The Impact of Regulation in a Model of Evolving, Fitness-Maximising Agents

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

September 2001

***** Corresponding author

We are grateful to British Telecommunications plc for financial support towards this research

Abstract

The process of evolution of biological species has recently been analysed in formal models of complex systems. Individual species interact with each other, sometimes in co-operative symbiosis and sometimes in direct competition. The fitness level of each species evolves over time, and species whose fitness falls below a critical level become extinct. Variants of surviving species enter the eco-system in the niches vacated by extinct species.

The overall properties of complex systems such as these emerge from the interactions of the individual species. A key property is that of self-organised criticality. The system becomes tuned so that extinctions on any scale can occur at any point in time. The probability of observing an extinction of any given size falls away not just with the size but with the size raised to a power of itself. The mathematical expression which describes the relationship between the size of an extinction and its frequency is known as a power law. Models based on the principle of agents interacting with each other within a complex system provide a good description of the empirical fossil record on species extinction.

Complex systems such as these can be interpreted naturally in the context of the interactions of products in a market, or firms in an industry. Many of the interactions are competitive, but some are complementary. Firms operating under uncertainty seek to maximise fitness levels, whether for themselves or of their products. Firms or products which fall below critical levels of fitness become extinct.

In this paper we investigate the impact of regulation in a self-organised, complex system containing interacting agents. In particular, we examine the effect of reducing the proportion of complementary interactions between agents, and replacing them with directly competitive ones. This is the principle which underlies a great deal of regulatory activity in the West. We find that, initially, as the degree of competition is increased, the properties of the system do not change a great deal. Beyond a certain point however, there is a qualitative change in these properties. The frequency of large extinctions of agents rises sharply, and the average lifetime of agents falls markedly. Agents do not survive for long enough to build up fitness levels, and the overall fitness level of the system is reduced. The over-enforcement of competitive behaviour damages the system as a whole.

1. Introduction

The concept of self-organised criticality in complex systems is being applied increasingly to a wide variety of apparently disparate systems in the natural, biological and social sciences. 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. A good, non-technical discussion is given by Buchanan [1].

In a recent review, Turcotte [2] cites many examples of natural systems which might have this property, such as landslides, earthquakes and forest fires. Findings in the fossil record of the extinction of biological species suggest that the co-evolution of species in the ecosystem gives rise to self-organised criticality. Drossel [3] gives a detailed review of this empirical and analytical literature.

In the human world, several recent papers have found evidence of this phenomenon in the distributions of both links in web pages and in the growth and evolution of Internet providers (see, for example, Capocci et.al. [4]). It is also hypothesised that the evolution of human language possesses this property (e.g. Dorogovtsev and Mendes [5]). Ormerod and Mounfield [6] show that the distribution of economic recessions in the developed world is consistent with the view that the economy as a whole has the property of self-organised criticality. Ormerod and Smith [7] provide evidence that the extinctions of large firms in the twentieth century follows a pattern which is very similar to that of biological species extinction.

In this paper, we describe a general model of interacting agents operating in a complex system, which has the property of emergent, self-organised criticality. The agents interact in ways which lead to increases or decreases in their level of fitness. The model has a natural interpretation in terms of interacting products in a market or of firms in an industry. It is also very similar to one which has been used to account for the extinction record of biological species.

The main purpose of the paper is to investigate the sensitivity of the properties of the complex system to external regulation. We examine the impact on the complex system of a process of regulation which increases the number of competitive interactions between agents and decreases the complementary ones. An outcome of particular interest is the overall level of fitness which emerges for the complex system as a whole under different degrees of regulation.

Section 2 gives some background information on the concept of self-criticality. Section 3 describes the general complex system model of interacting agents, and Section 4 sets out its properties. Section 5 describes the impact of regulation on the properties of the model. Section 6 gives the conclusions.

2. Self-organised criticality, emergent properties and power law behaviour in complex systems

2.1 Self-organised criticality

The simplest conceptual model of self-criticality in complex systems was introduced by Bak et.al. [8] in 1987. Their model is that of a sandpile[1]. Individual, identical grains of sand are dropped slowly, at regular intervals, onto a flat surface. Initially, the slope of the pile is small. Occasionally, dropping a new grain will trigger a small avalanche of grains. As more grains get added to the pile, the pile becomes steeper. Large avalanches start to occur and the mean size of avalanches grows.

The system reaches a stationary state in which the average number of grains dropped onto the pile per unit of time is the same as the average number of grains leaving due to avalanches. Bak and colleagues analysed the properties of the system once the pile reached the critical state. Avalanches do not take place every single time a grain is dropped, so in some periods the pile will grow, and in others it will be reduced in size due to avalanches. But on average over time, its size will not change..

A key finding was that avalanches take place on all scales. This, itself, is a very important discovery. By construction, the proximate cause of an avalanche is always the same. Namely, the dropping of single grains onto the pile. But the consequences of these identical events vary dramatically. Sometimes, the pile grows by a single grain. Sometimes, the size of the pile remains the same, with the new grain displacing another grain from the pile. Sometimes, the pile is reduced by one grain, with the added grain displacing two others in an avalanche, and so on and so on.

In other words, when the system is in the so-called critical state, the 'common sense' connection between the size of an event and its consequences is broken. Events which by construction have the same size can have enormously different consequences. It is the interactions between the grains in the pile, how the pile is organised, which generates this property.

---

[1] Ironically, it has been shown subsequently [9] that real sandpiles exhibit behaviour which deviates to some extent from the Platonic idea of sandpiles in the Bak model. However, this does not detract from the powerful, general point which this model makes.

