DTD 5 ARTICLE IN PRESS

Available online at www.sciencedirect.com

SCIENCE DIRECT

ELSEVIER

Futures xx (xxxx) 1–8

FUTURES

www.elsevier.com/locate/futures

Complexity and the limits to knowledge

Paul Ormerod*

Volterra Consulting, 5 The Old Power Station, 121 Mortlake High Street, London SW14 8SN, UK

Abstract

Economies are systems in which the macroscopically observable quantities emerge from the effects of interactions amongst the individual constituents of the system. They exhibit key features of complex systems: short-term non-predictability, emergent properties and multiple possible histories.

The conventional approach to the control of the economy at the aggregate level requires the ability to both make reasonably accurate predictions of what will happen in the future in the absence of policy changes *and* 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.

The implication is not that governments should do nothing. The actions of governments clearly do have consequences, for better or for worse. But the conventional way of thinking, which requires a world which behaves like a dependable machine, needs to be abandoned.

1. Introduction

All theories are approximations to reality. Some, such as quantum physics, appear to be extraordinarily good approximations. But they are nevertheless approximations. Limits to knowledge are therefore not relevant just to complex systems, but to all forms of modelling. In this paper, my particular focus is upon modelling the economy, an important manifestation of a complex system.

* Tel.: +44 20 8878 6333.
E-mail address: pormerod@volterra.co.uk.

0016-3287/$ - see front matter © 2004 Published by Elsevier Ltd.
doi:10.1016/j.futures.2004.11.007

Key features of complex systems are:

- short-term non-predictability
- emergent properties
- multiple possible histories.

It is these properties which give rise to the particular limits to knowledge which characterise complex systems such as the economy. Generations of policymakers have been raised in the mechanistic view of the world, with the checklist mentality: to achieve a particular set of aims, draw up a list of policies, and simply tick them off. It is a comforting environment in which to live, being seemingly dependable, predictable and controllable. The planners of the Soviet Union believed this to be the case. But their economy ultimately could not compete with the more disordered world of capitalism, not as it is portrayed in conventional economics, but as it actually exists. The intricate interactions of millions of individual agents give rise to complicated behaviour of the system as a whole.

The implications for predictability and control of a capitalist economy at the aggregate level are the key theme of this paper. The question of predictability is discussed in Section 3, and that of control in Section 4. Before this, in Section 2, I discuss how the key features of complex systems can be found at the core of conventional, free market economic theory, even though very few economists perceive it in this way.

2. Complexity and free market economic theory

Conventional economics is widely perceived as requiring 'rational' agents to process efficiently large amounts of information—using 'rational' in the particular sense of acquiring and processing efficiently large amounts of information to carry out utility maximisation.

Suppose for a moment that we inhabit such a world. I offer from this world two examples of models which incorporate features of complex systems. First of all, financial markets, where equities, bonds, currencies and other financial assets are traded.

According to economic theory, rational agents are aware of all information relating to a share that exists. So the price will only change when fresh, unanticipated information arrives. If this could be predicted, the agents would do so, so this news by definition must arrive at random. The behaviour of the price of an asset can be described as

$$\ln p(t) = \ln p(t-1) + \varepsilon(t) \quad (1)$$

where $p(t)$ is the price and $\varepsilon(t)$ is an independent random variable. (It is usually assumed to be Gaussian, although we know empirically—ever since the pioneering work of Mandelbrot in the 1960s—that this is not the case.)

This appears to be a very good empirical description of how share prices actually behave. There is no short-range autocorrelation in asset price changes. Eq. (1) is not a perfect description, because the properties of $\varepsilon(t)$ are not properly understood. The Gaussian assumption under-predicts the number of large changes. A truncated Levy distribution is better, but it is not clear that the sample moments of the distribution are time-invariant, even with very large samples. Nevertheless, the model offers a good approximation to reality.[1]

The theory cannot be rejected by the data yet, in a very important part of modern economies, namely the financial markets, we experience a lack of predictability.

