the agent becomes extinct and is immediately replaced with a new agent.

A number of variants of the basic model have already been examined in the literature [15]. In this paper, we examine the properties of the model when it is populated by a small number of agents, which we can think of as representing very large firms in a capitalist economy.

We examine the relationship between the frequency and size of extinctions which is generated by the model. It is often possible to obtain a relationship between the two which can be characterised by a power law with an exponent $\cong -2.0$, consistent with actual empirical evidence. However, there are important qualifications to this statement.

The sensitivity of the model is considered with respect to the number of agents and both the distribution from which the external stress variable is drawn and its parameters. We investigate two sets of rule by which the stress tolerance of agents is updated. First, when only a fixed proportion of agents is updated each period. Second, when each agent is updated in each period.

In the former case, a power law gives a reasonable approximation to the relationship between frequency and size of extinctions generated by the model only when a very small proportion of agents are updated each period. Whilst this may well make sense in a biological context, it is not realistic in an economic one. Agents are able to react rapidly following changes in the external environment, as the events after September 11 illustrate.

The variant of the model in which the stress tolerance of all agents is updated by small amounts in each period gives results which in general approximate somewhat more closely to a power law relationship between the frequency and size of extinctions.

However, a power law fitted to the model data usually over-predicts the number of medium-sized and large extinctions. The exception is when both the external stress variable and the stress tolerances of agents change by only very small amounts in each period, but the exponent on the estimated power law is then substantially larger in absolute terms than that which is observed empirically.

We conclude that the Newman model of agent extinction in which extinction arises solely from external shocks does not provide a particularly good description of reality in an economic context.

References

1. P.Ormerod (1998), *Butterfly Economics*, Faber and Faber, London

2. L.A.N Amaral, S.V.Buldyrev, S.Havlin, M.A.Salinger, and H.E.Stanley (1998) 'Power law scaling for a system of interacting units with complex internal structure', *Phys. Rev. Lett.*, **80**, 1385-1388

3. P.Ormerod and C.Mounfield (2001), 'Power law distribution of duration and magnitude of recessions in capitalist economies: breakdown of scaling', *Physica* **A**, 293, 573-582

4. P.Ormerod (2002), 'The Economic Cycle as a Complex System Phenomenon: Power Law Distribution of the Duration of Recessions in European Economies', *Economics and Complexity*, forthcoming

5. 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

6. D.M. Raup (1986), Science, **231**, 1528

7. R.V. Solé and J. Bascompte (1996), Proc. R. Soc. London B, **263**, 161

8. M.E.J. Newman (1996) Proc. R. Soc. London B, **263**, 1605

9. 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

10. S.C.Manrubia and M.Paczuski (1998), 'A simple model of large scale organization in evolution, *Int. J Mod. Phys. C*, **9**, 1025-1032

11. M.E.J.Newman (1997), 'A model of mass extinction', *J. Theor. Biol.*, **189**, 235-252

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

13. P.Ormerod, H.Johns and L.Smith (2001), 'An Agent-Based Model of the Extinction Patterns of Capitalism's Largest Firms', www.volterra.co.uk

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

15. M.E.J.Newman and R.G.Palmer (1999), 'Models of Extinction: A Review', adap.org/9908002 submitted to Philosophical Transactions of the Royal Society.

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

17. 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

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

19. R.Overy (1995), *Why the Allies Won*, Jonathan Cape, London

Appendix 4

A Model of Agent Extinction: Combining Exogenous Shocks and Networks

Laurence Smith (lsmith@volterra.co.uk)*

Paul Ormerod (pormerod@volterra.co.uk)

Helen Johns (hjohns@volterra.co.uk)

Volterra Consulting Ltd

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

***** 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

We consider a model of the evolution and extinction of agents from the biological literature in which extinction arises because of changes in the external environment. In addition, the agents are placed on a network.

We examine two versions of the model. First, the extinction of any given agent leads to those agents which are directly connected to it also becoming extinct with a given probability, which is a parameter of the model. Second, the extinction of any given agents leads to those agents which are directly connected to it each having their fitness levels re-drawn at random.

We find that in general the model does not give rise to power law behaviour between the frequency and size of extinctions. For certain ranges of the parameters, an exponential distribution fits this data very well. Finally, the properties of the model are not particularly sensitive to the choice of topology which connects agents.

1. Introduction

This paper builds on recent work described in [1] to determine whether the Newman Model of mass extinction [2] could be translated into a useful model for considering the extinction rates of firms. For values of the parameters which make economic sense, and with a small number of agents, it was not really possible to generate a power law which matches the empirical evidence [3].

This paper considers whether these results still hold under extensions to the Newman Model where networks are introduced which connect firms. In the basic model of the evolution and extinction of agents, extinction arises solely because of changes in the external environment. The agents in the model are not connected in any way.

In this paper we consider versions of the model in which agents are connected by a variety of geometries. We consider two separate ways in which the connections operate. First, whenever an agent becomes extinct, all agents to which it is directly connected also become extinct with a given probability chosen from [0, 1]. Second, these neighbours simply have each of their fitnesses re-drawn at random.

2. Descriptions of the theoretical models

2.1 The basic Newman Model

The Newman model is populated by N agents, each of which is characterised by one number, x_i, which stands for its stress tolerance. Initially, these are chosen at random from a uniform distribution on [0, 1]. The second variable in the model is the level of external stress, η, which is chosen at each time step independently and at random from a distribution $p_{stress}(\ )$.

The model evolves in a sequence of steps. In each step, a value for the stress η is chosen, and all agents for which $x_i < \eta$ become extinct. A step in which m agents become extinct is referred to as an extinction event of size m. Each extinct agent is replaced immediately with a new one with a value of x_i chosen at random from a uniform distribution on [0, 1][16].

In addition, in each step, a small fraction, f, of all agents obtain a new value of x_i drawn from a uniform distribution on [0,1]. In an economic context, this can be thought of as a firm updating its strategy by a process of trial and error.

2.2 Newman Model with Small Continuous Updates (SCU)

Another version of the model (also detailed in [2]) updates the agents differently. Instead of drawing a random value for x_i for a fraction, f, of agents, all agents have their stress tolerance changed in every period by the addition of a small amount of random noise. I.e. if the stress tolerance of agent i at time t is $x_{i,t}$, at time t + 1 it will be $x_{i,t+1} = x_{i,t} + \varepsilon_t$, where ε_t is the random noise. In this paper we have kept ε_t as a normally distributed random variable with mean zero and variance σ^2, where σ^2 is small. We have labelled this version of the Newman Model as being the Newman Model with Small Continuous Updates (SCU).

2.3 Newman Model on a network

The basic Newman model is a model of exogenous shocks. There is no interaction between the agents in the model. This is obviously not the case for firms in an economy where firms may supply goods or services to each other or be competitors and thus events which affect one firm are likely to also affect others.

We connect firms on different types of network (see 2.4) and apply two different rules for how the extinction of one of a firms neighbours affects it.

[16] It can be shown analytically [4] that the results are not sensitive to the choice of this particular distribution

2.3.1 Grouped dependents

In this version, firms are treated as if they are organised so that a firms neighbours are highly dependent on it, for example they could be the major suppliers or customers.

When a firm becomes extinct, each of its neighbours become extinct with probability ζ_1.

2.3.2 Fight for survival

This version is not so harsh on firms' neighbours. When a firm becomes extinct, its neighbours draw a new random value of x_i . We can think of this as altering their own strategies following the extinction of an agent which is a close competitor or collaborator.