2.2 Power laws

Bak et.al. found that the relationship between the size of an avalanche and its frequency can be described in straightforward mathematical terms. This is the famous 'power law' relationship. Small avalanches are much more frequent than large ones, and the frequency with which a large one is observed falls away with the size of the avalanche raised to a power of itself. The probability of a truly enormous avalanche taking place, which wipes out most of the pile, is very, very small, but given sufficient time one will happen on this scale. Denoting the frequency of the event by F and its size by S, we have:

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

$$\text{or } \log(F) = \log(\alpha) + \beta\log(S) \qquad (1a),$$

where $\beta < 0$.

Systems which are in a state of self-organised criticality have properties which are characterised by power law relationships.

A number of studies show that a least squares fit of 1(a) to the biological species extinction data gives an estimate of β of around -2 (the classic study is by Raup [10]). Of course, because extinctions do not take place solely because of the critical state of the eco-system but also through external shocks such as comet and meteor strikes, the power law fit to the data is not perfect, though as Drossel [op.cit.] notes 'there seems to be good evidence that the size distribution of extinction events is not far from a power law with the exponent -2'.

As an example from the social sciences, Ormerod and Mounfield [op.cit.] examine the duration of recessions in 17 Western economies over the period 1871 - 1994. Defining a recession as a year in which real GDP growth was less than zero, there are 206 instances of a recession lasting one year, 88 examples of two-year recessions, and so on down to just one example when output fell in seven successive years. The frequency is related to the duration of the recession in years with an exponent of -1.7. Again, as with the biological example, a power law provides a good but not perfect empirical description of this process. The authors conclude that 'that there may be two separate processes going on in the process which generates data on capitalist recessions. When a recession arises, for whatever reason, agents appear to have some capacity to react quickly, which often prevents the recession from being prolonged beyond one year. Once this has happened, however, recessions can take place on all scales of duration, exactly as in the sand pile experiment'.

2.3 Emergent properties

These refer to properties of the system as a whole which emerge from the interactions of the individual agents which make up the complex system.

Clearly, in the case of inanimate objects such as grains of sand, by definition there can be no intent on the part of the individual agents to bring about the relationships which emerge at the level of the system as a whole. But even in systems in which agents can follow conscious, behavioural rules, properties of the system as a whole can emerge, even though no single agent intends to bring them about.

Economists are perfectly familiar with the concept of emergent properties, though they do not usually describe them in this way. General equilibrium theory, for example, is the central core of traditional economic theory. General equilibrium theory, as its name might imply, is concerned with the behaviour of *all* markets in an economy.

Individual agents in the system are postulated to behave in a completely self-interested way, and act to maximise their individual utility. No agent acts with any intention at all of influencing the overall properties of the system. Nevertheless, the system has emergent properties. No agent intends all markets to clear, but in general equilibrium this *emerges* from the reactions of agents to prices.

In other words, supply and demand balance in every single market, even though no individual agent intends this to happen. Further, in a single period world, this solution is also a Pareto optimum[2] . No agent can be made better off without making at least one other agent worse off. Again, agents do not intend to bring this about. The property emerges from their individual actions.

3. The general model of evolving, interacting agents and its economic interpretation

The model contains N agents. It describes how individual agents evolve over time, and how they interact with each other.

In this model, all pairs of agents are connected to each other. In an economic context, the connections can be thought of as the way in which the net impacts of the overall strategies of firms lead their individual products to interact in a particular market at any point in time. In other words, the agents in the model are products or brands. Alternatively, we might think of the agents as representing firms as a whole, and the matrix describes how their overall strategies impact on each other in an industry. The identification of agents in the model as products gives the model more general applicability, as we discuss below, though the identification of agents with firms as whole also bears scrutiny. Quite independently, Ormerod [13] developed an agent-based model of economic growth, using the same principle of agents being connected by a random network.

[2] In a multi-period world, this is in general not true: Newbery and Stiglitz [11]

The model is based upon the model of the extinctions of biological species developed by Solé and Manrubia [12]. The biological interpretation of the model is set out in [3, 12].

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 species i affects every other species 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's overall strategy applied to product i on product j, and not just the impact of relative price.

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 products benefit from each other's presence in a market. The situation could arise when, for example, two firms are independently opening up a new niche in a market. Marketing activity such as advertising by each firm creates greater awareness of the new kind of product, from which the brands of both firms can benefit. It is of particular importance in so-called network industries.

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 the fitness level of its products in a particular market (or its own overall fitness level, depending upon the particular interpretation of the model). 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 (or product) will survive and the inefficient ones perish (the classic statement of this is Alchian [14]).

The overall fitness of an agent, whether seen as a product or a firm, is measured by the sum of its connections to all other agents[3]. More exactly, it is the sum of influences on

[3] these include the connection of the product/firm to itself, as it were, the J_{ii}. A product or firm may possess qualities which lead to positive or negative effects on its own fitness. For example, a product may attempt

each agent of all agents. 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, and extinction of size m is defined to have taken place.

The overall fitness of the system is simply the sum of the fitness levels of the individual agents at any point in time. This is measured at the very start of each step of the model, before any species is made extinct. So, for example in a period when there is a large extinction, the overall fitness of the system will be low.

Extinct species are immediately replaced by new entrants. These are based upon, but are not identical to, surviving agents. Occasionally, a replacement will differ substantially from the agent on which is modelled because of the random terms in the replacement rule. In general, however, a new entrant will be fairly similar to a successful, surviving agent. In terms of the formal rule, 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.

This seems perfectly reasonable when we regard the model as representing the interactions of products in a particular market, Even in mature markets, particular products or brands fade out almost all the time, and new ones are tried out. These new ones are usually variants of products or brands which are successful, although there is a small probability that they may not be. There is no guarantee that the new entrants will be successful, for they are not completely identical to the successful one on which they are based.

