Economics, management and complex systems

1. Introduction

There are several key features of complex systems which indicate that the economy is best analysed from the perspective of complexity science

Perhaps the single most important feature is that the macroscopically observable properties of a complex system emerge from the interactions of its constituent parts. In the context of economics, this implies that there is a need in any theoretical model for micro-foundations. In other words, a need for rules which describe the behaviour of the individual agents in the system, even when it is the macro properties of the system in which we are interested.

This view is shared by conventional economic theory, but the focus of such theory is to describe equilibrium situations. It is in essence a system of thought which is antithetical to the principles of complexity.

A further feature is a low (or even zero) ability to predict the state of the system at any given point in the future. There may very well be stable statistical distributions which describe the range of behaviours of the macroscopic factors, so that we can reasonably estimate the proportion of time which the system spends in any particular state. But we cannot predict consistently at particular points in time with any reasonable accuracy.

An important implication of this is that the understanding which individual agents have of the world is inevitably imperfect. They cannot be ascribed the cognitive powers of gathering and processing information which exist in conventional economic theory. This fits in very well with developments in economics itself in the late 20th/early 21st centuries.

From the conventional paradigm of the fully rational agent with full information and using a universal behavioural rule of maximisation, economics initially relaxed the assumption of full information, creating the concept of bounded rationality. Now, experimental and behavioural economics point to the use of limited information and rules of thumb, each one customised to particular circumstances

A final feature is that complex systems will typically exhibit multiple possible histories. By definition there can only ever be one actual history, but at any point in time the system has the potential to move in a variety of different ways.

A methodology which inherently captures these features is that of agent-based modelling with the agents connected on a network. Ideas/behaviour spread or are contained across the network, which may either be fixed or may evolve.

Section 2 of the paper considers agent behaviour and whether the assumption of economic rationality can be justified. Section 3 discusses the inherent predictability of key macro-economic variables, and section 4 presents evidence on non-Gaussian outcomes in the social sciences. Section 5 examines aspects of firm behaviour from a complex systems perspective, including an illustrative example of an agent based model with interacting agents.

2. How do agents behave?

2.1 Rationality

The understanding which individual agents have of the world is inevitably imperfect. They cannot be ascribed the cognitive powers of gathering and processing information which are implicit in conventional economic theory, even when agents are operating under bounded rather than full rationality.

The assumption of full rationality requires all agents to be able not only to gather *all* relevant information prior to making a decision, but to then be able to process it in a way which enables an agent to make the best possible decision for him or her, given a fixed set of tastes and preferences. The concept of bounded rationality relaxes only one of these key features of agent cognition, namely that of the possession of all relevant information. Agents are still presumed to take the optimal decision given the set of information available to them.

There is now a very large literature in the field of experimental/ behavioural economics. The work of the 2002 Nobel Prize winners, Vernon Smith and Daniel Kahneman[1], makes clear that in general agents do not behave according to the postulate of economic rationality. Kahneman, for example, states unequivocally in his Nobel lecture that 'humans reason poorly and act intuitively'.

Their conclusions are reinforced by, for example, the 2010 book by Bardsley et.al. *Experimental Economics: Rethinking the Rules*[2]. The six authors all have distinguished pedigrees in experimental economics. Two in particular, Loomes and Sugden, have been involved with this research programme almost from its very outset some two decades ago. The book provides a comprehensive list of almost 500 scholarly references which ranges across the entire field of experimental economics.

Many of the key results were discovered in fairly simple experiments in the early years of the whole enterprise of experimental economics. For example, consumer preferences appear in general to be non-transitive. In other words, if I prefer A to B and B to C, then transitivity requires me to prefer A to C. But this logical postulate is frequently not observed in reality. Further, agents' decisions are influenced by irrelevant alternatives. In other words, preferences expressed by agents between a set of alternative choices can be influenced by the introduction into the set of alternatives which are worse than any of the alternatives already on offer. So, for example, introducing a product which has both a higher price and worse quality than existing products can affect the decisions which people make. Preference reversal is widespread, in other words the preference ordering of a pair of alternatives depends on the process used to elicit the preference. These are just some of the examples, all of which violate key assumptions of conventional economic theory.

---

[1] See their respective lectures: D Kahneman 'Maps of bounded rationality: psychology for behavioral economics', *American Economic Review*. 93, 1449-1475, 2003; V Smith ' Constructivist and ecological rationality in economics', *American Economics* Review, 93, 465-508, 2003

[2] N Bardsley, R Cubitt, G Loomes, P Moffat, C Starmer and R Sugden , *Experimental Economics: Rethinking the Rules*, Princeton University Press, 2010

Despite the existence of this large amount of empirical evidence, the postulate of rationality is still held widely in the economics profession. There are many reasons for this, but it is useful to reflect upon just one in the current context. The recent financial crisis, for example, was simply not anticipated by the central banks, Treasuries and international institutions around the world. This theme is picked up again below in more detail. But how can this be, we might ask, if agents are presumed to form expectations about the future in a rational way? This question is dealt with easily by the true believer.