2.4 Types of network

There are many different types of network on which firms can be placed. These networks have a number of different properties. The most important property of networks is the density or average number of connections. In this example, this is the number of neighbours of each firm. Other properties include the extent to which connections from firms who are neighbours overlap.

Anticipating the results below we will not spend much time describing different types of network. Descriptions can be obtained from [5] and references therein.

The main network we will be using is the small world network as described in Watts and Strogatz [6]. To create a small world network we start with a ring of firms where each firm is connected to its k nearest neighbours (k being a multiple of two). Then with probability p, each connection is rewired to a random firm.

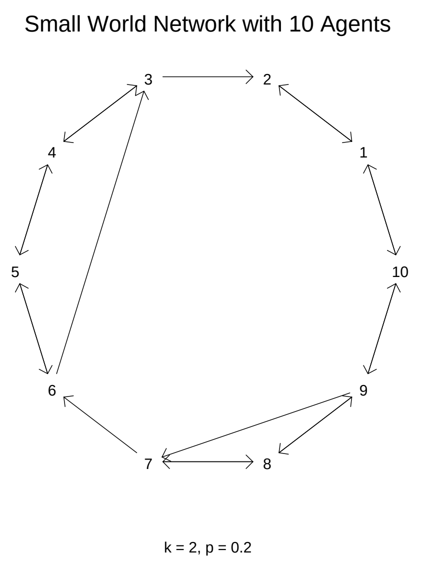

Figure 1: Example small world network with 10 agents, k = 2, p = 0.2

In the example above we can see that each firm (numbered 1 to 10) started off connected to the firms numbered one higher and one lower. In the rewiring process the connection from 2 to 3 was rewired to go from 6 to 3, and the connection from 6 to 7 was rewired to go from 9 to 7[17].

Watz and Strogatz give examples of real life networks which have similar properties to small world networks, the most famous being the collaboration graph of actors in feature films (the so called Kevin Bacon Graph). Since we do not have any evidence for the actual graph of connections between firms, we calculate most of the results below using a small world network. The results were tested for their sensitivity to different types of network.

3 Results

Due to the large number of parameters for the Newman Model, possible distributions for levels of stress and stress tolerances and possible networks there are a huge number of possible results and so it is not possible to report all of them here. In these

[17] The rewiring process is random, so there could be a situation where the connection from 4 to 5, say, was chosen to be rewired then, with probability of one over the number of firms, the new connection could also be form 4 to 5.

results we aim to limit the number of parameters by only using those which make economic sense.

We show a subset of the total results obtained, which are representative of the results in general and highlight the effects or non-effects of changing the inputs.

All the results presented are concerned with the frequency of extinction events of various sizes. We do not include results for periods when 0 extinctions took place. Each reported set of results is obtained by taking the mean frequencies of extinction events from 100 separate runs of the model (each run having an individually generated network). Each run of the model lasted for 100,000 periods. Experimentation with larger numbers of runs, or runs lasting for a greater number of periods show no significant change to the results.

The number of firms in each run was fixed at 100, this being the number of firms studied in [3] which is the real life evidence we have for rates of firm extinctions.

For the versions of the model where a proportion f of agents have their stress tolerances redrawn from the underlying uniform distribution each period, f was fixed at 0.2. This means that firms change their strategy on average only every 5 periods compared to the external stresses which occur every period. In terms of actual firms this seems to be the minimum acceptable level for f. This compares with typical values for f in [7] of 10^{-6}. See the discussion for more on this choice.

3.1 A typical result

First we show a typical result for the frequency of mean extinction events for each of the four models described above.

We use a small world network with k = 4 and p = 0.05. The external stress distribution is exponential with mean = 0.05.

The first result we show (figure 2) is for the Newman Model on a network with grouped dependents (as described in 2.3.1).

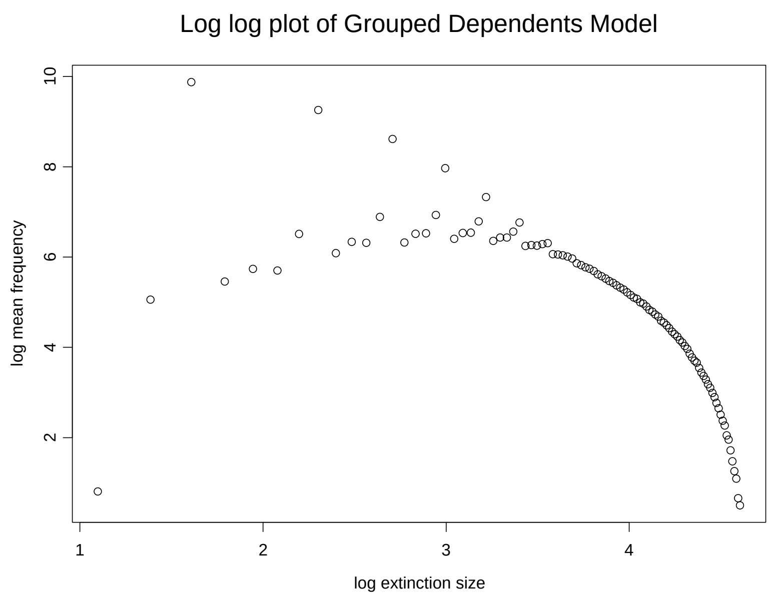

Figure 2: Mean results over 100 simulations of the Newman Model with Grouped Dependents. Performed on a small world network with N = 100, k = 4, p = 0.05. The stress distribution was exponential with mean 0.05.

From the Left hand side of the graph, at five point intervals there is a jump in the number of extinction events of that size. This is due to the fact that when an agent becomes extinct, all its four neighbours also become extinct and thus five agents become extinct at once. Not all extinction events are multiples of five however since some firms that become extinct might have neighbours which overlap and also in the rewiring process some connections might overlay one another.

However, it is clear from the graph that the results do not follow a power law.

We next look at the variations in results over the hundred runs of the model, each run being performed on a different draw of a small world network with the same parameters.

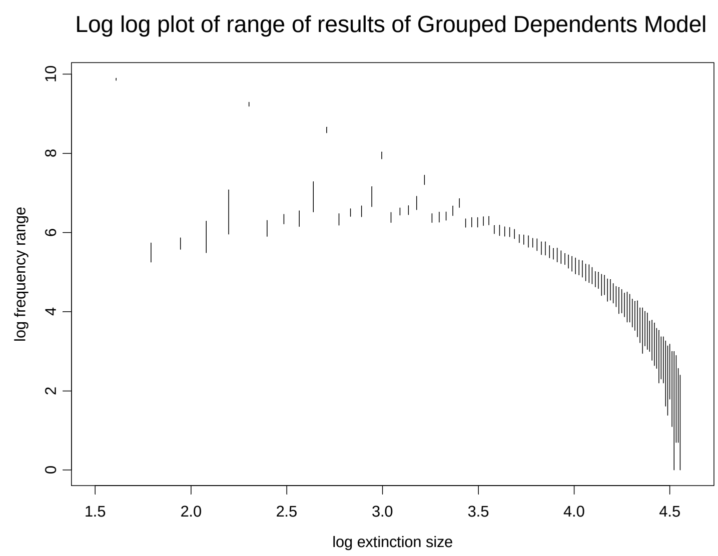

Figure 3: each line is drawn between the minimum and maximum results from 100 runs of the model for each extinction size.

We can see that the range of results for each extinction size is small[18]. Looking at the observations in more detail, they are generally normally distributed about the mean.

We next show the result for the grouped dependents model with small continuous updates with the same inputs.