Even if we think of the agents as firms as a whole rather as individual products or brands, the rule is not implausible as a description of behaviour in substantial parts of the economy. In this context, the economic interpretation obviously has more relevance in industries in which new entry is relatively easy. Entry does not appear to be too difficult in quite wide sections of the economy, particularly for firms with an existing presence in other markets who are carrying our brand extensions. For example, it is almost always the case in new industries [15]. Entry may be facilitated into more mature industries by regulatory change, (e.g. airlines and energy supply), by rapid technological innovation (e.g. the undermining of IBM by the development of the PC), or by a combination of both (e.g. telecommunications, financial services).

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

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

---

to occupy a niche for which 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.

i) for each species 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 species is measured by $f_i = \sum_j J_{ji}$, and any species for which $f_i < 0$ is deemed to be extinct. If m species become extinct, an avalanche of size m is deemed to have taken place

iii) an extinct species is replaced by a new entrant into the system, which is very similar to that of one of the surviving species. Specifically, a surviving species k is chosen at random to replace each extinct species j, and the linkages J_{ij} and J_{ji} are replaced with $J_{ik} + \eta$ and $J_{ki} + \eta$, where η is chosen at random from a small interval $[-\varepsilon, +\varepsilon]$.

4. Properties of the general model

4.1 The relationship between the size and frequency of extinctions of agents

The model is solved over a sequence of iterated steps. In the results reported, each version of the model is solved over 50,000 iterations, and the first 10,000 iterations are discarded in order to remove any transitory effects arising from the initial choice of the J matrix of connections (exactly as in [12]). The solution process is repeated over 500 separate solutions of the model in order to establish its properties.

The size distribution of extinction events in the model is described well by a power law. In other words, denoting the size of an extinction event by S (the number of species which become extinct in that step of the solution) and its frequency by F, an estimated least squares regression of the form

$$\log(F) = \log(\alpha) + \beta\log(S) \quad \text{(where we expect } \beta < 0\text{)}$$

gives a good description of the properties of the model. The frequency with which an extinction event is observed falls away with a power of its size.

The value of β depends upon the number of agents in the model, and for N sufficiently large it is close to the value of -2 observed for species extinctions in the fossil record[4]. In an economic context of the number of products in a market, it is probably more realistic to focus upon relatively small values of N. Table 1 sets out the summary statistics for values of β obtained over 500 simulations of the model, for N = 25, 50 and 100.

[4] though [3] notes that the properties of the model are not yet understood for very large N

Table 1 Summary of least squares estimates of power law exponent, 500 simulations

Number of agents			β			
N	Minimum	1st Quartile	Median	Mean	3rd Quartile	Maximum
25	3.062	3.295	3.366	3.363	3.436	3.6
50	2.729	2.877	2.922	2.92	2.964	3.09
100	2.239	2.383	2.427	2.427	2.469	2.599

Figures 1 (a) to (c) plot the values of F and S which emerge in typical individual solutions of the model for N = 25, 50 and 100 respectively, along with the relevant line of least squares fit.

For interest, the pattern of extinctions over the steps of a typical solution of the model with N = 100 is plotted in Figure 2. As can be seen, relatively long periods of 'quiet' behaviour with a low level of extinction are interspersed with occasional bursts of large extinctions.