Rational expectations do not require that an agent's predictions about the future are always correct. Indeed, such predictions may turn out to be incorrect in every single period, but still be rational. The requirement is that on average over a long period of time, expectations are correct. Agents are assumed to take into account all relevant information, and to make predictions which are on average unbiased. Deviations from perfect foresight in any given period are an inherent feature of this behavioural postulate, but such deviations can only be random. If there were any systematic pattern to the deviations, the agent would be assumed to incorporate the pattern into his or her expectations. Again, on average over a long period, such expectations are correct.

It will be apparent that the theory is difficult to falsify to someone who really believes in its validity. Even the most dramatic failure to predict the future, such as the 2008 financial crisis, can be explained away as a random error. A rational expectations enthusiast can still continue to maintain the correctness of the theory by simply assuming that over some (theoretically indeterminate) period of time, on average agents' expectations prove accurate.

An assumption of the theory is that, as part of the set of information being processed, the agent is in possession of *the* correct model of the economy. Indeed, on the logic of the theory itself, if the model being used to make predictions were not correct, the forecasts would exhibit some sort of bias, some systematic error, and agents would realise that it was wrong.

It might reasonably be argued that it is difficult to subscribe to the view that agents understand the correct model of the economy given that economists themselves differ in their views as to how the economy operates. For example, in the autumn of 2008, many prominent American economists, including a number of Nobel Prize winners, vigorously opposed any form of bail-out of the financial system, arguing that it was better to let banks fail. Others, including decision makers at the Federal Reserve and Treasury, took a different view entirely.

The response of the academic mainstream has been to insist that there have been strong moves towards convergence within the profession on opinions about macroeconomic theory. By implication, anyone who takes a different view and is not part of this intellectual convergence is not really a proper economist. Olivier Blanchard, Chief Economist at the International Monetary Fund, published an MIT discussion paper in August 2008[3] on the state of modern macroeconomics. He concluded 'the state of macro is good'. The state of macro is good! Just three weeks before the financial crisis nearly brought capitalism to a halt!

2.2 Game theory

An important strand in modern economics is game theory. Uncertainty surrounds most economic decisions, and game theory appears to be an attractive way of dealing with it. In certain very limited

[3] Blanchard OJ (2008), 'The state of macro', MIT Department of Economics Working Papers, 08-17

contexts game theory and the concept of Nash equilibrium can be useful. Players, whether people, firms or governments are assumed to act rationally and seek to find a strategy that means that they themselves are as well off as they can possibly be, given how everyone else is behaving. Consider, for example, the game of noughts and crosses[4]. The outcome of this game should always be a draw since most combinations of moves will lead to this conclusion. In technical terms, the game has multiple Nash equilibria.

But beyond the confines of children's games, the concept of game theory is much less useful. Substantive assumptions about the pay-off matrix must be made before game theory can even begin to offer an account of any real world situation. The informational demands placed on agents by this are, of course, such as to render game theory useless in most practical situations.

More simply, in real life people do not appear to recognise Nash equilibrium strategies. A clear example of how an apparently simple game proves hard to play in practice is given by the *Price is Right*, a very popular television game show in America and many other countries. The rules are very straightforward and easy to remember. In other words, players have full knowledge of the rules. At all times, each player knows the state of the game. In addition, we can be sure that all those who actually get to play the game on television are devotees. They will have previously watched many previous episodes, shouting out advice or derision at the contestants from the comfort of their television rooms at home, and have had every opportunity to consider good strategic moves.

Tenorio and Cason (2002)[5], worked out analytical solutions for the Nash equilibrium strategy in *every* possible play in the game. Even more interestingly, they went on to compare these with the outcomes of what actual players did in some 300 editions of the programme. They discovered that, except where the Nash strategy is trivially obvious as it is, for example, in noughts and crosses, most of the time most of the players did not find it. Sometimes, their actual strategies were far removed from the optimal Nash decision.

The Price is Right is not a difficult game. The dimension of the problem might not seem to be large *a priori.* The rules are clear. There is no uncertainty about the situation in which a decision has to be made. Each contestant is in possession of full information about it. Yet in practice, people with every incentive to succeed, usually failed to compute the Nash equilibrium.

The disjuncture between how people ought to behave according to game theory and how they actually do behave is not a modern discovery. As Philip Mirowski makes clear in his book *Machine Dreams*[6], experiments at RAND established this almost as soon as games such as the Prisoner's Dilemma had been invented over fifty years ago. Indeed, Merrill Flood, its inventor, soon abandoned work on game theory altogether for exactly this reason.