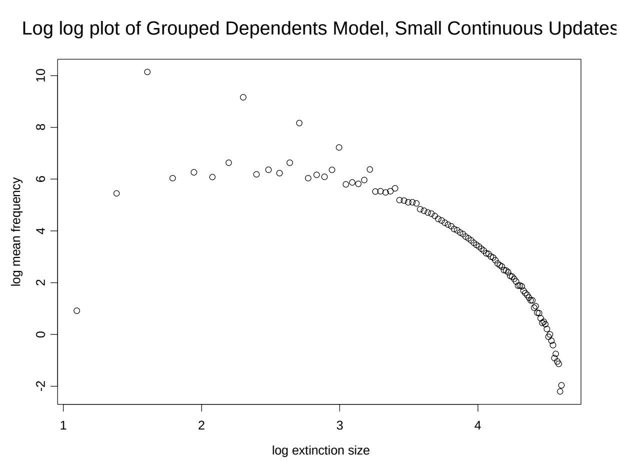

[18] Due to the infrequency of large extinction sizes the range of results increases as the extinction size increases.

Figure 4: Mean results over 100 simulations of the Newman Model with Grouped Dependents and Small Continuous Updates.

These results appear similar to the results for the standard Newman Model on a network but there are smaller numbers of large extinctions. Again they do not follow a power law.

We now turn our attention to the Fight for Survival model described in 2.3.2. In this case we plot the results for the model with and without Small Continuous Updates on the same graph.

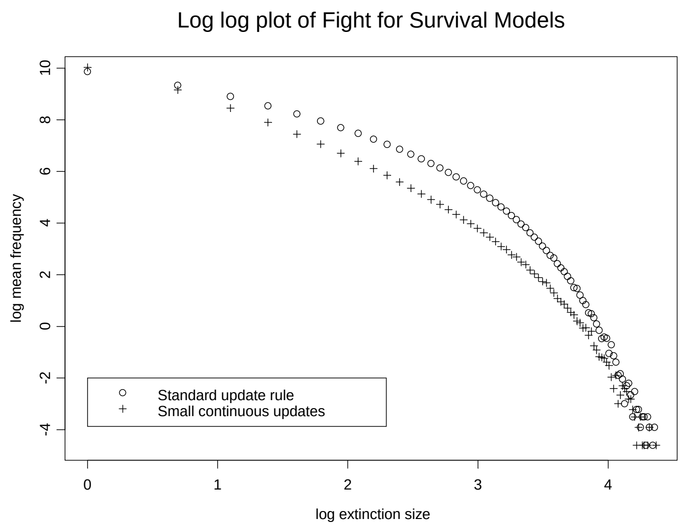

Figure 5: Results for the fight for survival models on a small world network.

This does not have the jumps in survival at intervals of 5 due to the fact that firms neighbours only have their strategies redrawn rather than also becoming extinct. Again the results do not follow a power law[19].

[19] They are closer to a power law - see discussion.

3.2 Varying the type of network

In this section we see the effect of different types of network on the results. Firstly we consider four different classes of network: small world, random, k-nearest neighbour and torus. These networks are described in [5].

We run each of these networks such that each firm has on average four neighbours.

The following results are from the Grouped Dependents model with standard updates [thus the results for the small world network are the same as in Figure 2].

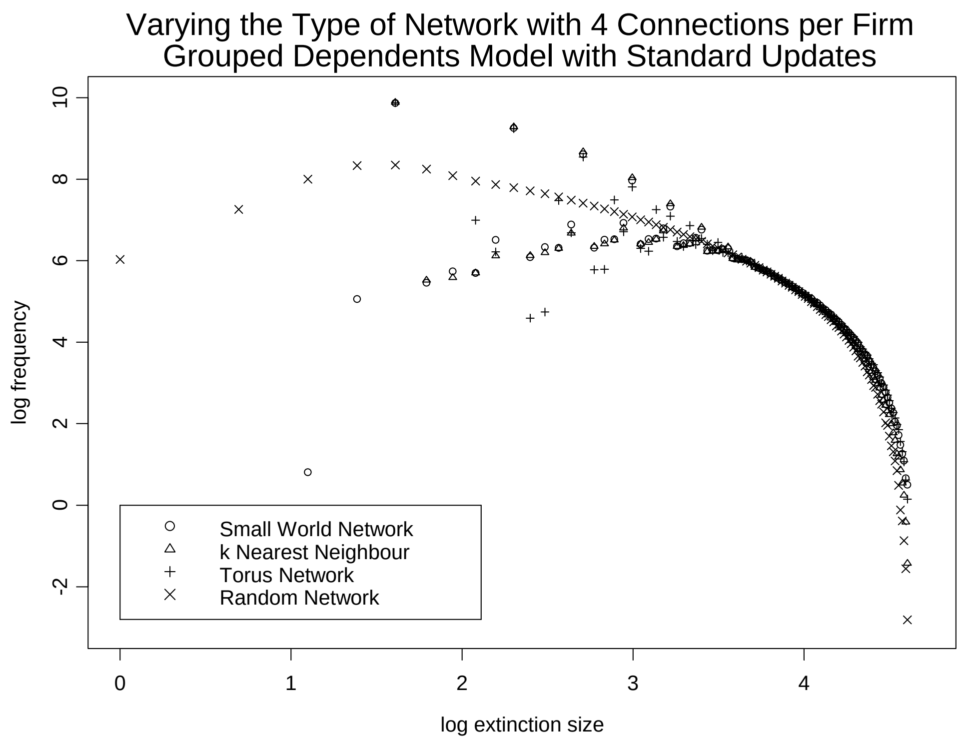

Figure 6: Examining the effects of different types of networks on the results of the Grouped Dependents Model with standard updates.

The results are very similar for large extinction events. For smaller extinction events the results differ due to the number of connections for each firm. In the torus network and k-nearest neighbour networks, each firm has exactly four neighbours, in the small world network, almost all firms have four neighbours, whereas in the random network, the number of neighbours is a binomial distribution with mean four. Thus the range of numbers of neighbours explains the results showing a much smoother curve for the random network. All other networks give very similar results, with differences only occurring due to the amount of overlapping connections.

The four different networks have very different properties apart from each having the same average number of connections per firm. Thus from this and other results we have examined it is possible to conclude that it is essentially only the number of neighbours of each firm which affects the results.

The other three adaptations of the Newman Model studied in this paper give similar results.

3.3 Varying the number of neighbours per firm

In this section we consider the effect of changing the number of neighbours of each firm. We do this for a small world network (and the results for other networks are similar). In the figure below we have used the Fights for Survival Model with small continuous updates.