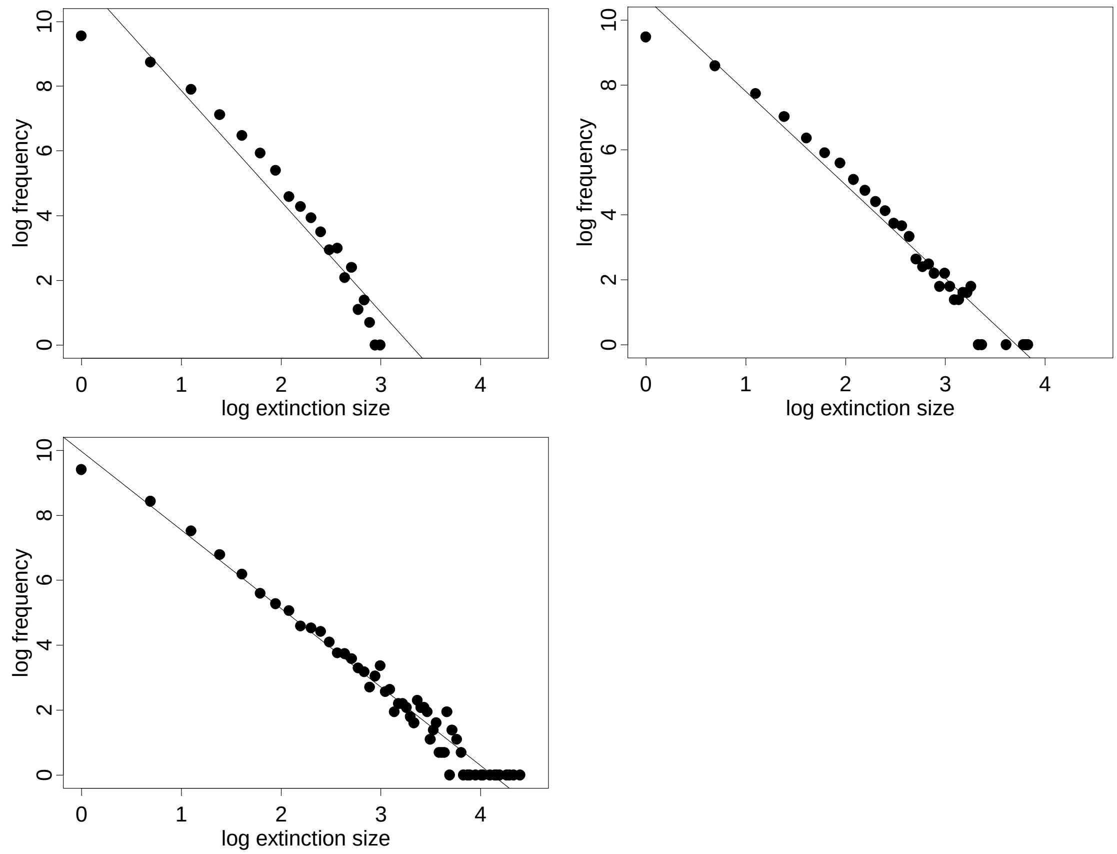

Figures 1(a) to (c). *Plot of log of frequency against log of extinction size in typical solutions of the general model, using iterations 10,001 to 50,000. N = 25 in (a) at top left; N = 50 in (b) at top right; N = 100 in (c) at bottom*

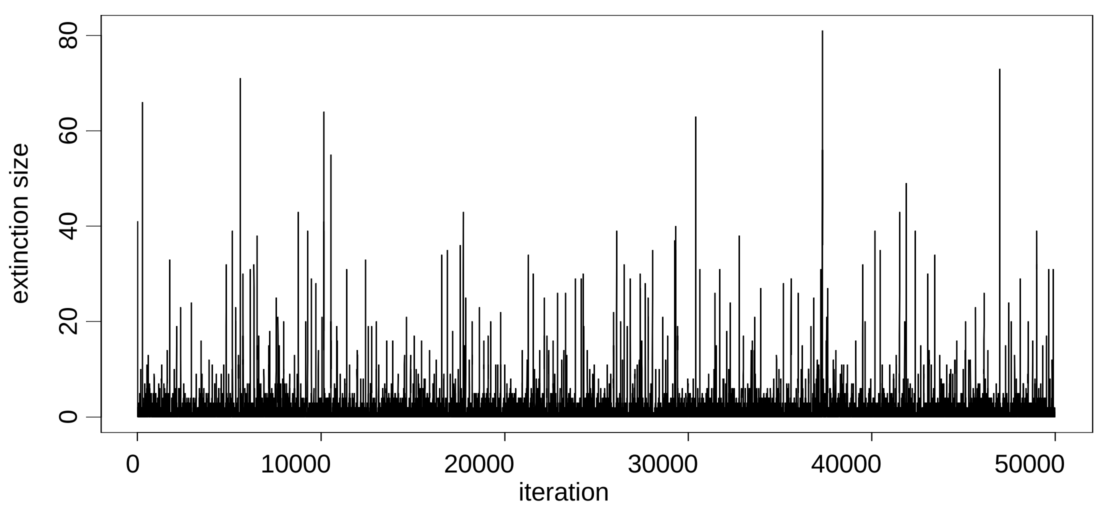

Figure 2 *Plot of extinctions over 50,000 iterations in a typical solution of the model with N = 100.*

4.2 The overall fitness of the model

We define the mean overall fitness of the model as follows. First, for each solution of the model, compute the average of the overall fitness of the system from iteration 10,001 to iteration 50,000. Second, divide this by the theoretical maximum figure for the system (i.e. N^2), in order to compare solutions with different numbers of agents. The mean overall fitness is the mean of this number over 500 separate simulations of the model.

The mean overall fitness does vary slightly with the number of agents in the model, and in general the smaller the number of agents[5], the higher the level of overall fitness - relative to the theoretical maximum - tends to be. This is due to the fact that the critical state is reached more quickly when N is small, leading to a greater frequency of extinctions. As total fitness increases after an extinction event, this has the effect of pushing up the mean overall. For N = 25 it is 0.167, for N = 50 it is 0.143 and for N = 100 it is 0.118. (For a given number of agents, the average fitness level in each solution of the model varies very little across 500 solutions).

The overall fitness of the system during a typical solution of the model with N = 100 is plotted in Figure 3. The overall level of fitness does vary over time. Periods immediately following large extinctions tend to have relatively high overall fitness, as new products are rushed in to fill the gaps opened up by the eliminations. These new entrants are variants of a surviving product, which by definition at that point in time has a fitness above zero. In other words, with the model in a self-critical state, extinctions essentially play the role described by Schumpeter in his phrase 'creative destruction'. Weaker agents are eliminated and replaced by agents which, on average, have higher levels of fitness.

