The Keynesian micro-foundations of the business cycle: some implications of globalisation

Paul Ormerod, Volterra Consulting Ltd., London e-mail pormerod@volterra.co.uk tel: (+44) (0)208 878 6333

I am grateful to Rod Gibson in particular for the Appendix, and to Michael Campbell and participants in the conference held at Caius College, Cambridge on 9-11 September 1999 for comments on an earlier draft.

Summary

This paper sets out a theoretical model of the business cycle based upon individual agents (firms) operating on Keynesian principles. The agents are heterogenous, and follow simple non-rational rules of behaviour which take into account each other's behaviour.

The model is deliberately parsimonious, and the agents do not exhibit learning behaviour.

Nevertheless, simulations of the model show that it offers an empirical account of the properties of the US business cycle which is decidedly superior to that of Real Business Cycle models.

It is not necessary to invoke any form of exogenous shock in the model, and the cycle is purely endogenous, essentially arising from the existence of heterogenous agents operating under uncertainty.

An important feature of the model is that the firms are different in terms of size, with the size distribution being based upon American evidence. Other things being equal, the model implies that the greater the degree of concentration of size, the greater is the amplitude of the cycle. This implies that, in so far as globalisation produces an increase in concentration, it will be associated with a greater degree of variability of output over the course of the cycle.

This implication would be offset if globalisation reduced the degree of uncertainty which agents felt they faced.

Even calibrated on US post-war data, the model suggests that occasional long and/or deep recessions are an inherent feature of the economy. A simulation of the model over a period of 1,000 years suggests that a recession in which the cumulative fall in output is more than 5 per cent takes place on average approximately every 50 years.

1. Introduction:

In this paper, we develop a microeconomic model of short-term output growth based upon interacting, heterogeneous individual agents operating under uncertainty on Keynesian principles. The aggregate output growth series which emerges from their activities has time-series properties which, on a standard range of criteria, are very similar to the cyclical properties of post-war US GNP growth.

It is not necessary to invoke any form of exogenous shock in the model, and the cycle is purely endogenous, essentially arising from the existence of heterogenous agents operating under uncertainty.

Section 2 of the paper sets out the economic basis of the model, and Section 3 describes the simulation properties of the model, particularly with respect to comparisons with actual American output data. Section 4 of the paper considers the implications, with particular reference to globalisation.

2. The microeconomic model of interacting Keynesian agents

All the agents in our model are companies, which is very much in keeping with Keynes's view that the primary source of fluctuations over the course of the business cycle is the corporate sector. The model is populated by heterogeneous individual companies operating under uncertainty. The firms are of different size, and reflect the relative size distribution of the 500 largest companies in the US.

Each individual firm decides its own rate of growth of output and its own rate of change of sentiment about the future. The model evolves in discrete time steps, and in each of these steps (periods) the majority of agents update their previous decisions on output and sentiment.

A key feature of the model is that firms interact with each other by taking account of the decisions of other companies in their own decisions on the output and sentiment variables.

The key question which every firm must decide during any particular period is the rate of growth of the output which it will produce in the next period. Once this decision has been made, the firm is stuck with it. When in the process of time the company arrives in this next period, it is allowed to decide a different growth rate for the period afterwards, but not to revise the previous decision on this period's growth. Obviously, this is somewhat artificial, but it is not completely unreasonable. For in the very short-run, there are often substantial costs involved in altering previous decisions about how much to produce. Contracts have been placed with suppliers, the workforce has been alerted as to how much effort will be needed, indeed employees may have been either taken on or sacked depending upon the circumstances, the marketing programme will be committed, and so on and so forth.

In these circumstances, agents act according to a straightforward rule in order to decide how quickly output should either be expanded or contracted in the next period. They are very short-sighted, and look no further ahead than this. In other words, they are satisficers and not maximisers.

One factor which weighs in a company's decision on how much to alter the amount produced in the next period is the rate of growth at which output is actually changing during the current period. There are costs and difficulties of altering the amount which a firm is producing by large amounts, whether up or down. So a certain amount of inertia is built into the system.

But, in addition and more importantly, in deciding the rate at which its own output is to change in the next period, each firm pays great attention to the general level of sentiment, the degree of optimism or pessimism, about the future, and how this is changing. This is very much in the spirit of Keynes, who set great store on the role of expectations and the general level of confidence in determining the outcome of the economy, both in the short-run and, through decisions on investment, in the longer-term.