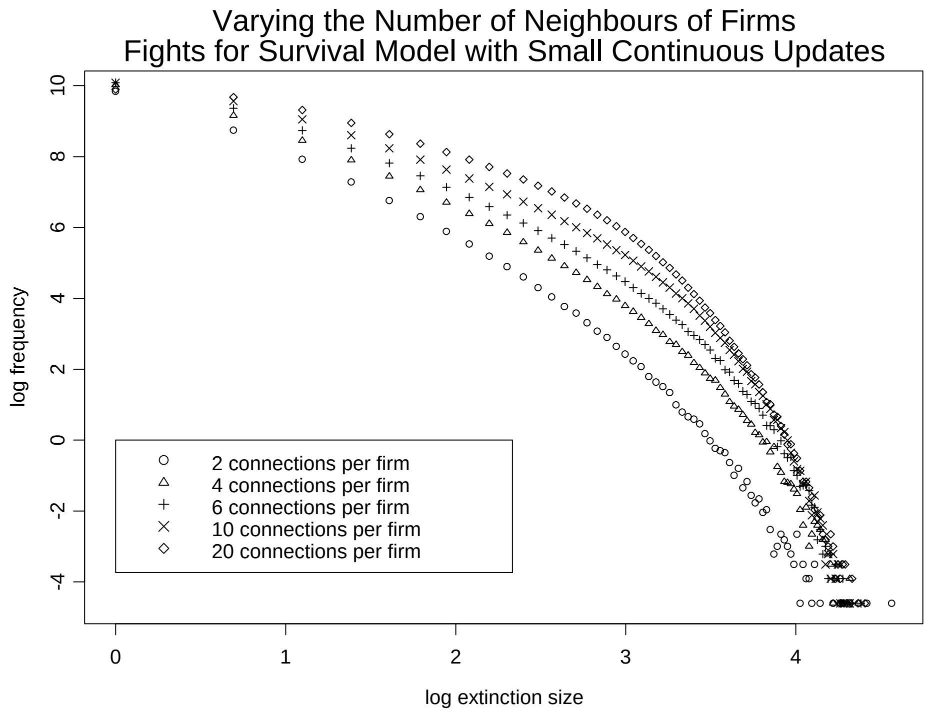

Figure 7: Varying the number of neighbours of firms with the fights for survival model with small continuous updates on a small world network.

We see that increasing the number of neighbours of each firm increases the frequency of large extinctions and makes the results look less like a power law.

3.4 Using stress levels which are distributed normally

As a final comparison, we consider the results of the model when using a normal stress distribution rather than an exponential. Here we use the normal distribution with zero mean and a standard deviation of 0.05 and compare it with the exponential distribution with mean 0.05. The model used is the Grouped Dependents Model with standard updates so the exponential results are the same as in Figure 2.

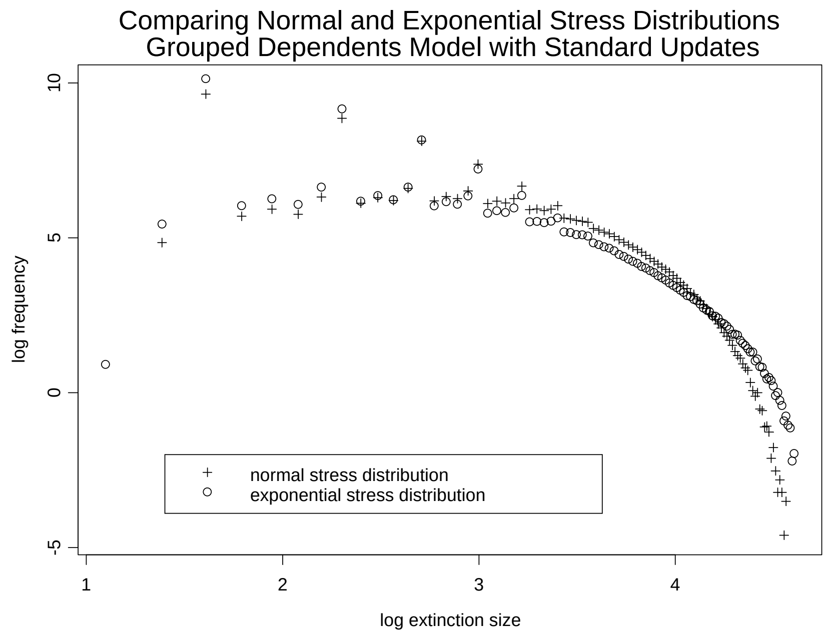

Figure 8: Comparison of results using a normal stress distribution with mean 0 and standard deviation 0.05 and an exponential stress distribution with mean 0.05

We can see that the frequencies of extinction sizes are similar but that the results from using the normal distribution are more curved.

4 Discussion

The idea behind this paper was to see whether a power law distribution for the frequency of extinction size events could be obtained by placing variants of the Newman Model onto a network. The results above show that it is not the case.

When we look at possibilities for the results following other distributions, we find that the distribution fits much closer to an exponential. The following graph shows the results of plotting the Grouped Dependents Model with the standard update rule along with a fitted exponential distribution. The results are the same as in Figure 2 except that the data has been binned by combining each set of five points to smooth out the jumps observed in Figure 2.

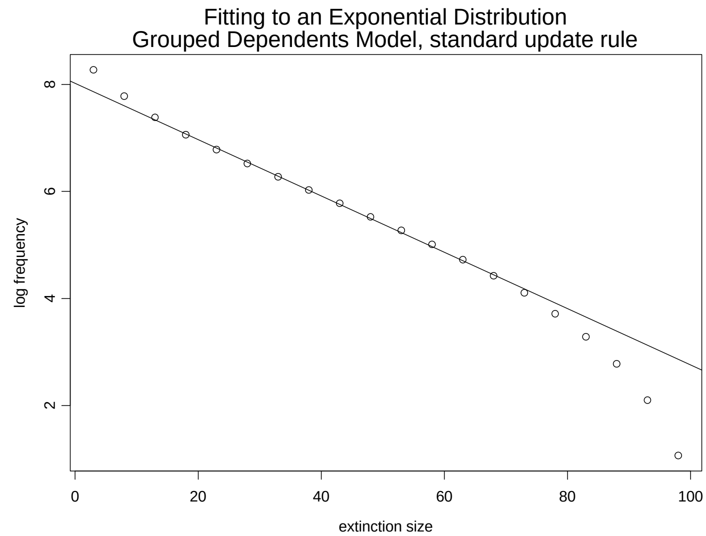

Figure 9: Results from the Grouped Dependents Model with standard updates with the line being the fit of the middle ten points to an exponential distribution.

We can see that whilst the distribution is not exponential, the middle ten points fit much better to an exponential distribution than to a power law.

The fit becomes more like an exponential distribution when we consider the Fights for Survival Model with Small Continuous Updates as shown below.

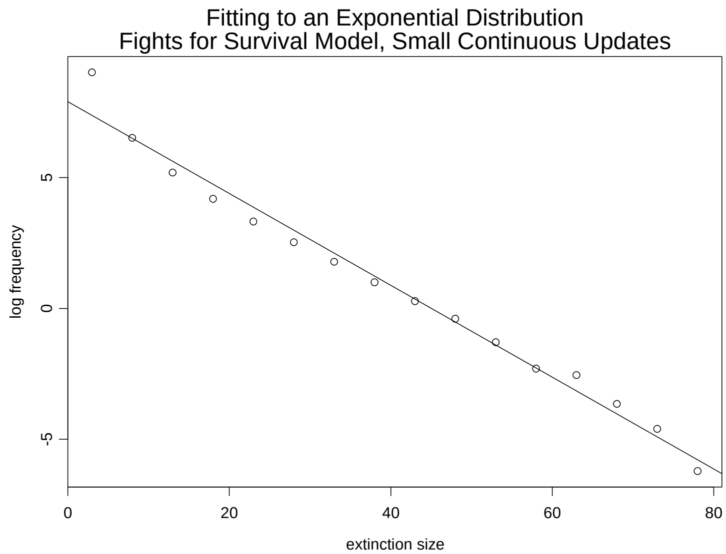

Figure 10: Results from the Fights for Survival model with Small Continuous Updates with the line being the fit of all the points to the exponential distribution.

If we consider the limit where all firms have their stress tolerances updated by a draw from the uniform distribution every period, then in this case it is simple to show that the results for the frequency of extinction sizes will have the same distribution as the distribution of stress levels.

By placing the firms on a network and having the extinction of a firm affect the firm's neighbours means that more agents will have their stress tolerances updated each period than would otherwise occur in the standard Newman Model. Thus we should expect (and we do observe) the distribution of the results to tend to the distribution of stress levels.