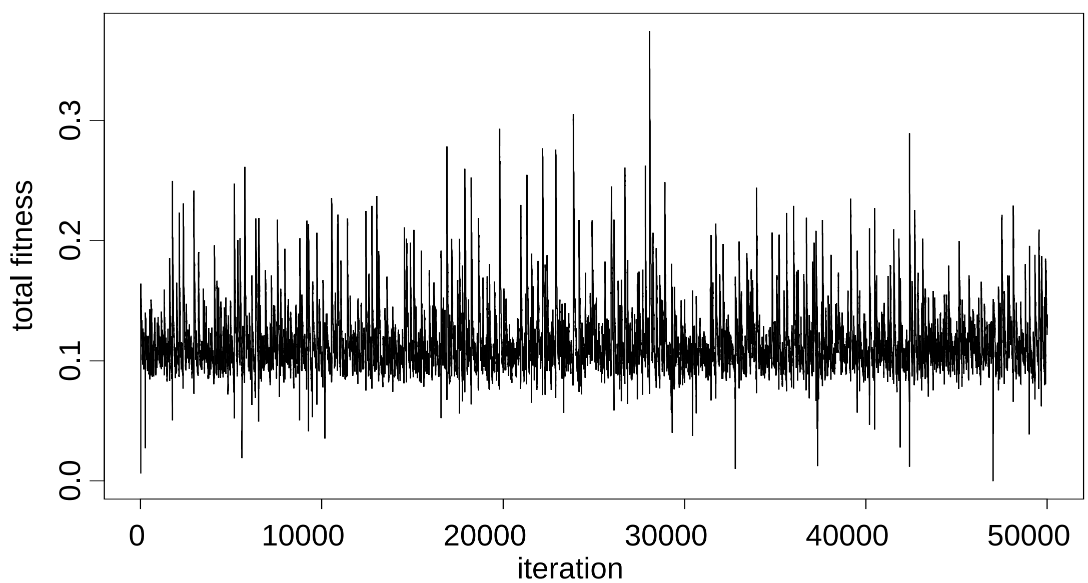

Figure 3 *Fitness of the overall system in a typical solution of the model for N = 100*

[5] this statement appears to be true for N as low as 10. Below N = 10, solutions of the model almost always end in total extinction. In the limit when N =1, the mean overall fitness of the system is of course zero.

5. The impact of regulation on the properties of the general model

5.1 How regulation is introduced into the model

In the general model, pair-wise connections between agents of the form J_{ij}, $J_{ji} > 0$ are permitted. In other words, complementary behaviour exists. The main purpose of economic regulation is to promote competition, and so the focus in this section is mainly on versions of the model which place restrictions on the number of J_{ij}, $J_{ji} > 0$ connections which are permitted.

We report here the results of versions of the general model when the interactions between pairs of agents are mainly or all of the form $J_{ij} > 0$, $J_{ji} < 0$ (or vice versa), which is the usual form of competition between firms. Allowing pair-wise connections of the form J_{ij}, $J_{ji} < 0$ leads to the properties of the model being even more different from those of the general version, with complete extinction of the entire set of agents being a not infrequent occurrence, especially when connections of the form J_{ij}, $J_{ji} > 0$ are restricted to be a low proportion of the total. Clearly, if an interaction of the form J_{ij}, $J_{ji} < 0$ holds for any length of time, the probability of both firms becoming extinct is high.

We retain as far as possible the rules from the general model described above. We retain the rule for initialising the matrix of connections between agents, J_{ij}. Given that we allow any transient behaviour arising from initial conditions to work itself out of the system before analysing the results, this should not affect the results obtained.

We also retain the rule by which firms seek to increase the fitness of their products over time. To recall, one of the elements J_{ij} is chosen at random for each agent in each period, and a new random value assigned to it. In other words, we retain the implicit assumption that firms are following maximising behaviour in the face of uncertainty.

The rule which is altered is that for replacing extinct agents. When new entrants are being created, the replacement J_{ij} are at first drawn in accordance with the normal rules of the general model. They are then checked for their sign relationship. If $J_{ij} > 0$ and $J_{ji} > 0$, then the sign of one of them switches sign with probability 1- P_{+ve} , where P_{+ve} is the probability that a positive-positive relationship will be allowed to be remain. We found it necessary to apply the same principle when $J_{ij} < 0$ and $J_{ji} < 0$. When same-sign positive relationships are excluded but same-sign negative ones are not restricted in any way, solutions of the model almost always end with total extinctions of all agents. So, if $J_{ij} < 0$ and $J_{ji} < 0$, one of them switches sign with probability 1- P_{-ve} . The general model obtains when $P_{+ve} = P_{-ve} = 1$.

This approach enables us to investigate what happens when regulation becomes more and more intense. In other words, we can examine the impact on the properties of the model when the regulator intervenes to exclude a specified percentage of same-sign connections between agents[6].

Table 2 shows for information the percentages of the connections of the forms J_{ij}, $J_{ji} > 0$, J_{ij}, $J_{ji} < 0$ and $J_{ij} > 0$, $J_{ji} < 0$ (or vice versa) over 50 solutions of the general model. The draw of the initial values for the J_{ij} will of course give 25, 25 and 50 as the respective percentages on average. But once any transient behaviour is eliminated from the system, the percentage of couplings of the form J_{ij}, $J_{ji} > 0$ rises to, on average approximately 36 per cent. This is because of the rule for introducing new agents into the model, which bases them on surviving and therefore (at the time) more successful species.

Table 2

% connections for which J_{ij}, $J_{ji} > 0$	**% connections for which** J_{ij}, $J_{ji} < 0$	**% for which** $J_{ij} > 0$, $J_{ji} < 0$ (or vice versa)
35.7	16.8	47.5

5.2 The impact of competition-enhancing regulation on the mean overall fitness of the system

We obtained 500 solutions of the model for values of $P_{+ve} = P_{-ve} = 0.9$, 0.75, 0.5, 0.25, 0.1 and 0. When the value 0.9 is chosen, for example, when a same-sign connections arises between a pair of agents, the probability of one of the signs being changed is 1 - 0.9 = 0.1. In other words, high values of P_{+ve} and P_{-ve} correspond to low levels of intervention by the regulator. We report here the results for N = 100, though the qualitative nature of the differences between these and those of the general model is not affected by the choice of N.

The effect of regulation is to lower quite unequivocally the mean overall fitness of the system. In the case of the general model, when $P_{+ve} = P_{-ve} = 1$, the fitness averaged over 500 solution of the model as a proportion of its theoretical maximum is 0.118. When the regulator intervenes as much as possible, excluding same-sign connections between agents completely, this falls to 0.0012[7]. For $P_{+ve} = P_{-ve} = 0.5$, it is 0.034.

