

EPOS 2006

Epistemological Perspectives on Simulation

II Edition

University of Brescia, Italy

October 5-6, 2006



Assessing agent-based models in the social sciences

Paul Ormerod and Bridget Rosewell*

1. Introduction

Agent-based models appear strange to many social scientists, and especially to economists. In general, they lack explicit analytical solutions and rely instead on simulation to discover their properties. Individual agents are usually heterogeneous, defying simple, easily generalizable descriptions. Probabilistic behaviour of such agents is normally incorporated in the models. It is not surprising that there is suspicion of this modelling approach, since it looks, indeed is, so different from the deterministic models in which most economists are trained.

We are writing this paper from our perspective as economists, although the points we make are relevant more generally to the validation of agent-based simulation models.

* Volterra Consulting, London, UK.

Both conventional economic models and agent-based evolutionary ones face a fundamental problem of validation. Because we can rarely undertake fully controlled experiments, the result of any empirical testing in social and economic science is inevitably to a degree ambiguous. No theory in economics or sociology can be justified or tested simply by reference to the data. We have only to look at the enormous efforts spent in time-series econometrics in specifying the consumption function, or at attempts to identify in cross-sectional data the elasticity of female labour supply with respect to the real wage, to realize that conventional econometric ‘testing’ has not taken us very far.

Moreover, predictive power is necessarily limited in a world where practical experiments are few and far between. Indirect tests are the most that can be hoped for. This is analogous to much of the current situation in theoretical physics. While the discovery of quantum mechanical rules was capable of experimental testing and refinement at each stage, the development of string theory has been not so testable. It has proceeded by theoretical insight, the ability to integrate issues previously seen as conundrums and some indirect experimental results (Greene, 2000).

Such approaches have relied on the ability to stand alone in both logic and plausible realism. Social science is not in the position of Einstein who, when asked what he would think if observations failed to confirm his prediction concerning the bending of light said, ‘I shall be very surprised’. But it is in the position of those who suddenly realized that string theory integrated quantum mechanics and gravity and predicted, rather than assumed, the existence of the graviton.

In any system there are three elements that need to be considered in assessing model configurations, and this seems to be general across scientific effort. The first is the rules of behaviour that the model system postulates. These rules may themselves rest on other theories and be developed in a particular context. In physics, this will refer to rules covering such elements as force, mass, energy, momentum; in biology perhaps molecules or perhaps creatures; in social science it is people, groups, firms and so on.

The second is the facts that the theory is designed to explain. In physics, very different theories are applied to macroscopic forces, visible to us in the world, and microscopic quantum fluctuations. In biology, modelling the evolution of a finch’s beak will imply different models over different time scales, than the models required for the development of a pharmaceutical. In social science, we may need different theoretical constructs for the business cycle than for the evolution of industrial structure. Of course, in every case, theoreticians are always yearning for the Theory of Everything – but a partial theory of something is better than a misleading complete construct.

The third is the precision of prediction. Clearly a theory which produces a clear and refutable prediction is preferred – but not necessarily right. In considering modelling approaches, the solutions the models proffer, their testability and relevance must also be considered.

This paper looks at each of these issues in turn. It illustrates them with reference to two particular agent-based models which have been developed by us in different contexts and to illustrate two different issues. One looks at the business cycle, and the other at the dynamic process of competition. We hope that this will

help make concrete some of the issues of model assessment in social science which are currently not properly taken into account.

2. Rules of Behavior

It is important that the behavioural rules of agents in any theoretical specification of a model of a complex system should be plausible in themselves. In other words, the rules chosen within any particular model should be capable of justification in a wider context than that of the model under immediate consideration.

We cannot at this stage hope to obtain *general* rules of agent behaviour which are valid in all contexts. The rules will be specific to explaining the task in hand. This is a familiar issue in physics, for example, where different sets of rules are used to describe the behaviour of problems at the everyday level (for example, how to build a bridge which stays up) than are used to describe behaviour of particles at the quantum level. This example illustrates the extreme difficulty of obtaining completely general rules or a Theory of Everything.

Consider, for example, the neoclassical representative agent, maximizing utility under conditions of perfect information. As a behavioural rule for agents it lacks plausibility in many contexts. There is a large literature in the other social sciences on the limits to knowledge processing by individual agents. There is a rapidly growing empirical literature within economics itself on this subject (an early and important example is Loomes (1991) and more recent ones are Fehr and Fischbacher (2002) and Tenorio and Cason (2002)). In the context of general equilibrium, there is the classic theoretical demonstration by Radner (1968) that existence of general equilibrium under uncertainty can only be demonstrated provided that each agent has access to literally an infinite amount of computing capacity.