It is therefore our prediction (yet to be tested) that if we applied a power law stress distribution, we would observe results which were a better fit to a power law. Economically, we do observe that firms frequently change their strategies in response to stress. Although there are no measurable results, it does not seem an unreasonable idea that stress levels on firms might follow a power law as large numbers of naturally occurring results do. This could well explain the power law for extinctions of firms observed in [3].

5 Conclusions

In this paper we have examined whether variations of the Newman Model of extinctions on networks can explain the power law behaviour observed in the extinction rates of firms. The results we have obtained show that this does not seem to be the case.

We examine the model in which agents are connected by a variety of different networks. But, essentially, the results are very similar whichever network is specified, and it is the average number of connections per agent which matters.

6 References

1. P. Ormerod, H. Johns and L. Smith (February 2002), ‘Patterns of Agent Extinction under External Shocks’. Available at www.volterra.co.uk.

2. M. E. J. Newman (1997), 'A model of mass extinction', *J. Theor. Biol.*, **189**, 235-252

3. 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.

4. M. E. J. Newman and R. G. Palmer (1999), 'Models of Extinction: A Review', adap.org/9908002 submitted to Philosophical Transactions of the Royal Society.

5. Ormerod and Smith (2000), ‘Access to Financial Services in the UK and the Topologies of Social Networks’. Paper published by Britannia Building Society. Available on www.volterra.co.uk.

6. D. J. Watts and S. H. Strogatz, 'Collective Dynamics of "Small-World" Networks', *Nature*, vol. 393, 4 June 1998.

7. K. Sneppen and M. E. J. Newman, ‘Coherent Noise, scale invariance and intermittency in large systems’, (1996), cond-mat/9611229.

Appendix 5

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

***** 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

We develop a theoretical 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.

The model contains N agents, and all pairs of agents are connected to each 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 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.

The properties of the model conform closely to the empirical evidence.

The paper raises the possibility that there are general mechanisms at work which account for extinctions of agents.

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[20]. 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[21].

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

[20] 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.

[21] 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

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.

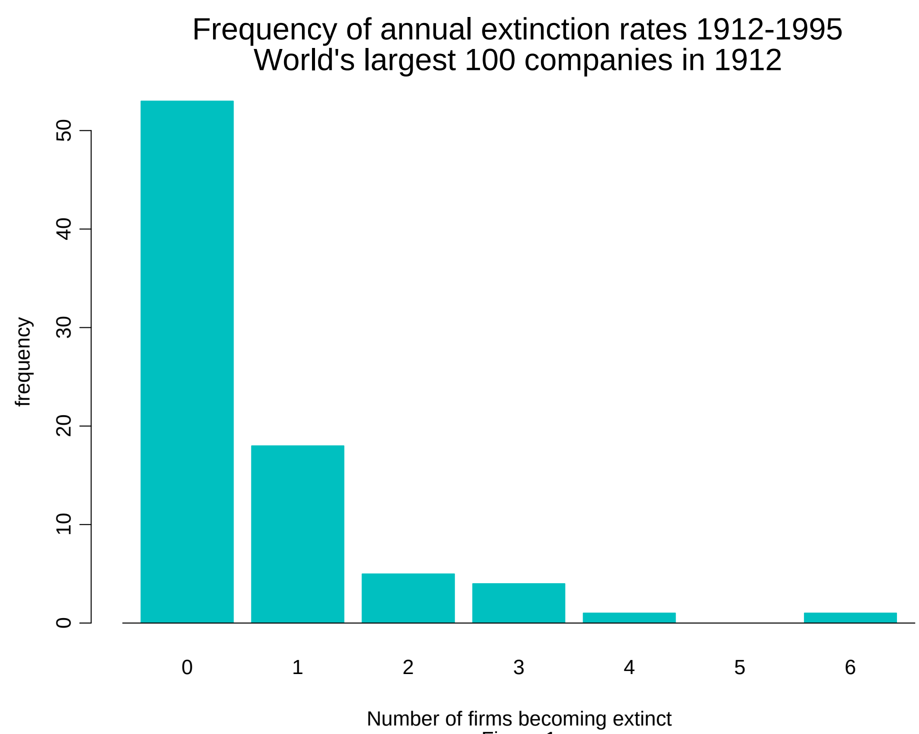

Figure 1

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[22] 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 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

[22] 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

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}[23].

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 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.

[23] 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

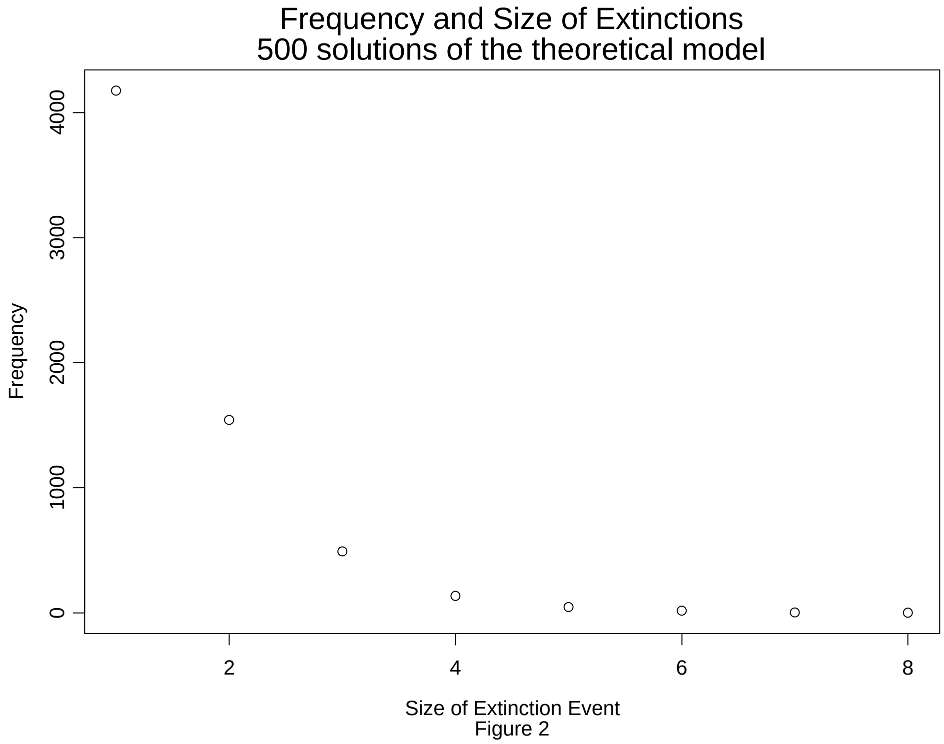

Figure 2: plot showing the power law relationship between extinction size and frequency. The frequency data show the number of times over the 500 solutions of the model in which the particular number of agents become extinct in any given period

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

Appendix 6

External Shocks in a Model of Endogenous Agent Extinctions

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

***** 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

We examine a model of agent extinction which attributes extinction events to the interconnectedness of species within an ecosystem. Both biological extinctions and the extinctions of companies in an economy follow a power law pattern which is replicated by the model. This paper considers whether applying external influences (shocks) on the system alters the extinction frequency distribution. It is found that under this extension to the model, its properties remain essentially unchanged.

1. Introduction

Analysis of the fossil record has shown that the sizes of extinction events follow a statistical distribution Φ which is compatible with a power law fit:

$$\Phi(S) \approx \alpha S^{-\beta}$$

where S is the size (in terms of species or families) of an extinction event and α and β are constants [1]. Several models have been created to try to simulate this power law behaviour [2]. These follow one of two different approaches; by either proposing that the observed pattern of extinctions arises from the internal mechanisms of the system (endogenous causes) or from purely external factors (exogenous causes).