Two examples will suffice. Flood offered RAND secretaries a choice. One of them was given the option of either receiving a fixed sum of money ($10, say), or receiving a total of $15 provided that agreement could be reached with another secretary as to how this money was to be divided between them. One Nash solution is that the two split the marginal difference. In other words, they divide the extra $5 between them so that they get $12.50 and $2.50 respectively. Obstinately, in practice most secretaries

---

[4] British English, tic-tac-toe in American English

[5] R Tenorio and T Cason, 'To spin or not to spin? Natural and laboratory experiments from "The Price is Right"', *Economic Journal*, 112, 170 – 195, 2002

[6] P Mirowski, *Machine Dreams: Economics Becomes a Cyborg Science*, Cambridge University Press, 2002

appealed not to the new idea of the Nash equilibrium but to the concept of fairness, as old as humanity itself. They divided the total amount exactly equally, $7.50 each.

The second is even more interesting. Flood carefully devised a pay-out system in the Prisoner's Dilemma in which the best option for both players was not the usual co-operative one. The Nash equilibrium was unequivocally for both players to defect. To play the game, he recruited distinguished RAND analysts John Williams and Armen Alchian, a mathematician and economist respectively. They were to play 100 repetitions of the game. They each knew about von Neumann's work, but not about the Nash equilibrium, which had only just been discovered. Both were asked to record their motivations and reactions in each round.

The Nash equilibrium strategy ought to have been played by completely rational individuals 100 times. It might of course have taken a few plays for these high-powered academics to learn the strategy. But Alchian chose co-operation rather than the Nash strategy of defection 68 times, and Williams no fewer than 78 times. Their recorded comments are fascinating in themselves, and a single aspect will have to suffice us here. Williams, the mathematician, began by expecting both players to co-operate, whereas Alchian the economist expected defection. But as the game progressed, co-operation became the dominant choice of both players.

In other words, even leading academics who had been involved in game theory research, but who were not yet aware of the newly discovered concept of the Nash equilibrium, behaved most of the time in a way contrary to the predictions of Nash's theory.

Nash was immediately told of these results, and his reaction is quoted at length by Mirowski. Many of the points are technical, but the most dramatic by far is the following: ' It is really striking how inefficient the players were in obtaining rewards. One would have thought them more rational'. In other words, his theory predicted a particular kind of behaviour. The players did not follow it and, clearly, the mistake lay with them and not the theory. Two very clever people, intimately familiar with game theory in general, had persistently chosen a non-Nash strategy. But the theory simply could not be wrong, because that is how rational people ought to behave!

3. How predictable is the economy?

Most of the results above relate to individuals. Could it be the case that institutions such as central banks, the International Monetary Fund or national Treasuries have knowledge which is superior to that possessed by the typical individual?

Certainly, the track record of forecasting macroeconomic variables such as next year's growth in GDP does not suggest any special knowledge on the part of the authorities. For example, at the start of 2008, decent growth was predicted both for Europe and the US in 2009[7]. Even as late as August, the general view was that there would still be positive growth in 2009. But in fact, the West was already in recession in August 2008!

This was not simply a one-off error in an otherwise exemplary forecasting record. The major crisis in East Asia in the late 1990s was, for example, completely unforeseen. In May of that year the International Monetary Fund (IMF) predicted a continuation of the enormous growth rates which those economies had experienced for a number of years: 7 per cent growth was projected for Thailand in 1998, 7.5 per cent for Indonesia and 8 per cent for Malaysia. By October, these had been revised down

[7] See, for example, chart 4 in the *Bank of England Quarterly Bulletin,* October 2008

to 3.5, 6 and 6.5 per cent respectively. But by December the IMF was forecasting only 3 per cent growth for Malaysia and Indonesia, and zero for Thailand. Yet the actual outturns for 1998 for these countries were spectacularly worse, with output not growing but falling by large amounts. The fall in real GDP in 1998 was -10 per cent in Thailand, and -7 and -13 in Malaysia and Indonesia respectively.

Over the past forty years in particular, a track record of forecasts and their accuracy has been built up. Economists disagree about how the economy operates, and these disagreements are reflected in, amongst other things, the specification of the relationships in macro-economic models. But, over time, no single approach has a better forecasting record than any other. Indeed, by scientific standards, the forecasting record is very poor, and a major survey of macro-economic forecasting[8] concluded that there is no real evidence which suggests that accuracy has improved over time.

As examples of the one-year ahead forecasting record for GDP growth, for the US economy recessions have not generally been forecast prior to their occurrence, and the recessions following the 1974 and 1981 peaks in the level of output were not recognised even as they took place[9] . In general, the forecasting record exhibits a certain degree of accuracy in that the average error over time is smaller than the size of the variable being predicted. But the error is still large compared to the actual data, and most of the accurate forecasts were made when economic conditions were relatively stable.