However, in limited contexts, it might not be a bad behavioural rule to use. This is particularly the case in situations in which a) there is only a restricted set of information to process and b) the consequences of actions taken now for the future are small, for it is the introduction of the future which complicates enormously the task of the representative agent. So, for example, in a model which contains agents choosing between brands of a particular product in a supermarket, it might be quite reasonable to assume that agents act according to the postulates of standard consumer demand theory. There is, at this stage of the purchasing process, a strictly limited amount of information to process and the choice of one brand of a fast moving consumer product rather than another will have only marginal implications for the future well-being of the agent making the choice.

In other contexts, different rules might be more appropriate. In particular, other agents – firms for example – need to be considered. In each case, we are seeking to postulate rules which are consistent with what has already been established to work in more limited contexts and to achieve internal plausibility for the context under investigation. To illustrate these considerations, we turn to the motivation for the rules chosen in our two model examples.

2.1. An agent-based model of the business cycle

This model addresses the issue of the business cycle in the United States. It focuses on the rate of growth of real output both of individual firms and of its aggregate (GDP). A detailed description is in Ormerod (2002, 2004).

The business cycle is assumed to arise primarily in the corporate sector, so the agents in the model are taken to represent firms. Although it is not completely accurate to say that the sole reason for the business cycle is the activity of firms, it is a reasonable simplification to make (see, for example, Burns and Mitchell, 1946).

Firms are assumed to operate at different scales of activity. This is well known to be the case. (As it happens, without going into the details here, this is a key reason for the existence of the business cycle.)

As a simplification, the model is populated by a limited number of agents (500). In other words, it is assumed that the business cycle arises primarily through the activities of the 500 largest firms. Again, this seems a reasonable simplification to make. Decisions by General Electric are of more immediate consequence for output in the US than those of, say, a small retailer.

2.1.1 *The behavioural rules*

The model evolves on a step-by-step basis, and in each step, or period, each firm decides two things:

- its rate of growth of output for that period
- its degree of optimism or pessimism about the economic conditions in which it is operating – ‘sentiment’ for short.

The rate of growth of output each agent decides to have depends upon three things. First, its rate of growth in the previous period. Second, the general level of sentiment about the future (where this is the weighted sum of individual agents’ sentiment). Third, uncertainty which is *specific* to each individual agent, and which varies from period to period. Agents are uncertain about the general level of sentiment (they do not have perfect information), and agents are uncertain about the implications of any given level of overall sentiment for their own decisions. Agents act in a heterogeneous way in the face of such uncertainty.

The level of sentiment each agent has also depends upon three things. First, its level in the previous period. Second, the rate of growth of overall output. Third, uncertainty which again is specific to each agent, and which varies from period to period. Agents are uncertain about the implications of any given rate of growth of overall output for their own level of sentiment. The connection between the level of sentiment of an agent and the rate of growth of overall output is consistent with Keynes’ description (1936) of the ‘trade cycle’ in chapter 22 of the *General Theory*.

Indeed, the overall theoretical perspective of the model is very Keynesian. Agents do not follow complicated behavioural rules involving multi-period optimization. In this model, they are myopic and follow simple rules of thumb. Fur-

ther, uncertainty plays a key role in the behaviour of agents. Finally, agents behave in heterogeneous ways in the face of uncertainty.

The specific behavioural rules are not obtained empirically. Rather, they are set in the context of an overall theoretical approach which we believe gives a reasonable description of how firms actually behave. In economic theory, the activity of business is in principle very easy: just discover your cost and demand schedules and maximize subject to these. In practice, business is a difficult and demanding activity.

It is not necessary, in our view, to test these behavioural rules directly. Rather, they can be tested, especially against alternative rules, by investigating whether the solutions to a model incorporating such rules provides a better representation of the key features of the actual data than models which incorporate different rules.

However, in other circumstances, we might wish to be able to observe more directly the parameters of a particular behavioural rule. Our second example illustrates this.

2.2. The evolution of market structure and competition

This model addresses the dynamic evolution of market structure and competition. The specific focus is on the consequences of new entrants into a market in which there is initially a single monopoly supplier.