Figure 4 plots the relationship between the mean overall fitness of the system and the degree of regulation, expressed by the choice of P_{+ve} and P_{-ve}.

---

[6] The retention of the updating rule from the general model does mean that we cannot specify absolutely exact percentages of the total number of connections which are of the forms J_{ij}, $J_{ji} > 0$ and J_{ij}, $J_{ji} < 0$, even averaged over a large number of simulations of the model. However, we are able to obtain very good approximations to such percentages. For example when we want to exclude same-sign connections completely, the number of such connections which remain in the solutions is less than 1 per cent of the total. The advantage of retaining the updating rule from the general model is that the changes to the rules of the general model are minimised.

[7] recall that if only positive same-sign connections are excluded, the solutions usually end in total extinction of all agents

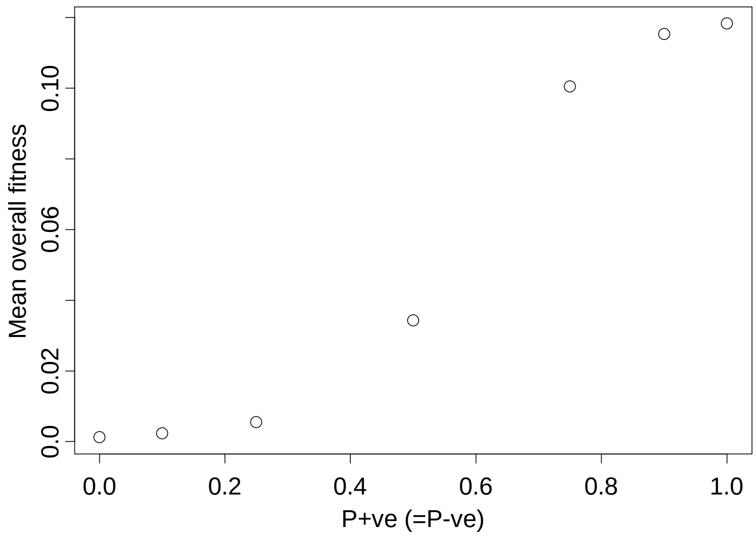

Figure 4 *Plot showing the variation in mean overall fitness with various proportions of complementary and competitive behaviour allowed to remain. When* $P_{+ve} = P_{-ve} = 1$, *the system acts as the general model, with any type of relationship between firms permitted. When* $P_{+ve} = P_{-ve} = 0$, *only competitive interactions between firms are allowed.*

As can be seen, regulation initially has little effect on the overall fitness of the system. However, beyond a certain point, overall fitness drops sharply. In practice, of course, it would be extremely difficult to identify when this was about to happen. Apart from the general reduction in fitness brought about by regulation, regulation carries a particular risk. Namely, that the degree of regulation may be such that a relatively small increase in regulation would have substantial impact on the overall fitness of the system, but the regulator would find it very hard to establish that the system was in this position.

5.3 Competition-enhancing regulation and the pattern of extinctions

The key feature of the results is that the power law relationship between the size and frequency of extinctions gradually breaks down and ceases to offer a description of the data. In particular, as the degree of regulation increases, the number of large extinctions of agents rises substantially. In other words, the property of self-organised criticality which the system possesses is destroyed.

For relatively low levels of intervention by the regulator, the relationship is not altered dramatically, but the impact of regulation rises rapidly. The relationship between the size and frequency of extinctions, along with the least squares estimate of the exponent of a power law, is plotted for typical solutions of the model with N = 100 in Figures 5 (a) - 5 (g). These show results for, respectively, $P_{+ve} = P_{-ve}$ = 1, 0.9, 0.75, 0.5, 0.25, 0.1 and 0. Figure 1(a), of course, is the general version of the model described above in section 4.

An economic interpretation of the results is that compelling firms to compete too fiercely drives profit margins down to low levels, so that firms lack the ability, and possibly the incentive, to invest for the longer-term. Of course, this is simply one possible interpretation, and this process is not contained explicitly within the model as it stands. And the concept of fitness extends beyond that of simple profitability, embracing concepts such as the X-efficiency of the firm, willingness to innovate, and so on.

In the general specification of the model, when P_{+ve} and P_{-ve} are both close to 1, the pattern of extinctions follows a power law, reflecting the self-organised critical state of the system. As noted above, extinctions play a positive role in such systems. Agents with low levels of fitness are weeded out, as it were, and replaced by ones which on average have higher levels of fitness.

However, as the self-organised state becomes broken up as P_{+ve} and P_{-ve} are reduced. There is a qualitative change in the behaviour of the system, and extinctions begin to play a more destructive role. The increasing frequency of very large extinctions means that the average lifetime of agents is shortened markedly, and individual agents are unable to build up fitness levels. The process is therefore self-reinforcing. The less agents are able to build up fitness levels, the closer their fitness will be to zero, and the greater the probability of a very large extinction taking place. Further, replacement agents also have relatively low levels of fitness, even though they are based on surviving species.

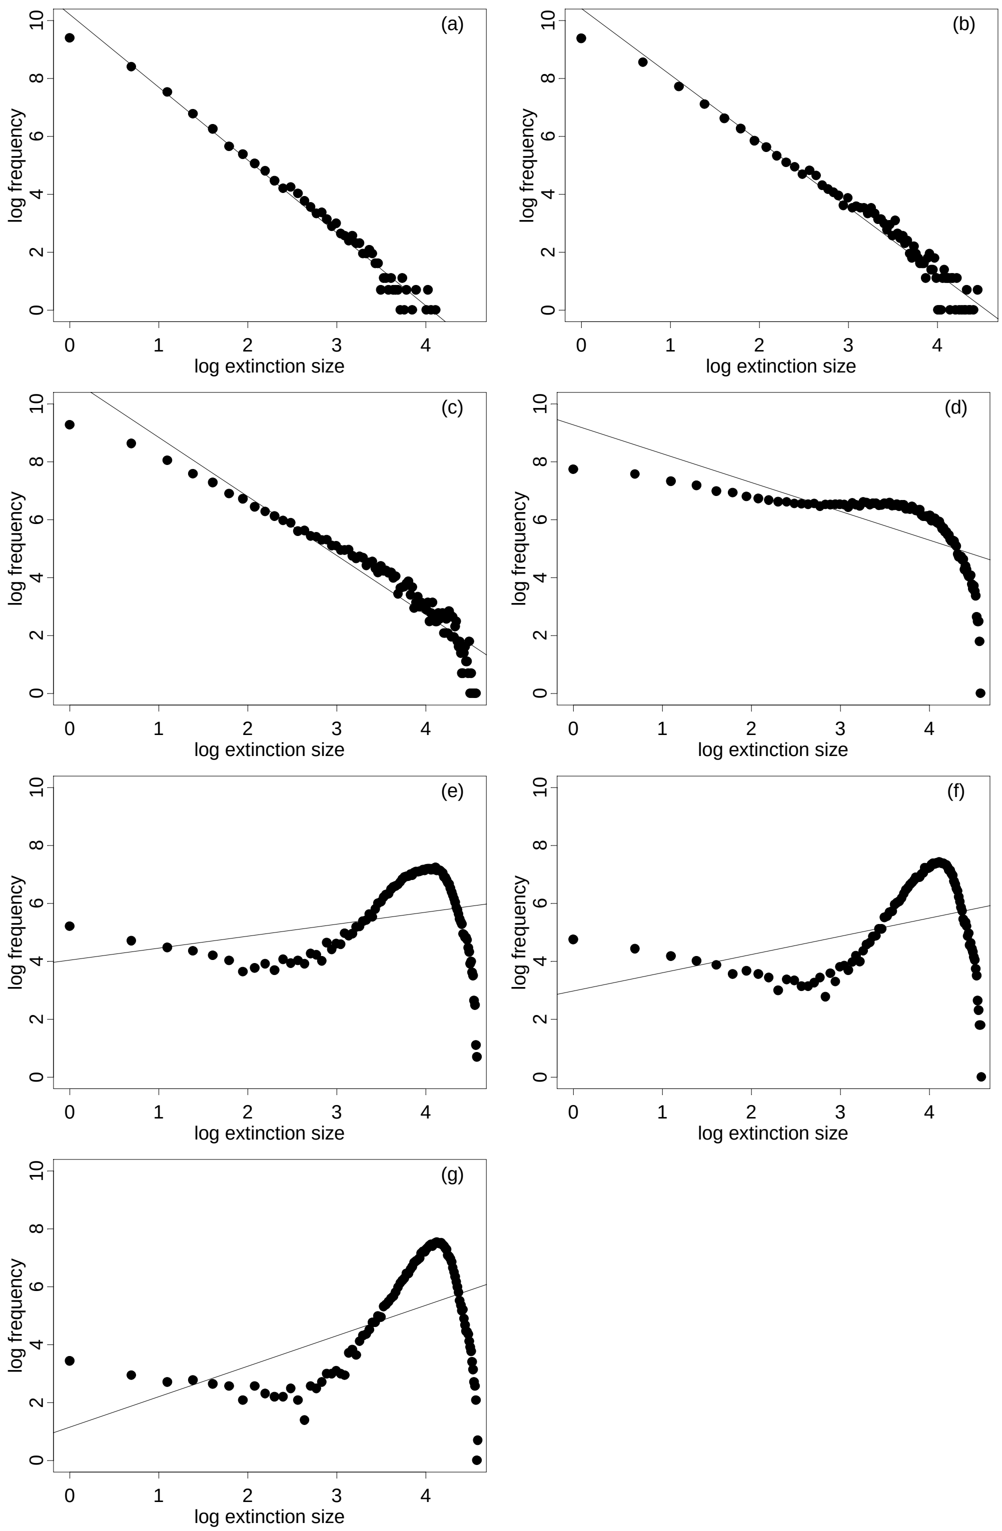

Figure 5 *Plots showing the frequency of extinction size pattern for various values of* P_{+ve} *(* $= P_{-ve}$*). (a)* $P_{+ve} = 1$ *; (b)* $P_{+ve} = 0.9$ *; (c)* $P_{+ve} = 0.75$ *; (d)* $P_{+ve} = 0.5$ *; (e)* $P_{+ve} = 0.25$ *; (f)* $P_{+ve} = 0.1$ *; (g)* $P_{+ve} = 0$

5.4 The impact of other regulatory interventions

The general model can be used to examine other ways in which regulation might affect the interactions between products and/or firms. We report briefly on two types of possible intervention, and discuss in turn how the properties of the general model are affected.

First, the regulator might intervene in the nature of new entrants into the system. Specifically, we examine the effect of varying the parameter η. The replacement rule is that at each extinction new entrants are based on the template of a surviving agent, with a random value, η, being added to the relevant J_{ij}. The smaller is η, the closer a new entrant becomes to an existing product.

The effect of varying η is very minor, and the properties of the model are essentially identical when $\eta = 0.001$, for example, and when $\eta = 1$.

It must be emphasised, however, that this does not represent regulation which forces *all* products to become more (or less) similar to each other.

The second intervention has entirely predictable results. We postulate here that the regulator intervenes to compel firms to keep products in the market event though their fitness suggests that they should become extinct. We might think, for example, of a communications company being obliged to offer services at a universal rate to remote areas, or a financial services company being required to offer services to low income households.

The way in which we introduce this is that, when an agent ought to become extinct (when its fitness falls below zero), it only does so with a fixed probability. The probability is held fixed for all the iterations of any individual solution of the model, and it is applied at each iteration. In the general model, this probability is obviously equal to one.

If the agent remains in the model, the updating rule continues to be applied to it. So in the next iteration, its overall fitness will change. If it remains below zero, the probability of extinction is again applied. But it might rise above zero, in which case it survives naturally.

Figure 6 plots the mean overall fitness of the system over 500 iterations of the general model with N = 100 for the parameter of extinction probability set equal to 1, 0.9, 0.75, 0.5, 0.25, 0.1 and 0.01.

As can be seen, the impact of regulation is to reduce the overall fitness of the system. However, in contrast to Figure 4, the effect of regulation is felt gradually over a wide range of the degree of regulation. The effects only intensify sharply at high levels of regulation.

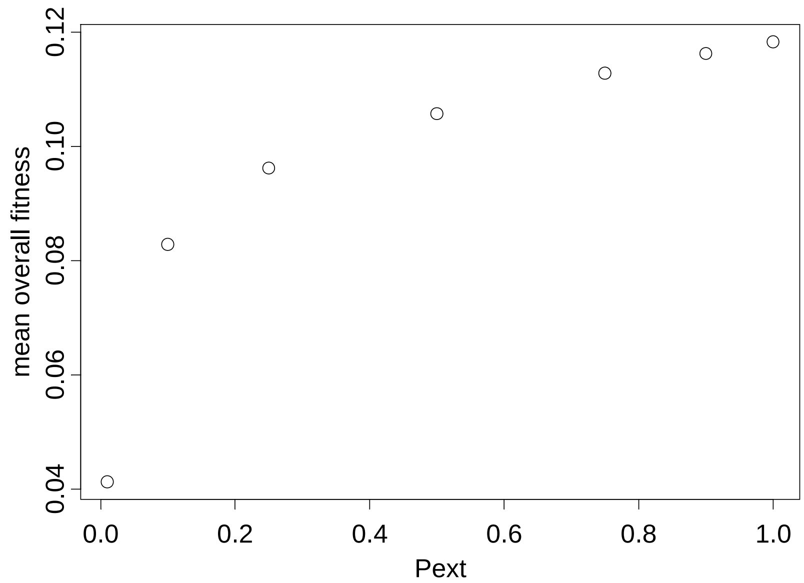

Figure 6 *Plot showing the variation in mean overall fitness as probability (Pext) that an agent which should become extinct will be allowed to do so is decreased.*

6. Conclusion

The process of evolution of biological species has recently been analysed in formal models of complex systems. Individual species interact with each other, sometimes in complementary symbiosis and sometimes in direct competition. The fitness level of each species evolves over time, and species whose fitness falls below a critical level become extinct. Variants of surviving species enter the eco-system in the niches vacated by extinct species. Such complex systems provide a good description of the empirical fossil record on species extinction. There is evidence of similarity of the empirical power law which relates the size of an extinction to its frequency in the extinction records of both biological species and firms in the economy

The outcomes of interactions between products in a market or firms in the economy on each others' fitness levels can be analysed in the same way. Firms carry out strategies, which have impacts on other firms. Many of these interactions between individual products or firms are competitive, but some are complementary.

Firms operating under uncertainty seek to maximise fitness levels, whether or themselves or of their products. Further, firms or products which fall below critical levels of fitness become extinct.

In this paper we investigate the impact of regulation in a general, complex system of this kind. In economic theory, it is regarded as desirable that firms in the same industry should compete with each other, and should not enter into complementary behaviour. This is the principle which underlies a great deal of the legal regulation of competition in industries which exists in the West. The particular focus of the paper is therefore to examine the effect of reducing the proportion of complementary interactions between agents, and replacing them with directly competitive ones.

The key finding is that as the degree of competition is increased by regulation, the properties of the system alter markedly. The frequency of large extinctions of agents rises and the overall fitness level of the system is reduced. The over-enforcement of competitive behaviour damages the system as a whole.

The results identify a particular risk associated with enforcing competitive behaviour. Initially, the enforcement of competitive relations between agents has little effect on the overall fitness of the system. Fitness falls, but only slowly. However, beyond a certain point even small further increases in regulation begin to have sharp adverse effects on the fitness of the system. In practice, it would be very difficult to identify when this point was being reached.

References

1. M.Buchanan (2000), *Ubiquity*, Weidenfeld and Nicholson, London

2. D.L.Turcotte (1999), *Rep.Prog.Phys.*, 62, 1377

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

4. A.Capocci, G.Calderelli, R.Marchetti and L.Pietronero (2001), 'Growing Dynamics of Internet Providers', cond.mat/0106084

5. S.N.Dorogovtsev and J.F.F.Mendes (2001), 'Language as an Evolving Word Web', cond.mat/0105093

6. 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

7. P.Ormerod and L.Smith (2001) 'Power law distribution of lifespans of large firms: breakdown of scaling', Volterra research paper, 2001, available at www.volterra.co.uk

8. P.Bak, C.Tang and K.Wiesenfeld (1987), *Phys.Rev.Lett.*, 59, 381

9. H.M.Jaeger, C.Liu and S.R.Nagel (1989), *Phys.Rev.Lett*, 62, 40

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

11. D.Newbery and J.Stiglitz (1982), 'The choice of techniques and the optimality of market equilibrium with rational expectations', *Journal of Political Economy*, 90

12. 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

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

14. A.Alchain (1950), 'Uncertainty, evolution and economic theory', *Journal of Political Economy*, **LIX**

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