As long ago as the 1920s, Irving Fisher, the most distinguished American economist of the early decades of the 20th century, argued that the business cycle – the short term fluctuations in GDP growth - is inherently unpredictable. He believed that movements over time in the volume of output were 'a composite of numerous elementary fluctuations, both cyclical and non-cyclical' (*Journal of the American Statistical Association*, 1925) and quoted approvingly from his contemporary Moore, who wrote that 'business cycles differ widely in duration, in intensity, in the sequence of their phases and in the relative prominence of their various phenomena'.

In such circumstances, it would be virtually impossible to distinguish this type of data from data which was genuinely random in terms of its predictability. There are too many factors, and not enough data with which to identify their separate impacts. As noted above, the actual macroeconomic forecasting record is certainly compatible with this view.

Mounfield and Ormerod (2000) formalised Fisher's insight[10]. Essentially, they formed a delay matrix of time-series data on the overall rate of growth of the economy, with lags spanning the period over which any regularity of behaviour is postulated by economists to exist. They used methods of random matrix theory to analyse the correlation matrix of the delay matrix. This was done for annual data from 1871 to 1994 for 17 economies, and for post-war quarterly data for the US and the UK. The properties of the eigenstates of these correlation matrices are similar, though not identical, to those implied by random matrix theory. This suggests that the genuine information content in economic growth data is low, and that the time-series data on GDP growth is very similar to genuinely random data.

The poor forecasting record of GDP growth by economists appears to be due to inherent characteristics of the data, and cannot be improved substantially no matter what economic theory or statistical

[8] Fildes, R. and Stekler, H. (2002) The state of macroeconomic forecasting, *Journal Of Macroeconomics*, 4, 24, S. 435-468.

[9] Economic data, except in financial markets, does not appear immediately, and it can be several months before a preliminary estimate of the level of output in a given period becomes available.

[10] P.Ormerod and C.Mounfield, (2000),'Random Matrix Theory and the Failure of Macro-economic Forecasting', *Physica A,* 280, 497-504

technique is used to generate them. Over what is thought of as the time period of the business cycle in economics, in other words the period over which any regularity of behaviour of the growth of GDP might be postulated to exist, the genuine information content of correlations over time in the data is low.

The same technique can be applied to the *change* in the inflation rate, and the results are qualitatively very similar. Monetary authorities such as the Bank of England and the European Central Bank are each set a target rate of inflation which they have to try to achieve by the manipulation of short-term rates of interest. But the rate of inflation in, say, a year's time is inherently unpredictable. Indeed, we do not even know whether it will be higher or lower than it is at present, given that the changes in inflation are very similar to purely random data. So the monetary authorities are essentially attempting to target and control a random variable.

The conventional approach to the control of the economy at the aggregate level requires the ability to:

- make reasonably accurate predictions of what will happen in the future in the absence of policy changes
- have a reasonably accurate understanding of the impact of policy changes on the economy.

Neither of these is the case. There are inherent reasons why the ability to forecast with any reasonable degree of accuracy over time is severely limited, and why the ability to extract information from aggregate time-series data about the ways in which economic variables interact is also restricted.

Short-term lack of predictability is of course a key feature of complex systems.

4. Power laws and non-Gaussian outcomes

A second key characteristic is correlated behaviour amongst the individual agents of the system, which gives rise to distinctly non-Gaussian outcomes for the system as a whole.

There is a literature, stemming from the physical sciences, which tries to fit a particular kind of distribution, namely a power law. Before moving to a more general discussion of this, it is important to clarify a confusion which sometimes arises in fitting such relationships[11]. We often observe relationships reported between the size of an event and its ranking. So, for example, Zipf[12] reported a relationship between the number of times a given word is observed in a language and its rank number when all words are ranked by size. Given a variable y which orders a set of data by the size of the individual observations and the rank of each observation in this ordering, r, if the data follows a power law distribution, we will observe the relationship: $y = r^{-\beta}$. We might equally, however, examine not the size/rank relationship but the frequency distribution. In terms of the distribution of high-income earners, for example, we could perform a regression with the Zipf relationship so that a given income is proportional to the ranking of that income in the data set, or we could regress the number of people

[11] A detailed discussion of these points is available at http://www.hpl.hp.com/research/idl/papers/ranking/ranking.html

[12] GK Zipf, *Human Behavior and the Principle of Least-Effort*. Addison-Wesley, 1949

whose income is higher than this on income. But the two regressions[13] are simply different ways of looking at the same thing.

Perline[14] offers a detailed critique of the claim that power laws characterise many data sets in the social sciences. He notes that findings are often represented as though data conformed to a power law form for all ranges of the variable of interest. Perline refers to this ideal case as a *strong* inverse power law (SIPL). However, many of the examples used by Pareto and Zipf, as well as others who have followed them, have been truncated data sets, and if one looks more carefully in the lower range of values that was originally excluded, the power law behavior usually breaks down at some point. This breakdown seems to fall into two broad cases, which Perline calles here *weak* and *false* inverse power laws (WIPL and FIPL). WIPL refers to the situation where the sample data fit a distribution that has an approximate inverse power form only in some upper range of values. FIPL refers to the situation where a highly truncated sample from certain right-skew (and in particular, “lognormal-like”) distributions can convincingly mimic a power law. His paper shows that the discovery of Pareto–Zipf-type laws is closely associated with truncated data sets. Further, through detailed analysis of some reported results, he concludes that many, but not all, Pareto–Zipf examples are likely to be FIPL finite mixture distributions and that there are few genuine instances of SIPLs.