In the 1980s and 1990s, a wide range of industrial structures with a single, dominant firm was undermined. In part this was due to regulatory changes, such as in electricity and gas markets. In part, it arose from technological change, with the rise of the PC eroding IBM's position, for example. And in yet a further group of industries, such as telecommunications, both regulatory change and technology were important.

Competition policy in many countries remains fixed in the mind-set of a) comparing static equilibria and b) regarding perfect competition as the ideal. A model which is widely used is that of Cournot oligopoly. The fewer the number of firms in the market, the greater the price is assumed to deviate from the competitive ideal. A priori, the existence of a firm with a large market share is thought to be anti-competitive.

2.2.1 The behavioural rules

Initially, there is a single monopoly supplier. The model evolves on a step-by-step basis. A rule is specified for each of the following:

- an entry process for new firms
- the process by which firms gain/lose market share
- the process of how firms react to competition
- the process of how consumers choose between firms

The entry process of new firms is stochastic. The probability of whether new firms enter the market in any given period is a parameter of the model. And the number of firms which do enter is also a parameter. Firms operate under uncer-

tainty, and the simple rule for entry reflects this fact. In any event, the reasons why firms enter markets is understood imperfectly at best (see, for example, Carroll and Hannan, 2000).

Each firm offers the product at a particular price and with a particular level of quality, where both the price and quality of the product are in $[0, 1]$. The monopolist initially offers a price and quality of 1. The price of 0 corresponds to the lowest price at which the product can be offered, given the state of technology, and a normal profit made. Similarly, a value of 0 for the quality dimension indicates the *best* which is available.¹

New entrants offer a price and quality both chosen at random from $[0, 1]$. Firms are heterogeneous, and differ both in their ability to provide the product profitably at any given price, and in their perception of what price and quality they need to offer to gain a desired level of market share.

Each firm gains access to a fixed proportion of consumers, drawn at random from $[0, 1]$. We mean by this that a particular proportion of consumers become aware that a given firm is making an offer in the market. By definition, the monopolist has access to all consumers. Firms may differ in the proportion of consumers which they wish to target. Further, both firms and consumers operate with imperfect information. In practice, by no means all consumers are aware of the full range of firms which operate in any given market – a point illustrated in a seminar we gave to economists in 2001, when they were challenged to name, in their role as consumers, all the firms making offers in the UK domestic electricity and gas industries. None of them could. Equally, firms may wish to gain access to a large proportion of consumers, but their marketing campaign may be more or less successful than planned.

Each consumer reviews the price and quality offered by each of the firms on his or her network, that is, the firms of which any given consumer is aware. Consumers are heterogeneous in their preferences between price and quality, and each consumer is assigned a weight, w_i , in $[0, 1]$ and calculates the utility of the product offered by the j th firm by $w_i p_j + (1 - w_i) q_j$. The consumer switches to the lowest value of this on offer (remembering that low values of p and q are better than high values). However, each consumer does so with a fixed probability, drawn at the outset from $[0, 1]$. This reflects a further dimension of the heterogeneous nature of consumers. They differ in their sensitivity to price and quality, and they differ in the benefits they derive from a switch compared to the costs of switching.

Firms can react to the offers of competitors. Firms do not have complete information on the preferences of consumers, and calculate at each step of the model for all of the other firms $\omega p_j + (1 - \omega) q_j$, where ω is the average of the w_i across the consumers. They are able in principle to match immediately this (p, q) offer. However, each firm is allocated a ‘flexibility’ factor, drawn at the outset from $[0, 1]$. This gives the probability of a firm being able to match this (p, q) in any given step of the model. In practice, firms differ in their ‘X-efficiency’, and this rule reflects this fact.

The model allows both firms and consumers to react to price and quality, as in conventional economic theory. However, the model allows for:

- heterogeneity of agents
- uncertainty

- imperfect information

The rules of behaviour described here are in some ways less general than in the business cycle example. In that case, an individual firm's growth depended on its own past and growth and its expectation of the system's future growth. In this case, firms and consumers are given individual characteristics which vary – and how they vary will affect the solutions of the model.

For example, the degree of customer loyalty or the flexibility of individual firms in matching price and quality will have a marked impact on how the market structure evolves in practice. In this context, external data about these parameters can be used to refine the way in which the rules operate. This does not affect the rules of behaviour as such but would certainly result in greater precision of solution.

It might be the case that market research data were available on the willingness of consumers to consider switching product – this would inform the distribution on which the 'loyalty' parameter could be based. And of course, such willingness might well vary from industry to industry – to the 'classical' position where all consumers are always willing to switch with a probability of one.