How can we in any sense 'control' $p(t)$ as policy-makers? We may introduce new information, but—by assumption in this world—agents are rational, and if there is any pattern at all to this they will learn our behaviour and build it into the price. We might—by some unspecified means—be able to reduce the variance of $\varepsilon(t)$. This would reduce the amount of time which $p(t)$ spent at more extreme values, but we could not prevent it from assuming them at some (unpredictable) point.

My second example is taken from the very heart of conventional theory, so-called general equilibrium theory. The description of financial asset prices above is a theory of how one particular market behaves. General equilibrium theory, as its name might imply, is concerned with the behaviour of *all* markets in an economy.

The fundamental proposition of orthodox economic theory is that the price mechanism operates to ensure that demand will equal supply in every single market. Imbalances cannot persist, because they are smoothed away by the negative feedback generated by the price mechanism. In other words, if demand exceeds supply for a particular product or service, the price will rise and the demand fall, bringing it back into line with supply.

Perhaps the most outstanding intellectual achievement of conventional economics has been to formalise general equilibrium theory. A key aspect of this has been to establish the least restrictive set of conditions that must hold for the existence of equilibrium to be guaranteed. In other words, the conditions under which it can be proved there exists a set of prices such that demand and supply will be in balance in every single market.

It is not the purpose of this paper to enter into a critique of general equilibrium theory, and it is sufficient to note that we can state with absolute confidence that the conditions required to prove existence do not apply in reality. It has been proved that, in a multi-period world in which agents hold different beliefs about the future, each agent must have access to an infinite amount of computing power for the existence of general equilibrium to be proved. A non-technical summary of this and other theoretical problems associated with general equilibrium is given in [7].

Suppose, however, by a wild leap of the imagination, that these conditions *do* obtain. What does general equilibrium theory tell us? First of all, the system has emergent properties. Each individual agent is simply maximising utility given his or her (fixed) preferences and given the vector of prices. No agent intends all markets to clear, but this *emerges* from the reactions of agents to prices. In a one 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. It emerges from their individual actions.

[1] There is an enormous empirical literature on this, and some of the most powerful work has been done recently by physicists applying their techniques to this issue. A good summary is given by Mantegna and Stanley [3]. More recent papers are frequently posted on the world econophysics website (http://www.unifr.ch\econophysics).

[2] In a multi-period world, this is in general not true: Newbery and Stiglitz [5].

Further, it has become recognised [12] that there will usually be multiple solutions in general equilibrium. We cannot say, for given tastes and preference of agents, which of these will actually obtain. Peter Allen, in his paper in this special edition, defines a complex system as 'one that can respond in more than one way to its environment'. But we have this in general equilibrium theory, with rational, maximising agents.

So the world of rational, maximising agents can give us:

- non-predictability
- emergent properties
- multiple possible histories

3. Predictability and the business cycle

Governments of all ideological persuasions spend a great deal of time worrying about how the economy will develop in the short-term, over the next couple of years. If the anxiety levels of politicians were the only issue, few would be concerned. But our representatives do not merely contemplate the short-term future, they seek to influence it. Elaborate forecasts are prepared, not just by governments but by academic institutions and commercial companies. Advice is freely offered as to how the prospects for the economy can be improved, by an alteration to income-tax rates here, or a touch of public expenditure there. But the control that governments believe they have, in their ability both to make reasonably accurate forecasts and to understand the consequences of policy changes designed to alter the outcome, is largely illusory.

The idea that short-term fluctuations in the overall economy, the booms and recessions of what is called the 'business cycle', are intrinsically unpredictable is not new in economics. Milton Friedman argued in the early 1950s that short-term government intervention was just as likely to accentuate the fluctuations of the business cycle as it was to dampen them. In essence, he was very sceptical that governments could anticipate events with sufficient accuracy. By luck, some individual governments would get the timing of their interventions right and succeed in containing the strength of booms and slumps, but their unlucky counterparts would only succeed in intensifying the fluctuations in their economies.

The same conclusion was reached even earlier by Irving Fisher, the most distinguished American economist of the early decades of this century, using a more sophisticated argument. One of his many contributions was an article in the *Journal of the American Statistical Association* in 1925. In this, he argued that the business cycle 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', and 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.