The problems of truncation in data sets are particularly acute. For example, as Perline observes ‘it is in the nature of things the low end, or very commonly, all but the upper tail, of many kinds of data is hidden because of definitional fuzziness and the difficulties associated with measurement below some threshold. At the same time, it is frequently the high end that is most important or most likely to capture our attention’.

The reasons why power laws are particularly attractive to the physical sciences, whilst important, are nevertheless a diversion to the themes of this chapter, and the interested reader is referred to the Wikipedia entry on power laws for a clear introduction to this topic.

However, there is a fundamental difference between physical systems and human and social systems. In the latter, the component parts, the agents, can act with purpose and intent, unlike the component parts of the former, the particles.

As a modelling strategy, there is a great deal to be said for taking the ‘particle’ model as the ‘null model’ In other words, to set up a model in which the agents by definition have zero cognition, with no ability to gather or process information or to learn from the past. We initially see how far this model takes us, how far it is able to account for the phenomena under investigation, before starting to make it more realistic by ascribing weak cognitive powers to agents. This concept is discussed in much more detail in Bentley and Ormerod (2010)[15]. The contrasting approach of the standard social science model is to posit the fully rational agent as the null, and then make agents slightly less smart if the model needs to be refined.

---

[13] For a more reliable way than simple least squares regression of obtaining estimates of the exponent, β, see Clauset, A., Shalizi, C. R. and Newman, M. E. J. "Power-law distributions in empirical data". *SIAM Review*, 51, 661–703, 2009

[14] R Perline. ‘Strong, weak and false inverse power laws’, *Statistical Science*, 20, 68-88, 2005

[15] RA Bentley and P Ormerod, ‘Agents, intelligence and social atoms’ in M Collard and E Slingerland, eds., *Integrating Science and the Humanities*, Oxford University Press, forthcoming 2010

A particular success with the 'zero intelligence' agent approach is reported by Farmer et.al.[16] They use data from the London Stock Exchange to test a simple model in which minimally intelligent agents place orders to trade at random. The model treats the statistical mechanics of order placement, price formation, and the accumulation of revealed supply and demand within the context of the continuous double auction and yields simple laws relating order-arrival rates to statistical properties of the market. We test the validity of these laws in explaining cross-sectional variation for 11 stocks. The model explains 96% of the variance of the gap between the best buying and selling prices (the spread) and 76% of the variance of the price diffusion rate, with only one free parameter.

There are, however, few such examples and better models are usually obtained when a small amount of cognitive ability is ascribed to the component parts. The implication is that in the social science we should not have the same fixation with trying to discover power law properties at the system level. What *is* significant, however, is that we observe very generally right-skewed (heavy tailed), distinctly non-Gaussian.

For example, Ormerod (2010)[17] examines both the duration and size of economic recessions in 17 Western economies using annual data over the period 1870 to the present. Two definitions of recession are used. First, the duration of a recession is the number of consecutive years in which real GDP growth is less than zero. The size of a recession is the cumulative percentage fall in GDP during these years. Second, a recession is defined as a period of successive years during which the level of real GDP remains below its previous peak. The size on this definition is the cumulative sum of the percentage differences between the level of GDP in each of the recession years and the level of GDP at its previous peak.

On either definition, most recessions are very short, lasting only one year in around two-thirds of the cases. Power law fits to the data give relatively poor approximations, and both the size and duration of recessions are more clearly exponential. Two approaches were used to calibrate both the exponential and the Weibull distributions to the size data. First, estimation by non-linear least squares of the appropriate functional form. Second, a grid search of the parameters which maximise the p-value at which the null hypothesis that the actual data and the theoretical distribution are the same, again using the Kolmogorov-Smirnov test.

On both definitions of a recession and using both statistical approaches, the data are best approximated by the Weibull distribution with shape parameter less than one, indicating the probability of exit from recession is reduced as duration and size are increased. This is consistent with Keynes' concept of 'animal spirits', of the sentiment of agents, becoming depressed.