In the case of this model, then, we cannot test the model simply from its solutions. We will also need to test the parameters.

3. The facts to be Explained

A key part of the empirical assessment of the model needs to be made with reference to the ability of the model to reproduce the key facts of the issue which is being modelled. Often these are referred to in social science as 'stylized facts' – themselves simplified version of the more complex facts of the real world. This in itself suggests that the stylized fact is a fact which may not always be true. We will return to this point in our section on the explanatory/predictive power of a model. If a fact is not always true, does it matter if your model does not predict it?

Picking the key facts that require explanation is a central issue and has a major impact on the kinds of model approach chosen. If your theory suggests that perfect competition is an ideal, then you will wish to explain deviations from its predictions. When your model makes no such assumption, the facts that you wish to explain may look rather different. Our examples illustrate the debate that can be caused by which facts are taken as key.

3.1. An agent-based model of the business cycle

Business cycle theory is one area where there is disagreement about what constitutes the key facts to be explained. The time-series properties of real GDP growth in the United States are:

- a) positive but weakly determined low order autocorrelation, with other terms in the autocorrelation function being insignificant
- b) a weak concentration of the power spectrum at frequencies associated with those of the business cycle.

In general, real business cycle models are unable to replicate such features, leading Eichenbaum (1995), for example, to state that these models suffer from ‘first-order failure’. Real business cycle enthusiasts ignore these aspects of business cycle data and concentrate instead on the so-called ‘method of moments’, attempting to reproduce the relative variances and cross-correlations of different segments of GDP. We should say in passing, however, that the methodology of validation adopted by real business cycle theorists has much in common with that of agent-based evolutionary models. Unlike the trivial curve fitting exercises of time-series econometrics, the real business cycle approach does not attempt to replicate any particular series of past data which we actually observe, but to replicate certain underlying properties of such data. Where we part company with the RBC approach is on the questions of agent homogeneity and perfect information.

3.1.1 The key facts

The key facts of output growth are identified to be the following:

- positive but weakly determined low order autocorrelation, with other terms in the autocorrelation function being insignificant
- a weak concentration of the power spectrum at frequencies associated those of the business cycle
- in general, positive correlations between the growth rates of individual agents over the course of the business cycle
- a particular statistical distribution of the cumulative size of recessions

Given the first two points above, economists have been required to consider to what extent it is meaningful to speak of a business ‘cycle’ at all. Lucas (1977) noted that output changes across broadly defined sectors of the economy tend to move together over time, and argued that with respect to the qualitative behaviour of co-movements among [sectors], business cycles are all alike. The agents in our model are not aggregated into sectors, and the correlations between their individual growth rates will be lower than that observed at the sector level, because of competition between agents within sectors. But we still expect, a priori, positive correlations. This is widely recognized to be a central feature of the business cycle.

The fourth point above is discussed in Ormerod (2004), where the cumulative size of recessions in capitalist economies is shown to follow an exponential distribution

3.2 The evolution of market structure and competition

The model of the business cycle seeks to describe the evolution of a particular time series, in terms of the components of which it is composed. In the case of this model, it seeks to describe an evolution in a number of features of a market and to position such changes in time. Since the rules of behaviour are specified in a general context, the process of the evolution under examination is also taken as a general one at the outset. Thus the model is not attempting to explain a particular

set of facts, but a set which has been generalized – ‘stylized’ – from a variety of experiences.

In the case of this model, we therefore identified a set of potential outcomes, which we wished the model to explain, as follows:

3.2.1 The key facts

We have observed a number of features in the relevant industries:

- reductions in market price, both relative to the general price level and in many cases in absolute terms
- improvements in the quality of the offer
- the original incumbent retains a large market share
- a relatively small number of competitors succeeds in establishing itself in the market
- most new entrants fail and withdraw from the market

A further point is that competition and market structure evolve dynamically in time. We rarely, if ever, observe a static equilibrium.

These are rather different ‘facts’ from those usually taken to be the key elements in describing an industry. First, industry models do not generally look at dynamics at all, preferring to explain a particular snapshot. This means that the failure of entrants through time would not be considered at all.

Second, this would also imply that considerable emphasis is put on the number of firms in an industry and their combined and several shares. Indicators such as the Herfindahl index become key facts, rather than individual firm behaviour.