Parallels can be drawn between the behaviour of species interacting in an ecosystem and that of companies interacting in an economy [3]; the symbiotic and competitive relationships which evolve can be seen as analogous. Analysis of extinction rates of the world's 100 largest companies[24] over the period 1912 to 1995 has suggested that company extinctions also exhibit a similar power law distribution [3]. The question arises, then, of whether biological extinction models can provide an insight into the causes of firm extinctions, and if so, which approach should be taken.

Here we have followed the model of Solé and Manrubia [4], which is outlined in detail in section 2. Like many models attempting to simulate endogenous mechanisms, the model exhibits the property of self-organised criticality [5], with the implication that extinctions occur, without the need for external stimuli, on any scale as the result of systematic microscopic changes. Some smaller extinction events may "avalanche" into cataclysms because they involve the extinctions of certain keystone

[24] As measured in 1912

agents upon which many others are dependent. A discussion of the application of the Solé-Manrubia model to the empirical company extinction data is given in [6]. This paper continues by applying external shocks to the agents in the Solé-Manrubia model.

We have also examined the Newman model [7,8], which obtains an appropriate extinction power law distribution solely by considering external stresses on the population, without explicit links between agents.

The real world is evidently more complicated than the case where *either* endogenous *or* exogenous factors are the sole causes of extinction events, and it is highly likely that real-life agents in either ecosystems or economies are prey to a mixture of both. External factors could easily exacerbate the avalanching of an extinction event.

2. Description of the model

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. This matrix encapsulates the relationships between different agents as symbiotic (J_{ij}, $J_{ji} > 0$), mutually competitive (J_{ij}, $J_{ji} < 0$) or predatory ($J_{ij} < 0$, $J_{ji} > 0$).

The model can be summarised as follows.

1. J_{ij} is initialised

2. Each agent has one of its J_{ij} updated, i.e. assigned with a new value 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.

In order to simulate the effect of external shocks on the system a random variable $X(t)$ is subtracted from the fitness of each agent such that

$$F_i(t) = \left(\sum_{j=1}^{N} J_{ji}(t) \right) - X(t)$$

For the purposes of comparison, $X(t)$ is drawn from one of two distributions, either:

1. A truncated power law distribution[25] $P(X)$:

$$P(X) \begin{cases} = uX^{-w} & X_{\min} < X < X_{\max} \\ = 0 & \text{otherwise} \end{cases}$$

   This seems reasonable as the frequency of many natural and sociological cataclysmic phenomena (e.g. earthquakes, wars) appear to approximate to power law distributions [9].

2. A normal distribution characterised by a mean μ and standard deviation σ. If $\mu \sim \sigma$, agents will of course experience an appreciable number of beneficial shocks as well as harmful shocks.

We examine the effect of the model of

- the shock in each period being common to all agents
- each agent receiving its own specific shock in each period

---

[25] If X_{min} and X_{max} are specified then u can be calculated from the condition that $\int_{X_{\min}}^{X_{\max}} P(X)\,dX = 1$

3. Model results

3.1 Without shocks

The results presented here are for N = 100 agents and are averaged over 50 simulations each. Each simulation comprises of n=50,000 iterations of the model, of which the first n'=10,000 are discarded in order to eliminate behaviour arising from particular initial conditions.

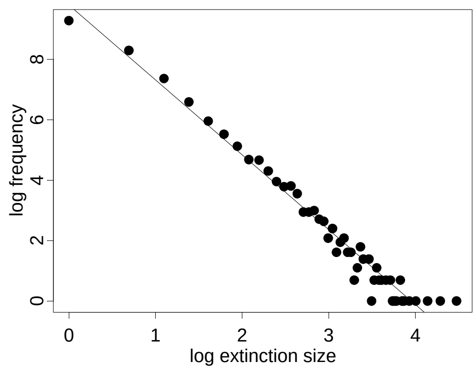

Figure 1 Log extinction size vs. log frequency of extinction size for a typical simulation.

Without external shocks, the model exhibits the desired power law relationship between extinction size S and frequency $\Phi(S) \approx \alpha S^{-\beta}$ (fig 1).

A measure of the total fitness of the system $\overline{F}$ is given as:

$$\overline{F} = \frac{\sum_{i=1}^{N} \sum_{t=n'+1}^{n} F_i(t)}{(n-n')N^2}$$

i.e. the mean agent fitness as a proportion of its maximum possible value. The average values of β and $\overline{F}$ over 50 simulations of the model are β= 2.43 and $\overline{F}$ = 0.118. The average sum of the square of the residuals of the line of best fit to the power law G = 11.1

We now examine the model results with shocks.

3.2 With common shocks drawn from a power law distribution

Fig 2 shows the log extinction size vs. log frequency plots for typical simulations using various power law parameters, while Table 1 shows the results averaged over 50 simulations for β, $\overline{F}$ and G.

Table 1

Power law parameters			**Results**		
X_{min}	X_{max}	w	β	$\overline{F}$	G
0.01	10.0	1.5	2.08	0.143	17.4
0.01	10.0	2.0	2.33	0.121	10.2
0.01	10.0	2.5	2.42	0.117	10.2
0.1	10.0	1.5	1.96	0.158	23.2
0.1	10.0	2.0	2.15	0.139	15.0
0.1	10.0	2.5	2.30	0.124	10.4
1.0	10.0	1.5	n/a[26]	n/a	n/a
1.0	10.0	2.0	n/a[3]	n/a	n/a
1.0	10.0	2.5	2.04[27]	0.151	23.8

The extinction size vs. frequency relationship departs from a power law the more probable the larger shocks (i.e. the lower the magnitude of w or the higher X_{min}). In these

[26] Total extinction occurs on every simulation.
[27] Total extinction occurs the majority of simulations.