There are now many examples of right-skewed distributions in the social sciences, regardless of whether they are strong, weak or false inverse power laws in the sense of Perline discussed above. All the data sets share the property that their distributions, whatever they may be, are distinctly non-Gaussian. This has been known to be a feature of the distribution of income and wealth since the time of Pareto around 1900. Decisive evidence on the right-skew distribution of firm sizes, for example, has been both available and well known in industrial economics for many years[18]. Plausible candidates in the

---

[16] JD Farmer, P Patelli and I Zovko, 'The predictive power of zero intelligence in financial markets', *Proceedings of the National Academy of Sciences*, 108, 2254-2259, 2005

[17] P Ormerod, 'Risk, recessions and the resilience of the capitalist economies', *Risk Management*, 0, 00-00, 2010

[18] For example, P.E.Hart and J.S.Prais , *Journal of the Royal Statistical Society*, 119, 150-191, 1956 J.Steindl, *Random Processes and the Growth of Firms*, London, Griffin, 1965,

economics literature to represent the empirical size distribution are the lognormal, the Pareto and the Yule. The main problem is in capturing the coverage of small firms. Recent attempts to do this, such as on the population of US firms[19], lend support to a power-law distribution linking firm sizes probability densities with the size ranking of firms. However, this may well be an as yet unexplained outcome of aggregation, because the findings seem not be robust with respect to sectoral disaggregation[20].

An innovative finding by econophysicists is that the variance of firm growth rates falls as firm size increases, although this too was anticipated in the early 1960s[21]. A further discovery is that the size-frequency relationship which describes the pattern of firm extinctions appears to be very similar to that which describes biological extinctions in the fossil record[22].

5. Modelling firm behaviour

5.1 Empirical evidence

An excellent insight of life inside a giant firm is given in Marlin Eller's book *Barbarians Led by Bill Gates*[23]. Eller was from 1982 to 1995 Microsoft's lead developer for graphics on Windows. Eller's introductory remarks are worth quoting at some length: 'There was a great disconnect between the view from the inside that my compatriots and I were experiencing down in the trenches, and the outside view… in their quest for causality [outsiders] tend to attribute any success to a Machiavellian brilliance rather than to merely good fortune. They lend the impression that the captains of industry chart strategic courses, steering their tanker carefully and gracefully through the straits. The view from the inside more closely resembles white-water rafting ."Oh my God! Huge rock dead ahead! Everyone to the left! NO, NO, the *other* left!"' . Eller goes on 'reality is rarely a simple story and is probably more like a *Dilbert* cartoon.'

The experience of Microsoft illustrates much more general points about the behaviour of firms within the complex system which is the economy. Windows now of course dominates the PC operating systems world. But its success was based far more on a series of accidents than on a far-sighted, planned strategy.

In the late 1980s, the main strategic goal of Microsoft was to link up very closely with IBM. In particular, the two companies were developing jointly a new operating system, OS/2. Windows merely limped along. Bill Gates staged a major publicity coup at the computer industry's biggest exhibition, COMDEX, in 1983. He announced that Windows 1.0 would be shipped in the spring of 1984. After immense effort, it finally appeared in November 1985. The reviews were blistering. The product size was huge relative to the capability of the personal computers which then existed. The *New York Times* observed that 'Running Windows in 512K of memory is akin to pouring molasses in the Arctic'. In Eller's blunt

[19] R.L.Axtell, *Science*, 293, 1818-1820, 2001

[20] G.Dosi, LEM Working Papers 2005/17, Sant'Anna School of Advanced Studies, Pisa, 2005

[21] L.A.N Amaral, S.V.Buldyrev, S.Havlin, H.Leschorn, P.Maass, M.A.Salinger, H.E.Stanley and M.H.R.Stanley, *J. Phys I France*, 7, 621-633, 1997

[22] W.Cook and P.Ormerod, *Physica A*, 324, 207-212, 2003 and C.DiGuilmi, M.Gallegati and P.Ormerod, *Physica A*, 334, 267-273, 2004

[23]

description: ‘the product was essentially useless’. The support team within Microsoft for Windows was cut back to a mere three people.

In contrast, great effort was being put into the relationship with IBM. In October 1988, the two companies launched OS/2 Presentation Manager, with Bill Gates proclaiming ‘[this] will be the environment for office computing in the 1990s’. Marlin Eller quotes Steve Ballmer, Gates’s number two, as saying ‘This is it, after this we’re not going to have any more Windows. It’s all OS/2’.

Windows 2 meanwhile had been launched, with little success. Only a couple of people were left within Microsoft to maintain the product. Sporadic development of the product still took place on the next version, Windows 3.0. But an article in the *National Review* summed up the view of the industry ‘Microsoft would cease development of its Windows software after the release of Windows 3.0… IBM’s OS/2 would become the main PC operating system for the 1990s’.

On May 22 1990, Windows 3.0 was made available to the public. It sold 2 million copies in the first six months.

The point is that, despite the enormous business abilities of Gates and his key players, they did not foresee that it would be Windows and not OS/2 which would fulfil this role. Windows was almost abandoned as a stand-alone product. Its support team was cut to virtually zero. And it proved a massive, overwhelming success. Success, like failure, comes in many guises.

It is the sheer complexity associated with many decisions which defies the orderly application of the rational calculation of economic theory. The number of possible permutations of outcomes is simply too great to be computed. The degree of uncertainty rarely permits the computation of the optimal, the unequivocally best strategy at any point in time.

A further practical illustration of the complex nature of the system in which firms operate is provided by Marc Levinson’s book[24] on how the humble shipping container transformed the world. Almost fifty years ago, in April 1956, a refitted oil tanker made the first ever container voyage from Newark to Houston.

From this modest start, the container has revolutionised economic geography, devastating traditional ports such as New York and London and enabling massive growth in obscure ones like Oakland and Felixstowe. Shipping costs have fallen so dramatically that the structure of world trade itself has been altered. Most trade used to be raw materials or finished products. Now it is mainly intermediate goods, with manufacturers able to source from almost anywhere, thanks to cheap transport costs. In turn, this has facilitated the massive economic growth of Asia. The container has enabled global supply chains and just in time production to become routine.

The most powerful and general insight of the book is set out in the final chapter: ‘time and again, even the most knowledgeable experts misjudged the course of events... almost nothing [the container] touched was left unchanged, and those changes were often not as predicted.’

For example, the leader of New York’s longshoremen warned in 1959 that containers would eliminate 30 per cent of his members’ jobs. Within 15 years, three quarters of them had disappeared. Even the

[24]M Levinson *The Box: How the Shipping Container Made the World Smaller and the World Economy Bigger,* Princeton University Press, 2006

inventor of the container himself, Malcolm McLean, made colossal misjudgements. At the time of the 1973/74 oil price shock, he had just ordered a new fuel-guzzling fleet, and he built a new squadron of slow but fuel efficient ships just before fuel prices fell sharply in the 1980s.

Governments in New York, San Francisco and Britain invested heavily in reconstructing traditional ports, yet the investment was obsolete almost before the last of the concrete had dried. Top American economists predicted that containerisation would be good for manufacturing in the metropolitan North Eastern states, enabling them to ship more cheaply to the South than could the landlocked Midwest. No one foresaw that the collapse in transport costs would enable entirely new competitors from elsewhere in the world to decimate the region's traditional industries.

This massive uncertainty about the future is an inherent feature of the world, which permeates both public and private sector decision making. Carroll and Hannan[25] take an ecological approach to understanding firms, and provide many interesting illustrations. A hundred years ago, for example, in the first two decades of the 20th century, over 2,000 firms attempted to make cars for the new American market. Over 99 per cent of them failed.

The endorsement of the book by Oliver Williamson, 2009 economics Nobel Laureate, brings out a further key point from their empirical examples : '...the authors adopt a demographic perspective in which variety among firms within industries becomes the object of analysis. Vitality resides in the differences--which has important ramifications for organization theory and for public policy toward business.' In other words, a key empirical feature of firms is their diversity. In the jargon of economics, the agents are heterogeneous. To non-economists, this may seem blindingly obvious, but the theoretical model of the 'representative agent', the single agent whose behaviour can proxy that of all agents in the economy, survives strongly in mainstream economics.

5.2 Failure and extinction

Failure and extinction is in fact a pervasive feature of firms, and one which is almost entirely neglected by mainstream economic theory. On average, just over 10 per cent of all firms, both in the US and Europe, become extinct in any given year. And even giant firms fail. Modern examples include Enron, WorldCom and of course Lehman's. Evidence on this is provided by the British economic historian Les Hannah[26] and by the American sociologist Neil Fligstein[27].

Fligstein's evidence is less detailed than Hannah's for our immediate purposes, though it contains much interesting material. His data set does not include evidence on whether a firm failed completely and ceased to exist as an independent entity. Rather, it focuses on whether or not a company was in the list of the largest 100 American firms at the end of each decade from 1919 to 1979. Only 33 out of the top 100 in 1919 remained in the list in 1979, and since then the attrition amongst the 1979 survivors has continued.

Fligstein notes that no fewer than 216 companies in total made it into the American Top 100 over the sixty year period. Some, such as Bethlehem Steel, WF Woolworth, Chrysler and Goodyear Tire and Rubber were in the list for the entire period. Others enjoyed their fifteen minutes of fame in a single

---

[25] GR Carroll and MT Hannan, *The Demography of Corporations and Industries*, Princeton University Press, 2000

[26] L Hannah, 'Marshall's "Trees" and the Global "Forest": were "Giant redwoods" different?', in NR Lamoreaux, DMG Raff and P Termin, eds., *Learning by Doing in Markets, Firms and Countries*, NBER, MA, 1999

[27] N Fligstein, *The Transformation of Corporate Control*, Harvard University Press, 1990

appearance, such as Atlantic Gulf and West Indies Shipping Line in 1919, Lehigh Valley Coal in 1929, Climax Molybdenum in 1939, Allied Stores in 1949, Kaiser Steel in 1959, International Utilities in 1969 and, anticipating the future, Rockwell International in 1979. International Business Machines (IBM) makes its first appearance in 1939, but otherwise computing firms such as Microsoft are absent, simply because for the most part they barely existed at the last date on Fligstein's list, 1979.

On average, over the individual decades from 1919-29 to 1969-79, 78 out of the top100 at the start of any decade were still there at the beginning of the next. But no fewer than 22 out of 100 were not. These are, or rather in most cases were, the giants of American capitalism. Operating on a massive scale, and possessed of enormous resources, almost one in every four were unable to remain in the top 100 for more than a decade.

Hannah traces the survival of the world's 100 largest industrial companies in 1912 through to 1995. The companies in the world's top 100 in 1912 represented the cream of capitalism. These were the survivors of a brutal era of competition, and had successfully survived the massive wave of mergers around the turn of the century. As Hannah points out 'They were, on the whole, firms that contemporary stock market analysts considered attractive and safe because of their consistently reliable record of generous but sustainable dividends. A population of the largest firms of ten years earlier would almost certainly show earlier exits and faster rates of decline than this population'.

Yet within 10 years, no fewer than 10 of these companies had disappeared. Over the course of the 20th century, 29 became bankrupt and in total 48 disappeared. Of the 52 survivors, only 29 remained in the world's top 100 in 1995. Hannah notes laconically; 'the tendency to over-emphasise successes, and to rationalise them *ex post* is chronically endemic amongst business historians and management consultants'. The latter group are particularly prone to the temptation of claiming to have found the unique formula for business success. Books proliferate, and occasionally sell in very large numbers, which claim to have found *the* rule, or small set of rules, which will guarantee business success. But business is far too complicated, far too difficult an activity to distil into a few simple commands, or even some of the more exotic exhortations of the business gurus.

Firms certainly act with purpose and intent, and have no intention of failing. But the complexity of the environment in which they are operating means that it is as if they were operating much closer to the zero intelligence particle model of agent behavior than to that of the fully rational agent.

5.3 An illustrative agent-based complex adaptive systems model

An example of a complex systems approach to modeling key aspects of firms which uses the 'zero intelligence' model of behavior is given by Ormerod and Rosewell[28].

The methodology used here, that of computer simulation of an agent based model, is the standard way of modeling complex systems in the social sciences. Conventional economics remains constrained by its insistence on obtaining analytical solutions to sets of equations. Analytical solutions are nice to have if you can get them, but they act as serious constraints on the types of model which can be built. Partial differential equations are, for example, routinely solved by numerical algorithms rather than brain power being wasted in an effort to obtain an analytical result. We have all moved beyond using the

[28] P Ormerod and B Rosewell, 'What can firms know?', *Proc. North American Association for Computational Social and Organization Sciences,* Pittsburgh, 2003

abacus or slide rule to perform calculations, and so we should embrace computer simulation as the best way to make progress in the social sciences.

There are two key ‘stylised facts’ at the system level which the model attempts to replicate. First, it has been known for some time that the probability of extinction is highest in the early life of a firm, but declines rapidly and is thereafter more or less invariant with respect to the lifespan of the firm[29]. Second, it has been shown recently, that the empirical relationship between the frequency and size of firm extinctions is described well by a power law[30], very similar to that observed in the paleontological record of the extinction of biological species[31].

The model contains N agents, and every agent is connected to every other. The model evolves in a series of steps. The rules of the model specify a) how the connections are updated b) how the fitness of each agent is measured c) how an agent becomes extinct and d) how extinct agents are replaced. The overall properties of the model emerge from the interactions between agents. The connections between agents can be thought of as representing the way in which the net impacts of the overall strategies of firms impact on each other. Both the strength and the signs of the connections vary. Each firm can be thought of as attempting to maximise its overall fitness level. In the model, the firm proceeds by a process of trial-and-error in altering its strategy with respect to other firms. The model is solved over a sequence of iterated steps, and at each step, for each agent one of its connections is chosen at random, and a new value is assigned to it.

Despite the fact that firms are postulated to act at random, the system wide properties which emerge from the model are very similar to those observed empirically on the distribution of firm extinctions with respect to age, and on the relationship between the frequency and size of extinctions.

After establishing this initial level of confidence in the model, Ormerod and Rosewell go on to add successively greater levels of purpose and intent to the behavior of firms, and see how far this process can go. They find that there are very considerable returns to acquiring knowledge, for even a small amount leads to a sharp increase in the mean agent age at extinction for agents with knowledge compared to those without. Indeed, they find that as both the amount of knowledge available to firms increases and as the number of firms capable of acquiring such knowledge rises, the lifespan of agents begins to approach the limiting, full information paradigm of neo-classical theory in which agents live for ever.

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

[29] For example Carrol and Hannan, op.c.it.
[30] Cook and Ormerod and di Guilimi et.al., op.cit.
[31] For example, B Drossel, ‘Biological evolution and statistical physics’, *Advances in Physics*, 50, 209-295, 2001