Finally, if models cannot cope with individual behaviour by agents, industry studies (particularly in the context of competition policy) pay immense attention to the boundary problem – who is in and who is out, which products substitute for which others and which markets which firm operates in. With individual behaviour, these become problems rather of deciding on a distribution of tastes, loyalty and market access and do not require such hard and fast judgements.

Thus far we have considered some of the issues in choosing and testing the behavioural rules of a model, and in selecting the facts that it is designed to explain. Finally, we turn to perhaps the central issue – how to judge success.

4. Predictive power

The aim in any model building exercise is to develop a model which:

- fits into a theoretical context
- has plausible rules of agent behaviour
- is able to replicate the relevant (stylized) facts

Even if a model has all three of these properties, it does not mean that it is therefore *the* model of the problem being addressed. But it is in a stronger position than models which do not possess these properties.

Neither does it mean that it will have predictive power. Much of the criticism levelled at evolutionary and agent-based models (whether separately or together) is that they fail to provide precise predictions. By contrast, it has been argued, traditional profit maximizing agents can provide models of, for example, auction processes which can successfully be used to guide action and policy in the real world.

Let us look at the solutions of the models under discussion to review the testability of their predictions and whether this is a relevant criticism.

4.1 An agent-based model of the business cycle

We have already argued that many models of the business cycle do not provide a full enough description of the key facts – if this model is capable of explaining both the ‘usual’ facts and the ‘missing’ facts it will clearly have predictive power missing from standard models and ought therefore to be preferred.

This model is able to replicate the key stylized facts of the US business cycle summarized above.

It implies that the business cycle is generated *endogenously* and, although external shocks can be added to the model, it is not necessary to posit the existence of exogenous shocks to generate the cycle. The three key reasons for the existence of the business cycle are a) the existence of uncertainty b) the heterogeneous reactions of firms to such uncertainty and c) the fact that different firms operate at different levels of output. A specific prediction which arises from this latter point is that the amplitude of the business cycle will be lower, the less concentrated is production amongst a small number of very large firms. Since around 1960, smaller firms have become more important in the US economy and the amplitude of the cycle has indeed dampened, so the prediction is consistent with this outcome.

Despite the scientific power of the model and its clear grounding in economic theory, it has been met by a mixture of hostility and incomprehension amongst most economists who have seen it. One such group is simply unable to grasp that scientific testing extends far beyond the simple curve fitting of time-series econometrics, an exercise which, as a matter of interest, is *not* considered to be ‘modelling’ at all in other sciences. For another, smaller, group, it is as if macro-economics consisted solely of an exegetical account of the writings of Keynes, and any attempt to improve or extend them is heresy. Finally, and most generally, most economists reject the model purely because agents in it do not carry out maximizing behaviour.

It is perhaps useful to consider why the model has attracted the attention of physicists and is published in the world’s leading journal of statistical physics. The main reason is as follows. Statistical physics is concerned with the interactions between individual agents (or particles) and the behaviour of the system as a whole which emerges from such interactions. Many such systems exhibit power law, or fractal, behaviour at the aggregate level. Evidence is growing for the existence of power law behaviour within economic systems – for example, the size distribution of firms (Axtel, 2001), the extinction patterns of firms (Cook and Ormerod, 2003, de Guilmi et.al. 2004) and the distribution of growth rates amongst firms (Amaral *et al.*, 1998).

A widespread explanation for the existence of power laws is that of self-organized criticality. However, Amaral and colleagues comment on the existence of power law behaviour in social and biological systems, that it is difficult to imagine that for all these diverse systems, the parameters controlling the dynamics spontaneously self-tune to their critical values. In a study of the distribution of the growth of firms in the US, they propose an alternative mechanism based on a) the complex evolving structure of the units making up individual firms and b) an evolution of these units according to a random process. The business cycle model offers an account of the existence of power law behaviour of the system as a whole (the duration of recessions) which also arises from similar principles. In other words, it is grounded in the micro-behaviour of the agents which comprise the system, rather than arising by mere serendipity.

4.2 The evolution of market structure and competition

This model is general in two senses. First, its rules of behaviour do not describe any particular industry or product situation. Second, its rules are probabilistic – every time the model runs there is a different outcome. Multiple runs can be used to determine both an average and a range of solutions.

On average, the model shows, for a certain choice of behavioural parameters, the speed with which price and quality might adjust and the number of entrants than might succeed. It also shows the range of potential outcomes – how many times the incumbent might either disappear or indeed see off all competition. These solutions are predictions of how often these outcomes might be observed in practice.

The model is able to reproduce the key stylized facts which are observed in the dynamic evolution of competition. A further stage of validation can also be carried out, although so far we have not had access to any relevant data set. The model described contains parameters which are not calibrated against any particular set of empirical evidence. A specific industry is likely to have a particular configuration of consumers with a particular distribution of loyalty parameters for example. In order to test the model more widely, the outcomes need to be tested against specific information on a variety of different observed behaviours. The model, if broadly true, will provide predicted outcomes which then fit a relatively wide range of behaviours and in particular the model will be able to show it provides for a wider range of real world outcomes than the standard model.

A problem which economists appear to have with this model is a normative criticism. Many economists believe that they can analyse an optimum solution – the best from the point of view of welfare. A description of the outcomes of a set of rules carries no such weight. It simply is. If the sort of model of behaviour described here is accepted, then there is little basis for all the policy prescriptions beloved of competition authorities. We suspect that this will mean continued suspicion of the results.

5. Conclusion

All models are simplifications and there should be a presumption to keep any model as simple as possible, in terms both of the description of the rules which are used and the number of such rules. It is important to be able to understand *why* a model of a complex system gives particular results, and this rapidly becomes very difficult in complicated models.

The question is whether the right simplifications are chosen and this is a theoretical as well as an empirical question. We have suggested a number of criteria on both counts.

First, we should choose clear rules of behaviour for agents which are relevant to the context in which they are being applied. They should be as general as possible and should be capable of testing if their parameters might vary in particular circumstances.

Second, the facts that the model is attempting to explain should both be carefully defined and clearly set out. In some cases, these may be stylized facts – a simplified or generalized description of the real world. Ideally, if facts are to be ignored, we should explain why – though perhaps it is more realistic to suggest that critics should search for such facts!

Third, the power of such models needs to be judged by its ability to provide solutions which match the key facts under investigation. Probabilistic models may also explain outlying solutions that might occur. Application of the model parameters to a number of cases, and adjustment of the parameters to different choice situations is a good way of testing models which are formulated in a general framework. Indeed, the more general the framework which can be shown to be robust to a number of situations, the more acceptable it should be.

Notes

¹Obviously, it makes no difference whether the best quality is defined as being a value of 1 or a value of 0.

References

Amaral, L.A.N., S.V. Buldyrev, S. Havlin, M.A. Salinger and H.E. Stanley (1998), 'Power Law Scaling for a System of Interacting Units with Complex Internal Structure', *Physics Review Letters*, **80**, pp. 1385-1388.

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

Burns, A.F. and W.C. Mitchell (1946), *Measuring the Business Cycle*, NBER.

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

Cook, W. and P. Ormerod (2003), 'Power Law Distribution of the Frequency of Demises of US Firms', *Physica*, **A**, 324, pp 207-212

Eichenbaum, M. (1995), 'Some Comments on the Role of Econometrics in Economic Theory', *Economic Journal*, **105**, pp. 1609-21.

Fehr, E. and U. Fischbacher (2002), 'Why Social Preferences Matter - The Impact of Non-Selfish Motives on Competition, Co-operation and Incentives', *Economic Journal*, **112**, C1-C33

C de Guilmi, M Gallegati and P Ormerod (2004), 'Scaling invariant distribution of firms' exit in OECD countries (with C de Guilmi and M Gallegati), *Physica A*

Keynes, J.M. (1936), *The General Theory of Employment*, Macmillan

Loomes, G., C. Starmer and R. Sugden (1991), 'Observing Violations of Transitivity by Experimental Methods', *Econometrica*, **59**, pp.425-439

Lucas, R.E., 1977, 'Understanding Business Cycles', in K. Brunner and A.H. Meltzer (eds.), *Stabilisation of the Domestic and International Economy*, Carnegie-Rochester Conference Series on Public Policy 5, North Holland, pp7-29.

Ormerod, P. (2002), 'The US Business Cycle: Power Law Scaling for Interacting Units with Complex Internal Structure', *Physica*, **A**, 314, pp 774-785

Ormerod, P (2004) 'Information cascades and the distribution of economic recessions in capitalist economies', *Physica A*,

Radner, R. (1968), 'Competitive Equilibrium Under Uncertainty', *Econometrica*, **36**, pp.31-58

Tenorio, R. and T.N. Cason (2002), 'To Spin or Not to Spin? Natural and Laboratory Experiments from *The Price is Right*', *Economic Journal*, **112**, pp. 170-195.