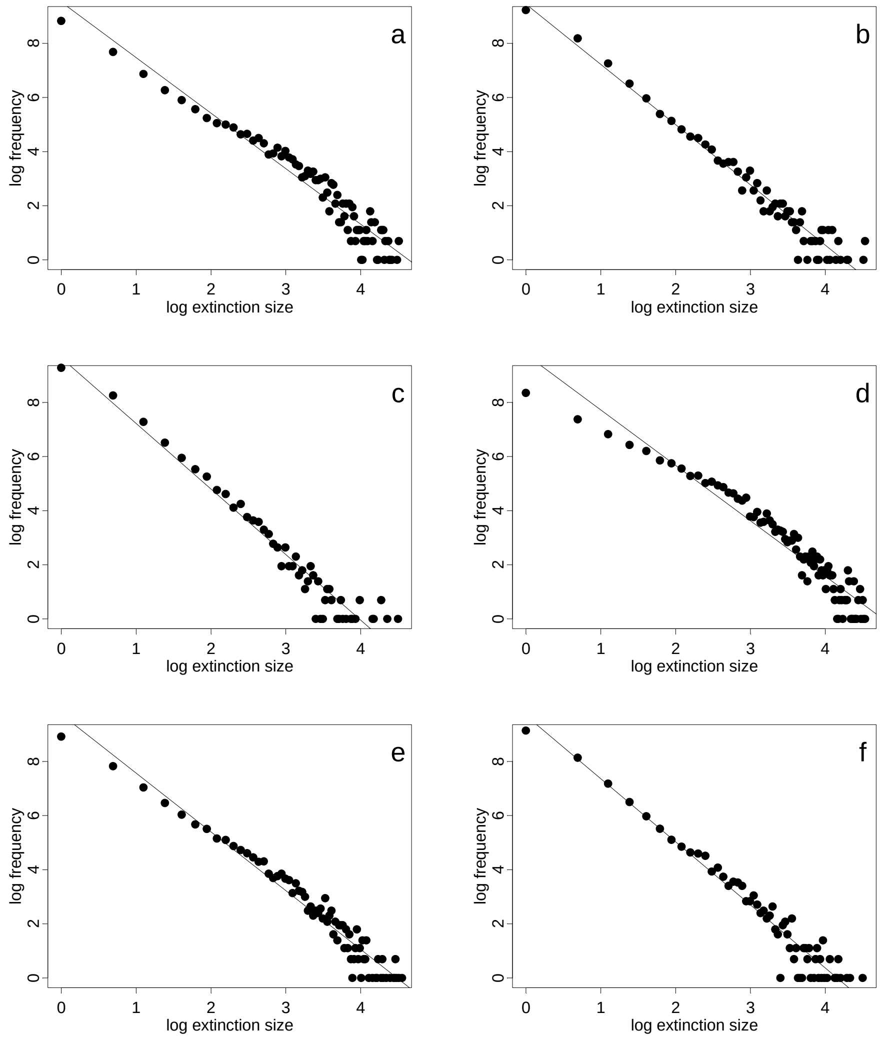

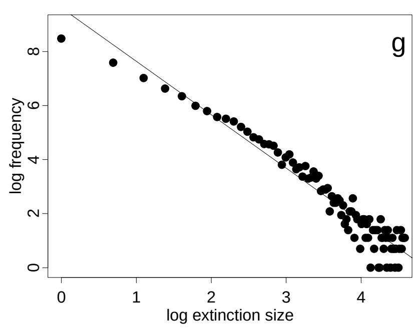

Fig 2 Log extinction size vs. log frequency of extinction size for typical simulations
(a) X_{min}= 0.01, X_{max}=10.0, w=1.5
(b) X_{min}= 0.01, X_{max}=10.0, w=2.0
(c) X_{min}= 0.01, X_{max}=10.0, w=2.5
(d) X_{min}= 0.1, X_{max}=10.0, w=1.5
(e) X_{min}= 0.1, X_{max}=10.0, w=2.0
(f) X_{min}= 0.1, X_{max}=10.0, w=2.5
(g) X_{min}=1.0, X_{max}=10.0, w=2.5

cases, the power law overpredicts the number of small extinctions which should be taking place.

The average fitness is increased with shock magnitude and large shock frequency. Applying shocks is equivalent to raising the fitness threshold for extinction. Only the very fittest will survive; therefore the parents of replacements will on average be fitter, thereby imparting a greater average fitness to the population.

Fig 3(a) shows the relationship (correlation coefficient ≈ 0.4) between extinction size and shock size for a typical simulation[28]. The positive correlation infers that larger shocks do indeed play a role in causing or exacerbating extinctions; however, the largest shocks are not necessarily coincident with the largest extinctions, many of which coincide with shocks of the smallest size. Smaller shocks are more numerous and so statistically more likely to coincide with a large extinction resulting from purely endogenous causes.

Fig 3(b) introduces a lag to the extinction size to see whether shocks cause extinction avalanches. The correlation coefficient between $X(t)$ and extinction size at iteration t +

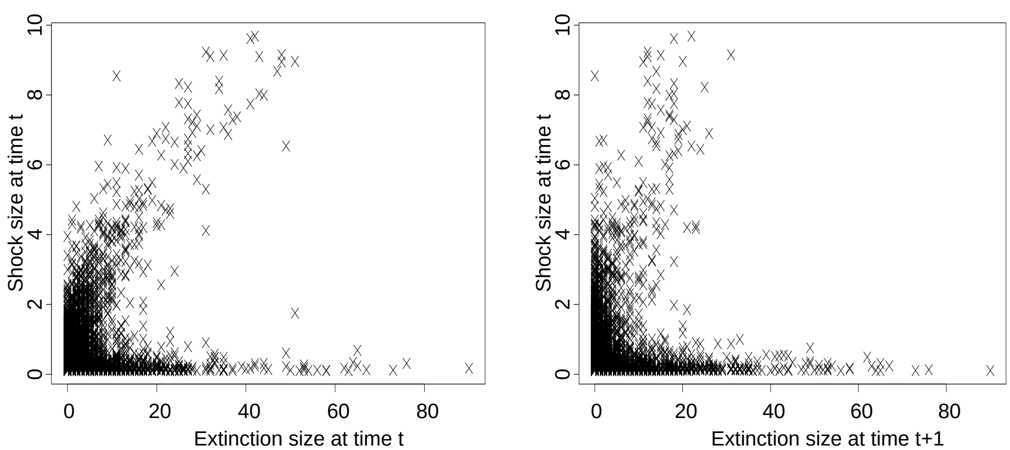

Figure 3 (a) Plot showing relationship between extinction event size and shock size for a typical simulation using power law shocks (b) For the same simulation, plot showing extinction size vs. the shock size from the previous iteration

[28] X_{min} = 0.1, X_{max} = 10, w = 2.5

1 is much smaller (≈ 0.16); however the large shocks are coincident with subsequent extinctions of moderate size.

3.3 With common shocks drawn from a normal distribution:

Fig 4 shows the log extinction size vs. log frequency plots for typical simulations using various power law parameters, while Table 2 shows the results averaged over 50 simulations for β, $\overline{F}$ and G.

Normal disbn parameters		**Results**		
μ	σ	β	$\overline{F}$	**G**
0	0.5	2.45	0.120	11.3
0	1.0	2.42	0.128	13.9
0	2.0	2.34	0.147	19.0
1	0.5	2.36	0.121	13.1
1	1.0	2.37	0.129	15.8
1	2.0	2.26	0.147	23.3
2	0.5	2.26	0.121	18.0
2	1.0	2.29	0.129	18.3
2	2.0	2.17[29]	0.147	26.8

Table 2

[29] Total extinction occurs the majority of simulations.

It appears that the larger σ is, the greater the deviation from the power law. Increasing μ also causes deviation, but increasing the likelihood of large shocks (i.e. increasing σ) is more influential. Again, $\overline{F}$ is increased by raising the likelihood of large extinctions, for the same reason argued above.