The track record of forecasting is certainly compatible with this view. Most of the evidence on economic forecasting accuracy relates to just one-year ahead forecasts.

In terms of predicting GDP growth one-year ahead, 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[3] [13]. Further, growth has generally been overestimated during slowdowns and recessions whilst underestimates occurred during recoveries and booms [14]. For the UK, the predictions of the Treasury over the 1971–1996 period have been at least as good as those of other forecasters, but the mean absolute annual forecast error for these one-year ahead predictions was 1.45% of GDP, compared to an actual mean absolute change of 2.10% [4]. In 13 European countries over the 1971–1995 period, the average absolute error was 1.43% of GDP, compared to the average annual change of 2.91% [6].

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 [14].

In a recent paper in *Physica* A, Craig Mounfield and I have formalised Fisher's insight from the 1920s [11]. Essentially, we form 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. We use methods of random matrix theory to analyse the correlation matrix of the delay matrix. This is 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 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, it should be said, applies to the *change* in inflation. Inflation is an important government target. Whilst the rate of change of prices—inflation—itself exhibits clear dependency on its own recent behaviour, the change in this variable shows very little pattern of this kind. So, for example, the Bank of England Monetary Policy Committee considers the current rate of inflation along with predictions of what it might be, and sets short-term[4] interest rates to try to achieve the government's target figure for inflation. But the change in inflation is as inherently difficult to predict as the change in output, the rate of growth of GDP.

A clear implication of this is that an approach to policy which is based upon anticipating the immediate future state of the economy in the business cycle, and taking

[3] 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.

[4] But not long-term ones. There is considerable misapprehension about the ability of the monetary authorities to control interest rates.

decisions now to try to alter the outcome, is essentially mistaken. This does not mean that action should not be taken once the position of the economy in the business cycle becomes clear. But attempts to anticipate events are unlikely to be successful.

4. Control of the business cycle

The standard instruments of macroeconomic policy—policy designed to influence the behaviour of the economy at the aggregate level—have come to be seen since the Second World War as variables such as public expenditure, taxation and interest rates. The view that governments, or monetary authorities, can set these in order to influence the course which the economy follows is still widespread.

Separate from the question as to whether future change can be anticipated with any reasonable degree of accuracy is whether the impact of changes in these policy variables is well understood. Changes in, say, tax or public expenditure undoubtedly have an effect on the course of the economy. However, despite a substantial research campaign spanning at least three decades[5], applied economists are by no means certain of the impact of such measures. The uncertainty can even extend to the *sign* of the effect.

Church et al. [2] report properties of the six leading macro-economic models of the UK of 1992 vintage, looking at straightforward issues such as the public spending multiplier. In other words, the amount by which total output in the economy rises/falls following a sustained increase/reduction in public spending. This is a key feature in Keynesian economics: the idea that governments can stimulate output by increasing public spending. In the first year of such a hypothetical sustained increase, all models give the same answer: total output rises by more than the increase in spending. But thereafter, they diverge. By year 3, once the various feedbacks between all the variables in the model have started to work through, the following answers can be found in at least one of the six models:

- the increase in total output remains higher than the increase in public spending
- the increase in total output remains positive, but is less than the increase in public spending
- total output falls despite the increase in public spending.

A technically more sophisticated exercise that demonstrates this point on policy properties was carried out by Bray et al. [1]. The team used four of the leading UK macro-economic models of the UK economy to carry out a policy optimisation exercise 5 years into the future. The broad objectives were to achieve low inflation, to reduce unemployment, to maximise growth, to keep government borrowing within certain constraints, to stabilise the exchange rate and to meet targets on the balance of payments. The available instruments spanned the range of conventional fiscal and monetary tools, being the standard rate of income tax, government expenditure and the short-term interest rate.

---

[5] The earliest such efforts were carried out as long ago as the 1940s, but it was the growth of computing power which expanded dramatically the ability of economists to do applied work of this kind.

The differences in the model results are not merely quantitative but are different in direction. Compared to the base projection, three of the models require interest rates to be lower, albeit by widely different amounts, and one requires them to be higher. As a percentage of GDP, government expenditure is higher in one model than in the base forecast, much lower in another, virtually unchanged in yet another, and both higher and lower over the 5-year period in the final model.

A discussion on why applied economists have been unable to achieve anything like a consensus on the effects of policy changes on variables such as GDP, inflation and employment is given in [8,9]. Essentially, this relates to the lack of information in the data—the high noise to signal ratio—referred to above in the context of economic forecasting. One way of thinking about this is as follows. Once we abandon the notion that economies are in equilibrium, our ability to construct relationships between variables which are statistically stable over time becomes problematic. Imagine a simple two-equation differential equation system. Even if only a mild degree of non-linearity exists, information obtained about the relationship between the movements of the two variables in one part of the vector field will not necessarily inform us about the relationship in other parts.

We do not need to invoke the idea that the behaviour of agents may not be time-invariant in order to obtain this result. And a feature of a complex system is that the behaviour of its component agents may *not* be time-invariant, as agents alter their behaviour in the light of what others do. So it is apparent why time-series econometricians can never really make progress in understanding how the economy behaves.

5. Conclusions

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.

The implication is not that governments should do nothing. The actions of governments clearly do have consequences, for better or for worse. But the conventional way of thinking, which requires a world which behaves like a dependable machine, needs to be abandoned.

Economies *are* complex systems in which the macroscopically observable quantities emerge from the effects of complex interactions amongst the individual constituents of the system. This type of analysis is in its infancy, but it offers the potential to provide a better understanding of how economies actually behave [10].

References

[1] J. Bray, S. Hall, A. Kuleshov, J. Nixon, P. Westaway, The interfaces between policy makers, markets and modellers, Economic Journal 105 (1995) 989–1000.

[2] K.B. Church, P.R. Mitchell, P.N. Smith, K. Wallis, Comparative properties of models of the UK Economy, National Institute Economic Review 1993; 145.

[3] R.N. Mantegna, H.E. Stanley, An Introduction to Econophysics, Cambridge University Press, Cambridge, 2000.

[4] C. Mellis, R. Whittaker, The treasury forecasting record: some new results, National Institute Economic Review 164 (1998) 65–79.

[5] D. Newbery, J. Stiglitz, The choice of techniques and the optimality of market equilibrium with rational expectations, Journal of Political Economy 90 (1982) 223–246.

[6] L-E. Öller, B. Barot, Comparing the accuracy of European GDP forecasts, National Institute of Economic Research, Stockholm, Sweden, 1999.

[7] P. Ormerod, The Death of Economics, Faber and Faber, London, 1994 (Chapter 4).

[8] P. Ormerod, Problems of time-series econometrics, in: S. Daniel, P. Arestis, J. Grahl (Eds.), Essays in Honour of Bernard Corry and Maurice Peston, vol. 3, Cheltenham, Edward Elgar, 1999.

[9] P. Ormerod, Post-orthodox econometrics, in: S. Dowm, J. Hillard (Eds.), Post Keynesian Econometrics, Microeconomics and the Theory of the Firm, Edward Elgar, Cheltenham, 2002.

[10] P. Ormerod, The US business cycle: power law scaling for interacting units with complex internal structure, Physica A 314 (2002) 774–785.

[11] P. Ormerod, C. Mounfield, Random matrix theory and the failure of macro-economic forecasting, Physica A 280 (2000) 497–504.

[12] J. Silvestre, The market power foundations of macroeconomic policy, Journal of Economic Literature XXXI (1993) 105–141.

[13] H. Stekler, R. Fildes. The State of Macroeconomic Forecasting', George Washington University, Center for Economic Research Discussion Paper, 1999, 99-04.

[14] V. Zarnowitz, P. Braun, Twenty-two years of the NBER-ASA Quarterly Outlook Surveys: Aspects and Comparisons of Forecasting Performance, National Bureau of Economic Research, 1992, Working Paper 3965.