In terms of the formal model, we let $X_i(t)$ be the growth rate of output of agent i at time t, and XBAR(t) the overall growth of output of the population at time t, weighted according to the sizes of the agents. In other words, XBAR(t) is simply the weighted sum of the individual $X_i(t)$.

The degree of optimism or pessimism about output growth in the future held by the i'th agent at time t is given by $Y_i^s(t)$, and the (weighted) aggregate of these individual decisions by $YBAR^s(t)$.

We use the superscript 's' to stand for 'sentiment' to emphasise that it is not the conventional concept of expectations. The $Y_i^s(t)$ should be thought of as expressing sentiment (optimism/pessimism) about change in the future, rather than being based in some sense on 'optimal' forecasts of the rate of growth of output at some specific point in the future. Keynes himself was deeply sceptical about the latter approach, arguing that 'our existing knowledge does not provide a sufficient basis for a calculated mathematical expression'.

In each time period, any firm can in principle decide to change its previous rate of growth of output, $X_i(t-1)$. The growth rate of output of agent i in this period is based on a combination of the firm's output growth in the previous period, and the aggregate sentiment in the previous period about future output growth, i.e.

$$(1-\alpha)\, X_i(t-1) \;+\; \alpha\, YBAR^s(t-1) \qquad (1).$$

This represents the rate of growth in the absence of any circumstances which are particular to agent i. The interaction between individual agents takes place through the YBAR variable, so that the sentiment of other agents about the rate of growth of output in the future affects the decision on the rate of growth of output taken by agent i.

An implication of (3) is that firms feel that there are constraints which operate on their output decisions. It is the aggregate state of sentiment, which could be deduced from newspapers such as the *Financial Times* or *Wall Street Journal*, which is a determinant of any revision to the growth of output of the ith agent. In other words, companies feel that, for example, demand may be a limiting factor on them in the future, and that one way of trying to judge the likely state of demand is via aggregate sentiments about growth in the future.

The general level of business optimism or pessimism about the future is inescapably linked with uncertainty, which in turn leads to the important property of the model that the individual agents each take different decisions at any point in time.

Given the existence of uncertainty about the future and the fact that the agents are heterogenous, each agent may interpret any given level or change in overall sentiment about the future differently, in other words may draw more or less optimistic conclusions from any given value of $YBAR^s(t-1)$.

We do not require that any individual agent is consistently more or less optimistic than any other over time with respect to these interpretations, *but simply that at each point* *in time, because of uncertainty, agents differ*. In terms of the $YBAR^s$ variable, there may also be a secondary source of uncertainty. Despite the large amount of information about sentiment available in the financial media, there is no single, published measure of $YBAR^s$ at any point in time, so firms may be uncertain, and hence differ in their interpretations, about the level of $YBAR^s$ at time (t−1).

In short, *uncertainty means that each heterogenous agent operates in a different way to every other*. There is uncertainty at any point in time about the precise level of business sentiment about the future. And there is uncertainty about what any perceived level of sentiment means for decisions by any particular firm about what its rate of growth of output should be in the period immediately ahead.

This is introduced formally into the model by modifying equation (1). The growth rate of agent i is set according to

$$X_i(t) = (1 - \alpha)\, X_i(t-1) \;+\; \alpha\, \{YBAR^s(t-1) + \varepsilon_i(t)\} \qquad (2)$$

where $\varepsilon_i(t)$ is a random variable with mean zero and variance v_1.

It is important to note, even at the risk of over-emphasising the point, that in each time period companies do *not* share the same ε. The variable ε is *not* a degree of uncertainty which is common to all firms, but each firm in each period has its own ε. In other words, ε must *not* be regarded as a common, exogenous shock which all firms experience.

The sentiment of the ith agent about the rate of growth of output in the future is derived on similar principles, which can be discussed more briefly. In the absence of circumstances particular to agent i, its sentiment is a simple function of the ith agent's sentiment in the previous period, and of the aggregate rate of growth of output in the previous period.

$$(1 - \beta)\, Y_i^s(t-1) \;-\; \beta\, XBAR(t-1) \qquad (3)$$

This differs from (3) in that there is a minus sign in front of the aggregate variable, XBAR. Other things being equal, the faster that aggregate output grows, the more pessimistic become sentiments about future output growth. This follows from Keynes's own definition of the business cycle, or what he called the trade cycle in chapter 22 of the *General Theory* : 'By a cyclical movement we mean that as the system progresses in, e.g., the upward direction, the forces propelling it upwards at first gather force and have a cumulative effect on one another but gradually lose their strength until at a certain point they tend to be replaced by forces operating in the opposite direction; which in turn gather force for a time and accentuate one another, until they too, having reached their maximum development, wane and give place to their opposite'. A mathematical approximation to this description is, of course, that of a simple oscillator, and hence the negative sign on XBAR(t - 1) in (3).

Interaction between agents takes place in (3) through the XBAR term. The decisions of all other agents on the rate of growth if output influences the sentiment of the ith agent about the rate of growth in the future.

Heterogeneity of agents is again introduced, by allowing agents to differ at any point in time on their interpretations of what any given value of XBAR actually implies. The sentiment formed about the rate of growth of output of agent i in the current period is therefore given by

$$Y_i^s(t) = (1 - \beta) Y_i^s(t-1) - \beta \{XBAR(t-1) + \eta_i(t)\} \qquad (4)$$

where $\eta_i(t)\}$ is a random variable with mean zero and variance v_2.

Again, it is important to note that in each period each firm has its own value for η, and this latter variable does *not* have a value which is common across all firms.

In summary, our model comprises equations (2) and (4)

$$X_i(t) = (1 - \alpha) X_i(t-1) + \alpha \{YBAR^s(t-1) + \varepsilon_i(t)\} \qquad (2)$$

$$Y_i^s(t) = (1 - \beta) Y_i^s(t-1) - \beta \{XBAR(t-1) + \eta_i(t)\} \qquad (4)$$

The key economic content of the model is the assumption that agents are heterogenous, which given the existence of uncertainty leads to them behaving differently. This can be seen as follows. Suppose we removed this aspect of the model, and worked instead with equations (1) and (3). Setting the sum of the weights used on individual agents equal to 1, these can be re-written as a simple pair of difference equations in XBAR and $YBAR^s$:

$$XBAR(t) = (1-\alpha)XBAR(t-1) + \alpha YBAR^s(t-1) \qquad (5a)$$

$$YBAR^s(t) = -\beta XBAR(t-1) + (1-\beta)YBAR^s(t-1) \qquad (5b)$$

The dynamics of this system of equations can be analysed quite readily by forming a matrix of its parameters and calculating the eigenvalues. For most economically meaningful pairs of values of α and β (ie: $0 < \alpha,\beta < 1$), the eigenvalues are complex but with real parts which lie between zero and one. Therefore the system of equations given by 5(a) and (b) gives rise in general to damped oscillations. It is the existence of uncertainty and the consequent introduction of the terms $\varepsilon_i(t)$ and $\eta_i(t)$ into equations (2) and (4) which gives to the model a pattern of behaviour which is quite distinct from this.

Despite its simplicity, the model does capture many of the aspects of economic behaviour which are usually associated with Keynesianism and, as we shall see, its solutions have properties which are very similar to those of actual business cycle data, without having to invoke exogenous shocks of any kind.

3. Model simulations and business cycle data

In this section, we report the properties of simulations of our model with those of the actual data series on quarterly US real GNP growth over the period 1947Q2 through 1997Q3. These are discussed in more detail in Ormerod and Campbell (1998).

The simulations are made over 200 periods, to give data series of similar length as the actual US output growth. The results we present are, in general, summaries of a total number of 1000 simulations, each carried out over 200 periods.

Our interest is in cyclical fluctuations in growth, and so the model solutions are compared with actual growth net of its mean value. (The actual data are seasonally adjusted. It would be very easy to introduce a seasonal element to our model, but this would not add in any meaningful way to its economic content)

In the first instance, we discuss the ranges over which both the aggregate output growth and the output growth of individual agents typically move, and examine the typical correlations between the rates of growth of output of individual agents over time. We then move on to discuss the time-series properties of the aggregate output growth variable.

Table 1 sets out the summary statistics for actual growth and for the average of the output growth variable, XBAR, in a 1000 simulations.

Table 1. Summary statistics of real US GNP growth (net of its mean) and the average of 1,000 model simulations

	Real US GNP growth	**Simulated growth**
Minimum	−0.0317	−0.0291
1st quartile	−0.0066	−0.0074
3rd quartile	0.0063	0.0073
Maximum	0.0321	0.0290
Standard deviation	0.0107	0.0109

There are two aspects of the simulations which concern the results for individual agents. The first of these points can be dealt with briefly. The aggregate output growth variable in the simulations moves within a typical range of around −0.03 to +0.03. The range for the typical individual agent is larger, but not dramatically so, being from −0.09 to +0.09, which seems realistic.

In terms of the second point to make, the model is set up with each of the agents representing a company and taking decisions on short-term output growth, so in its present form the relative volatility of various economy wide aggregates such as consumption and investment is not available.

However, a widely accepted property of business cycles is that output changes across broadly defined sectors move together over time. In its present, basic form, our model does not lend itself to an obvious aggregation of agents into groups representing, say, the car or alcohol producing industries. Any such aggregation would be purely artificial. But with the current model we can examine the cross-correlations in output growth between each of the agents in the population, and these are summarised in Table 2.

Table 2 Summary of cross-correlations between the period-by-period output growth of each of the 500 individual agents

Minimum	−0.08
1st quartile	0.19
Median	0.24
Mean	0.24
3rd quartile	0.29
Maximum	0.51

In other words, there is, almost universally, positive and statistically highly significant correlations between the period-by-period growth of output of individual firms. The correlations are relatively small, but this seems entirely realistic. For example, although firms within the same industrial sector will be affected in similar ways by developments in the aggregate economy, much of their marketing activity is devoted to struggles with their direct competitors over market share, which can and does fluctuate. So correlation of the short-run changes in output between firms in the same sector need not be high.

Real business cycle models have been criticised strongly in the recent literature for their inability to replicate key qualitative features of cyclical movements of the actual data. Examples of such criticism are Cogley and Nason (1995), Eichenbaum (1995), Rotemberg and Woodford (1996) and Watson (1993).

There are two serious shortcomings of real business cycle models in this context. Both the autocorrelation function and the spectral properties of their simulated data are quite different from those of the actual data (though given that the power spectrum is the Fourier transform pair of the autocorrelation function, it is entirely to be expected that deficiencies in one of these aspects are reflected in the other).

The simulated data from RBC models is *qualitatively* different from the actual American data. The actual data has low order positive autocorrelation, and then negative but insignificant autocorrelation at higher order. Simulations of RBC models, as noted for example by Cogley and Nason (op.cit.) produce data which is either complete white noise or is negatively autocorrelated at almost every lag.

In terms of its spectral properties, actual data has a power spectrum which is concentrated at the frequencies which correspond to those of the business cycle, noted by Cogley and Nason, for example, to be between 2.33 and 7 years per cycle, with maximum power of the spectrum at roughly 3.2 years per cycle. In general, RBC simulated data has a flat power spectrum, indicating that business cycle components are no more important than components of any other frequency.

Both the ACF and the power spectrum of our simulated data are very similar to that of actual data.

The first two autocorrelation coefficients of the actual quarterly data over the 1947-1997 period are 0.37 and 0.25 respectively, both of which are significantly different from zero on the usual criteria, given that an approximation to the standard error on the coefficients is $1/\sqrt{n}$, where n is the sample size. The third coefficient has a point estimate of 0.01, and the coefficients at four through eight lags are negative but insignificant.

The ACF of the actual data and the average of the ACFs of the simulated data are plotted in Figure 2. The simulated data clearly replicates *qualitatively* the actual data, exhibiting positive low order autocorrelation, and negative autocorrelation at higher lags. Over the first 8 lags, the coefficients of the ACF of actual data sum to 0.27, and those of simulated data to 0.29.

Autocorrelation function of actual and simulated data

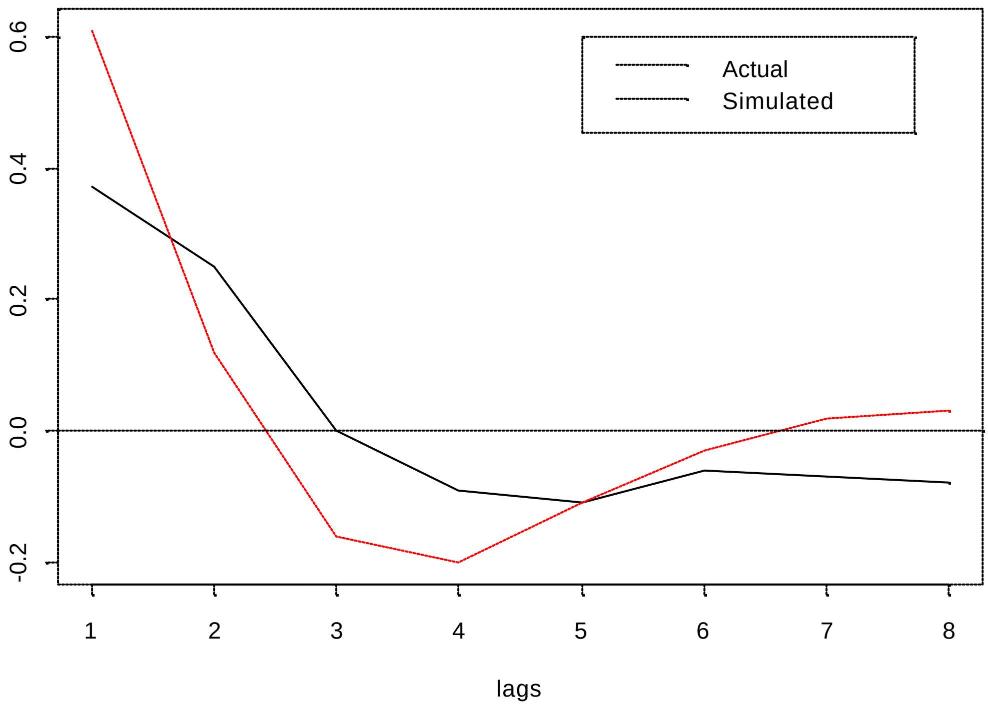

Figure 2

The spectrums of both the actual and simulated data are both concentrated at frequencies which correspond to those of the business cycle, with the simulated being somewhat more concentrated than the actual.

The degree of concentration of the data at business cycle frequencies is, however, not strong. An important reason why this is the case was advanced many years ago by Burns and Mitchell (1946), in their classic NBER work on the US business cycle, who argued that 'the sequence of changes is recurrent but not periodic; in duration cycles vary from more than one year to ten or twelve years'.

In summary, unlike real business cycle models, our simple model replicates the main qualitative features of actual US output growth. It is deliberately parsimonious in its structure, but nevertheless the calibrations of simulations of the model against actual business cycle data gives the model initial credibility.

4. Implications for globalisation

An important assumption of the model is that the individual agents have different weights. In other words, the firms are of different sizes. This merely reflects reality in the developed economies, where the very largest firms are very much larger than, say, the 500th largest firm, which itself may well be huge in terms of absolute size.

The implications can be seen when we consider a formal analytical solution of the model. Details of the derivation are given in the Appendix. (And even though such a solution can be derived, it is more convenient, given its complexity, to investigate the quantitative properties of the model by simulation)

The expression for overall output growth in the model can be reduced to a damped pendulum being driven by a stochastic forcing term:

$$\Delta^2 \overline{X}(t) + (a + b) \cdot \Delta\overline{X}(t) + 2ab \cdot \overline{X}(t) = F \qquad (6)$$

F[1] consists of terms like $\bar{e}(t) = \sum_i \frac{w_i(t)}{W(t)} \cdot e_i(t)$. By the rule of adding random variables, the variance of $\bar{e}(t)$ is going to be considerably smaller than the variance of $e_i(t)$. By the same reasoning, when considering a system with a larger number of smaller agents, the sum is over a greater number of random variables, hence leading to a smaller forcing term variance. Consequently, other things being equal, the amplitude of the cycle will be lower in an economy which has many small firms compared to one which is dominated by a small number of large ones.

To the extent that globalisation further concentrates the distribution of firm sizes, the model implies that, other things being equal, the amplitude of the business cycle will become larger.

This could be offset in principle if the level of uncertainty facing firms were reduced - in the model, this amounts to reducing the variances of ε and η. Even working with values of these calibrated on the post-war US experience, we can see that in the longer run, long and deep recessions are an integral feature of capitalist economies.

The model was simulated over 4,000 periods - 1,000 years - and the distribution of both the duration in years and the cumulative size of recessions was examined. A recession

[1] F of course has an implicit dependence on X_i through w_i. However the weights should change sufficiently slowly throughout a business cycle for this analogy with a damped pendulum to be useful.

is defined in this context as being a year in which the overall rate of growth of output is less than zero[2].

Figure 2 shows the distribution of the duration of recessions.

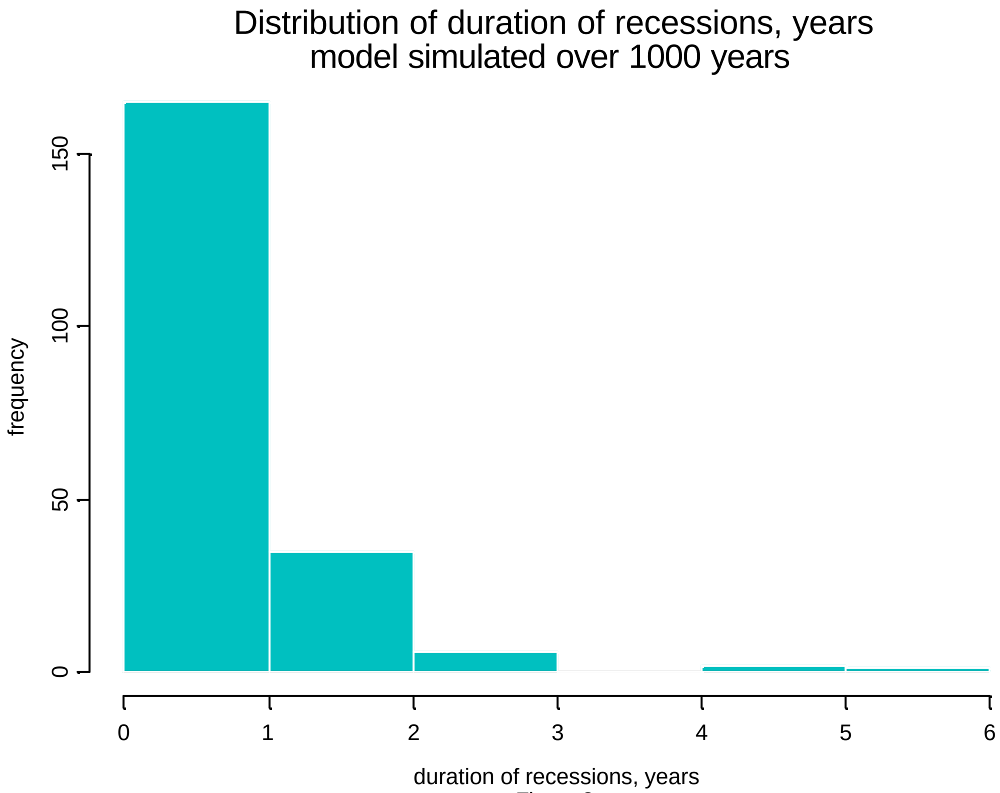

Figure 2

Recessions are rather frequent over the long-run, a feature which is certainly reflected in the experience of the capitalist economies over the past 100-150 years.

The cumulative size of the fall in output during a recession is shown in Figure 3.

[2] The model is a model of the cycle and not of long-term growth. So it is calibrated against the actual data purely in terms of the cycle ie: the actual data has its mean underlying growth rate removed. To obtain the frequency of recessions from the simulated model output, the mean annual growth rate of the US economy since the war (just under 2 per cent in real per capita terms) is added to all observations.

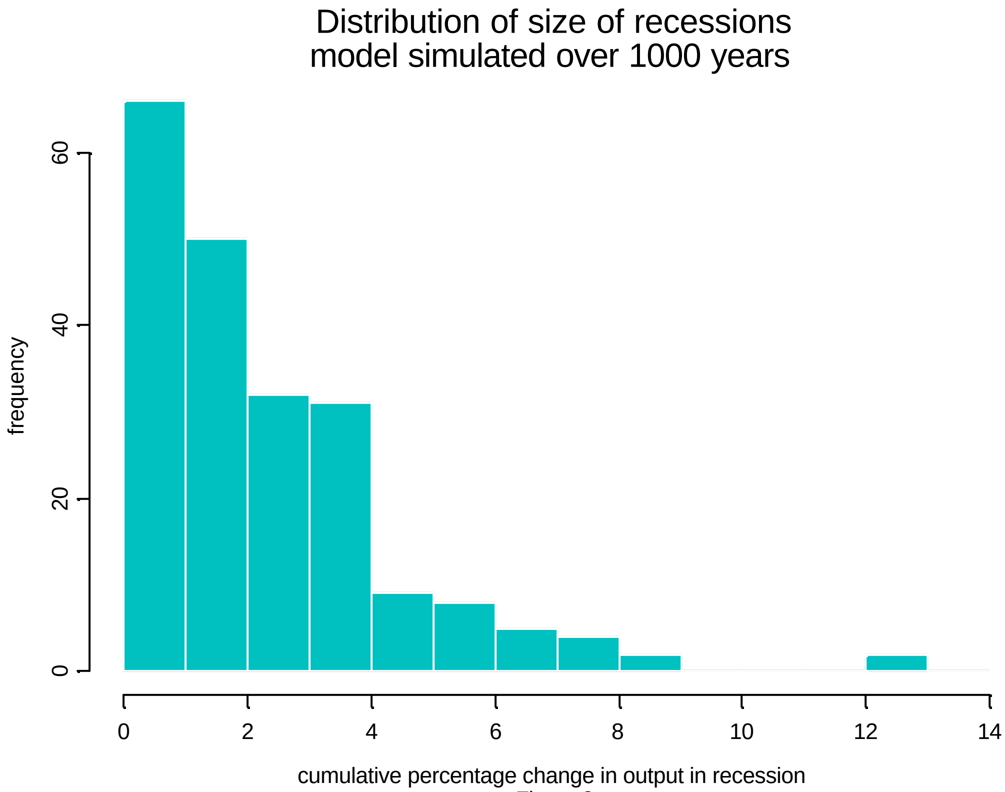

Figure 3

The data is grouped into bands of one percentage point, so the simulations of the model show just over 60 recessions in which the total fall in output was between 0 and 1 per cent, around 50 when the cumulative fall was between 1 and 2 per cent, and so on. Interestingly, there are 19 recessions in which the cumulative fall in output is over 5 per cent, approximately once every 50 years. This fits in rather well with historical experience in the capitalist economies.

To the extent that globalisation concentrates the size of firms, and to the extent that there is no offseting reduction in the overall level of uncertainty which companies face, the distributions of both the duration and size of recessions can be expected to become shifted to the right over time.

5. Conclusion

In this paper, we develop a microeconomic model of short-term output growth based upon interacting, myopic, heterogeneous individual agents operating on Keynesian principles.

The individual agents represent firms which differ in size, and are heterogeneous with respect both to decisions on output and to sentiment about future output growth. The decisions on the output growth of each agent are influenced by aggregate sentiment about the growth of output in the future, so that the decisions of others affects directly the behaviour of each agent. The sentiment on future output growth of each agent depends upon the previous value of his or her sentiment, and on the aggregate rate of growth of output. Again, through this latter term the decisions of other agents affects directly the behaviour of each individual agent.

We compare the properties of the output of 1,000 simulations of the model over 200 periods with that of the post-war US real quarterly GNP growth over the same number of periods, net of its mean. The model is calibrated so that the output growth variable has the same range of movement as that of the actual data.

A widely accepted property of business cycles is that output changes across broadly defined sectors move together over time. In its present, basic form, our model does not lend itself to an obvious aggregation of agents into groups representing different sectors. But the correlations between the growth rates of output of the individual agents in our model are generally positive and substantially different from zero, a result which is compatible with this particular stylised fact about business cycles.

The recent literature shows that the simulated data of real business cycle models does not replicate two key features of US short-term growth, namely the existence of significant positive autocorrelation at low order lags and then negative but insignificant autocorrelation at higher order lags, and the concentration of the power spectrum at frequencies which correspond to those of the business cycle.

In contrast to RBC models, our Keynesian model of interacting micro-agents does replicate qualitatively these key features of actual data.

Moreover, our model does not rely in any way on exogenous shocks. The cyclical behaviour which it exhibits in aggregate arises endogenously from the interaction of individual agents at the micro-level.

Our model is obviously a drastic simplification of reality. But the interaction of individual agents operating on Keynesian principles produces, endogenously to the model, a series for aggregate output growth whose qualitative properties are very similar to those of post-war US quarterly GNP growth.

An implication of the model is that, if globalisation further concentrates the size distribution of firms, the amplitude of the business cycle will increase. This could in principle be offset if policies could be found which reduce the overall level of uncertainty faced by companies.

The model suggests that recessions of considerable duration and size will occur quite naturally under capitalism, and that, other things being equal, these will become more frequent as the process of globalisation intensifies.

Appendix: derivation of analytical solution of the model

I am grateful to Rod Gibson for this appendix.

$$X_i(t+1) = (1-a)X_i(t) + a\{\overline{Y}(t) + e_i(t)\} \qquad \text{Eq 1}$$

$$Y_i(t+1) = (1-b)Y_i(t) - b\{\overline{X}(t) + h_i(t)\} \qquad \text{Eq 2}$$

Let the size of each agent be $w_i(t)$ and the total economy be of size $W(t) = \sum_i w_i(t)$

Then the growth rate of the whole economy is given by $\overline{X}(t) = \sum_i X_i(t)\frac{w_i(t)}{W(t)}$.

Each w_i is dynamically related to the corresponding growth rate by:-

$$w_i(t+1) = w_i(t)\cdot\{1 + X_i(t)\} \qquad \text{Eq 3}$$

Equation 1 can be partially solved by expressing it as an 'integral' equation

$$X_i(t) = (1-a)^t X_i(0) + a\cdot(1-a)^t \cdot \sum_{t=0}^{t-1}\{\overline{Y}(t) + e_i(t)\}\cdot(1-a)^{-(t+1)} \qquad \text{Eq 4}$$

The growth rate for the whole economy can now be reached by multiplying Equation 4 through by $\frac{w_i(t)}{W(t)}$ and summing over *i*.

$$\overline{X}(t) = (1-a)^t\,\overline{X}(0) + a\cdot(1-a)^t\cdot\left[\sum_{t=0}^{t-1}\overline{Y}(t)\cdot(1-a)^{-(t+1)} + \sum_{t=0}^{t-1}\sum_i \frac{w_i(t)}{W(t)}\cdot e_i(t)\cdot(1-a)^{-(t+1)}\right] \qquad \text{Eq 5}$$

An exactly similar equation exists for $\overline{Y}(t)$.

Dividing through by $(1-a)^t$ and taking differences leads to

$$\Delta\overline{X}(t) + a\cdot\overline{X}(t) = a\cdot\overline{Y}(t) + a\cdot(1-a)^{t+1}\cdot\Delta\left[\sum_i \frac{w_i(t)}{W(t)}\cdot\sum_{t=0}^{t-1} e_i(t)\cdot(1-a)^{-(t+1)}\right] \qquad \text{Eq 6}$$

Similarly for $\overline{Y}(t)$

$$\Delta\overline{Y}(t) + b\cdot\overline{Y}(t) = -b\cdot\overline{X}(t) - b\cdot(1-b)^{t+1}\cdot\Delta\left[\sum_i \frac{w_i(t)}{W(t)}\cdot\sum_{t=0}^{t-1} h_i(t)\cdot(1-b)^{-(t+1)}\right] \qquad \text{Eq 7}$$

Taking differences of Equation 6 and substituting Equation 7 and Equation 6 back in, gives:-

$$\Delta^2\overline{X}(t) + (a+b)\cdot\Delta\overline{X}(t) + 2ab\cdot\overline{X}(t) = F \qquad \text{Eq 8}$$

The left-hand side represents a damped pendulum, and the right-hand side a forcing term *F*. *F* is quite ugly and of course hides an implicate X_i dependence via w_i.

It is reasonable, at least for short stretches of time, to consider the size of agents to be relatively static. I.e. the weights are independent of time. This allows us to see the essential structure by simplifying the forcing term to:-

$$F = a\Delta\overline{e}(t) + ab\{\overline{e}(t) + \overline{h}(t)\} \qquad \text{Eq 9}$$

The bar represents the weighted sum of the components. *F* is a random variable (although serially correlated). The important point is that if the number of agents in the sum is made larger, then the variance of this forcing term will decrease, tending to zero for an infinite economy.

References:

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

Cogley,T. and Nason,J.M. (1995), 'Output dynamics in real business cycle models', *American Economic Review*, vol.85, no.3, pp.492-511

Eichenbaum,M. (1995) 'Some comments on the role of econometrics in economic theory', *Economic Journal*, vol.105, no.433, pp.1609-21

Keynes,J.M. (1936) *The general theory of employment, interest and money* Macmillan

Ormerod,P. and Campbell,M. (1998), 'A model of the US business cycle with interacting, non-rational individual agents', mimeo, Volterra Consulting, London

Rotemberg.J.J. and Woodford,M. (1996), 'Real business cycle analysis and forecastability', *American Economic Review*, vol.86, no.1, pp.71-89

Watson,M.W (1993), 'Measures of fit for calibrated models', *Journal of Political Economy*, vol.101, no.6, pp.1011-1041