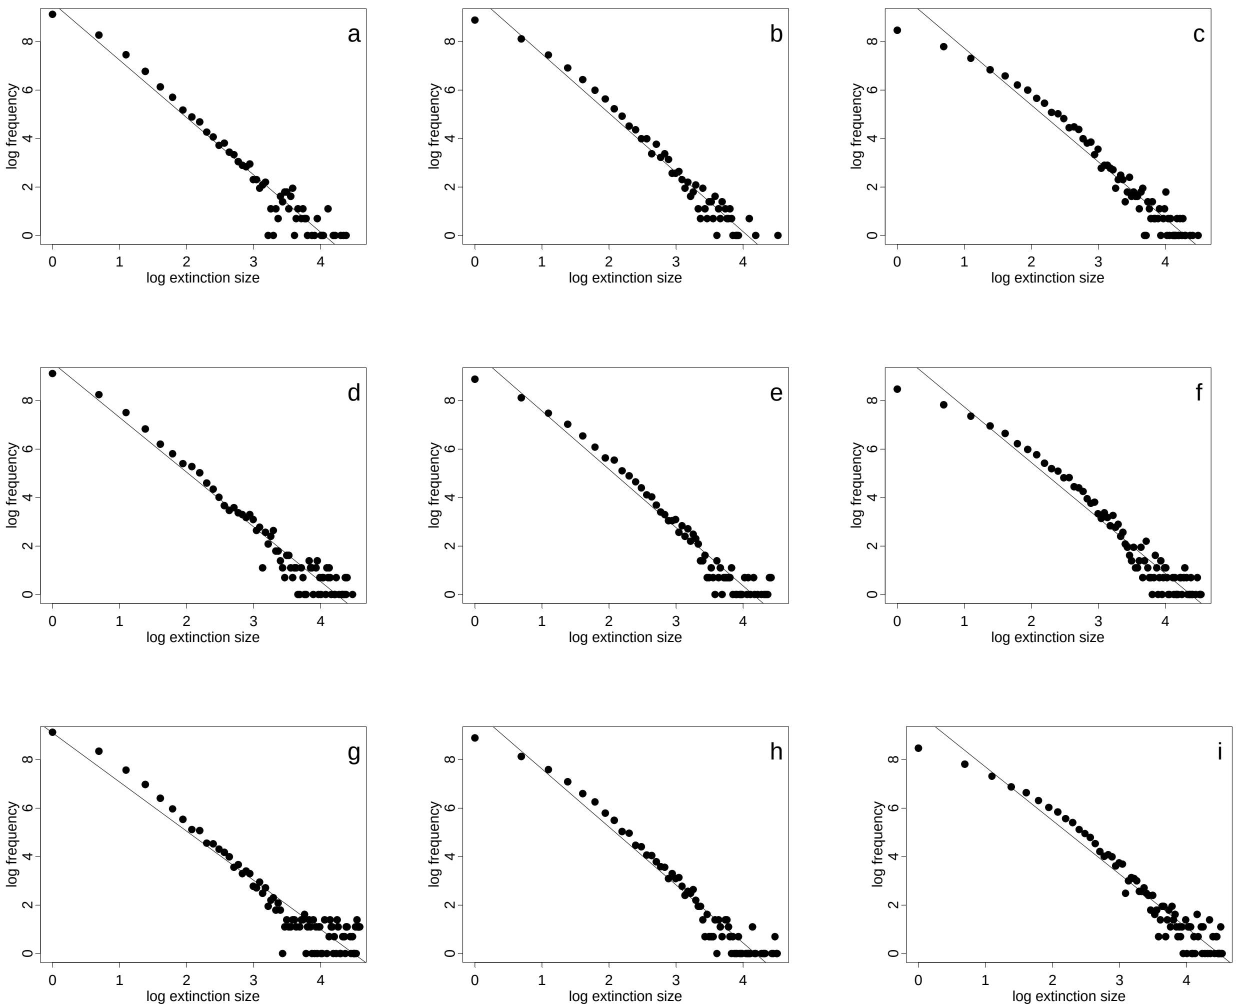

Figure 4 Log extinction size vs. log frequency of extinction size for typical simulations
(a) $\mu = 0$, $\sigma = 0.5$ **(b)** $\mu = 0$, $\sigma = 1$
(c) $\mu = 0$, $\sigma = 2$ **(d)** $\mu = 1$, $\sigma = 0.5$
(e) $\mu = 1$, $\sigma = 1$ **(f)** $\mu = 1$, $\sigma = 2$
(g) $\mu = 2$, $\sigma = 0.5$ **(h)** $\mu = 2$, $\sigma = 1$
(i) $\mu = 2$, $\sigma = 2$

3.4 Individual agent shocks drawn from a power law

Clearly, not every agent might react identically to an external shock. Here, $X(t)$ is replaced by $X_i(t)$, i.e. each agent receives an individual shock (drawn from a power law distribution).

Table 3

Power law parameters			**Results**		
X_{min}	X_{max}	w	β	$\overline{F}$	G
0.01	10.0	1.5	2.52	0.127	14.0
0.01	10.0	2.0	2.42	0.118	11.6
0.01	10.0	2.5	2.42	0.117	11.8
0.1	10.0	1.5	2.65	0.137	15.0
0.1	10.0	2.0	2.54	0.125	11.8
0.1	10.0	2.5	2.46	0.119	11.3
1.0	10.0	1.5	2.47	0.141	18.8
1.0	10.0	2.0	2.45	0.138	18.4
1.0	10.0	2.5	2.45	0.133	17.9

The relevant plots are not included in order to save space. In general, the power law underpredicts the number of small extinctions (in this respect applying individual shocks differs from applying a general shock). $\overline{F}$ increases, but to a much less marked extent, as fewer weak potential parents are wiped out.

3.5 Individual agent shocks drawn from a normal distribution

As above, each agent receives an individual shock, this time drawn from a normal distribution.

Normal disbn parameters		Results		
μ	σ	β	$\overline{F}$	G
0	0.5	2.42	0.117	11.3
0	1.0	2.51	0.121	10.7
0	2.0	2.66	0.131	10.8
1	0.5	2.42	0.117	10.4
1	1.0	2.43	0.122	12.0
1	2.0	2.57	0.132	12.7
2	0.5	2.29	0.120	16.7
2	1.0	2.34	0.122	14.8
2	2.0	2.50	0.132	13.6

Again, the power law underpredicts the number of small extinctions. $\overline{F}$ follows a similar, but less pronounced, pattern to the case when general shocks are applied. The goodness of fit to the power law is not greatly altered.

4. Conclusions

A model of endogenous agent extinction has been presented and modified so that agents experience in addition external shocks. We examine shocks drawn from power law and Gaussian distributions. Further, we examine a version of the model in which a shock is common to all agents in any given period, and a version in which each agent receives its own specific shock in each period.

In general, the basic properties of the model are not affected. The relationship between the frequency and size of extinctions which exists in the model without shocks is very similar to that of the model including shocks.

We show that the power law relationship between extinctions size and frequency normally shown by the model degenerates (but not sufficiently to negate the power law characteristics) when the expected value of shocks increases, or the likelihood of large shocks increases.

Adding external shocks to the model in the manner described here is equivalent to raising the threshold fitness for survival. The average fitness of the system increases after larger shocks have been applied. This is due to the fitness of parents increasing on average as large shocks wipe out weaker potential parents. This effect is not as marked when individual shocks are applied, as only a selection of weaker parents will be made extinct.

There is a moderate correlation between shock size and extinction size. Larger shocks do not necessarily lead to the largest extinctions because they may strike at a time when the system is particularly robust. There is little correlation between shock size and extinction size in the subsequent iteration. This is because the agents tipped over the extinction threshold by shocks may or may not be heavily influential on the others, i.e. they may or may not be keystone agents. Extinctions of moderate size appear to be common in the iterations following large shocks; although there will be many replacement agents with high fitness, some keystones may have been wiped out.

References

[1] R.V. Solé and J. Bascompte (1996), *Are critical phenomena relevant to large-scale evolution?*, Proc R. Soc. London B 263, 161

[2] B. Drossell (2001), *Biological evolution and statistical physics*, Advances in Physics 50:2

[3] 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

[4] R.V. Solé and S.C. Manrubia (1996), *Extinction and self-organized criticality in a model of large-scale evolution*, Phys. Rev E 54:1 R42

[5] P. Bak (1996), *How Nature Works: The Science of Self-Organised Criticality*, Copernicus Press

[6] P. Ormerod, H. Johns and L. Smith (2001), *An Agent-Based Model of the Extinction Patterns of Capitalism's Largest Firms*, www.volterra.co.uk

[7] M. E. J. Newman (1997), *A model of mass extinction*, J. Theor. Biol., 189, 235-252

[8] P. Ormerod, H. Johns and L. Smith (2002), *Patterns of Agent Extinction under External Shocks*, www.volterra.co.uk

[9] M. Buchanan (2000), *Ubiquity*, Weidenfeld & Nicolson

Comment: