Predictability and economic timeseries

Paul Ormerod (Post-Orthodox Economics and University of Manchester) and Michael Campbell (Post-Orthodox Economics)

1. Introduction

It is now nearly fifty years since the path-breaking article by Klein (1947), which used a very small econometric model of the US economy to carry out short-term forecasts. Since then, and especially during the past two to three decades, an intensive research programme has been carried out by the economic modelling community, involving the application of increasingly sophisticated time-series econometrics and of apparently more sophisticated economic theory.

Most existing macro-economic models, whatever their differences, are based upon the theoretical synthesis which became 'Keynesian' economics. Depending upon the various strengths of key linkages, these models are capable of generating a wide range of different policy simulation results, from ones close to the position of pure monetarism to ones in which fiscal policy can have a permanent effect on the economy. The criticisms which follow apply equally well, however, to the minority of models which describe themselves in alternative terms, such as 'supply-side'.

During the past thirty years or so, many short-term forecasts for GDP growth in the developed economies have been published, both by public and private sector bodies and a large literature exists on their evaluation. Three points are clear from this evidence. First of all, no single institution or methodology has a better record than any other. A particular approach may from time to time perform better, but in the longer run there are no unambiguous rankings of accuracy (see, for example, the major survey by Wallis (1989)).

The second, and related, point is that there is no substitute for genuine ex ante forecasts for evaluating any particular approach. Econometricians devote a great deal of effort to building models of the past which give good forecasts on hold-out samples, and there are many examples in the literature of such work. But the error term on data outside the period over which the equation was estimated often exhibits a non-zero mean and a variance considerably greater than that expected on the basis of in-sample performance (see, for example, the chapters by leading US and UK modellers in Ormerod (1979) discussing the essential task of choosing error terms with non-zero mean on key equations when carrying out such a forecast). Chatfield (1995), a statistician rather than an econometrician, notes the existence of this phenomenon more widely than with just economic data.

Further, the accuracy of even one-year ahead forecasts of GNP growth is, by any normal scientific criteria, very poor. Again, from time to time apparent improvements may appear to have taken place, so that McNees (1988), for example, in a survey of US forecasts claimed that 'annual forecasts for real GNP have improved over time'. But these putative improvements do not persist. The experience of the UK Treasury in the most recent ten year period is entirely typical. The average one-year ahead forecast error of GDP growth has been 1.5 percentage points, which is only very slightly lower than the actual annual average growth rate of the UK economy. The OECD (1993) note more generally for the G7 economies that over the 1987-92 period one-year ahead forecasts of growth in national output and inflation could not beat the naive rule that next year's growth/inflation will be the same as this year's. These forecasts were carried out by the governments of the G7 countries, the IMF and the OECD itself, a set of organisations embracing a range of ideological and economic views. Of course, the raw output of the models is adjusted by the forecasters, but such literature as exists on this topic suggests that the unadjusted forecasts of the models tends to be even worse (see, for example, Clements (1995) for a list of such references).

In short, genuinely ex ante one-year ahead forecasts of GDP/GNP growth in the Western economies, despite the very substantial research programme which has been dedicated to improving them, typically show errors whose first and second moments are similar to those of the actual data itself. This is not merely of academic interest, for a great deal of economic policy debate in the West is still conducted in terms of short-term prospects for the economy, and what the government could or should do to improve them. If short-term forecasting is not possible in any meaningful sense, short-term government intervention loses its validity. The key motivation of our chapter is to try to explain the phenomenon of the short-term forecasting record.

Further empirical evidence of the difficulty of constructing reasonable forecasts for, say, the growth of national output has emerged recently from the application of nonlinear estimation techniques to macro-economic data series. Potter (1995) and Tiao and Tsay (1994) investigated quarterly changes in real US GNP (national output) from 1947 through 1990, using the first difference of the natural log of GNP. Tiao and Tsay used a threshold auto-regressive model in which the data series was partitioned into four regimes determined by the patterns of growth in the previous two quarters. Potter also takes the threshold auto-regressive approach, partitioning the data into just two regimes on slightly different criteria with respect to past growth than Tiao and Tsay.

Both the above papers showed that nonlinear techniques were superior to linear auto-regressive representations of the data in terms of in-sample fit to post-war quarterly data on US GNP growth. However, the variance of the model error is barely less than the variance of the data, the former being 90 per cent of the latter in the Tiao and Tsay model and 88 per cent in Potter's best model. And this is the weakest possible test of forecasts, given that the fitted values of the model represent one-step ahead in-sample predictions.

The idea that the business cycle is intrinsically unpredictable is not new. Fisher (1925) suggested seventy years ago that the business cycle was inherently unpredictable because, in modern terminology, the dimension of the problem is too high relative to the available number of observations. He argued that movements over time in the volume of output were 'a composite of numerous elementary fluctuations, both cyclical and non-cyclical', and quoted approvingly from his contemporary Moore who wrote that '[business] cycles differ widely in duration, in intensity, in the sequence of their phases and in the relative prominence of their various phenomena'. In such circumstances, even though some deterministic structure exists, given the limited amount of data available, it would be virtually impossible to distinguish data generated by such a system from data which was genuinely random. In the 1930s, Slutsky (1937) argued that economic data was in fact generated by a series of random shocks, and that transformations of random data based upon the principle of moving summation could generate data which looked very similar to genuine time-series macro-economic data on output.

The main aim of this chapter is to examine the validity of the hypothesis that the business cycle is inherently unpredictable, in the sense that it is not possible to systematically generate forecasts of short-term growth of output which have errors whose variance is substantially less than the variance of the data.

In section 2, we describe the technique of singular spectrum analysis, which is designed to characterise the general properties of a time-series. In section 3, we apply the technique to quarterly data of post-war growth in real national output in both the United States and the United Kingdom.

2. Singular spectrum analysis

The Santa Fe Institute time series competition presented researchers with a number of unidentified time-series of data. From the outset, three distinct goals were specified, as Gershenfeld and Weigend (1993) point out in their introduction to a description of the competition. First, to forecast, or to 'accurately predict the short-term evolution of the system '. Second, to model or to 'find a description that accurately captures features of the long-term behaviour of the system'. A clear distinction was made between these two aims, and indeed was demonstrated in the results of the competition.

The third goal, described as 'system characterisation', is to determine fundamental properties of the system, such as the degree of randomness. It is this latter aim which is the focus of this chapter. It is a question which is prior to any attempt to represent the data by any particular model. Singular spectrum analysis (SSA) can be used to identify such general properties of a time-series.. A clear and accessible description is provided by Mullin (1993), and more formal accounts by Broomhead and King (1986) and Vautard and Ghil (1989).

The purpose of SSA is not, it must be stressed, to identify or build any particular model of the data. It provides information about the deterministic and stochastic parts of behaviour in a time series even when the series is short and noisy. Heuristically, the technique could be thought of as decomposing the data and identifying, should they exist, the principal periodic cycles in the data. In addition, it gives a measure of the signal to noise ratio of the data. The technique thus provides information on how far any particular model - whatever it might be - might be able to predict future movements in the data.

The task of obtaining a model, be it linear or nonlinear, to represent the data would be an entirely separate and subsequent undertaking. Even if singular spectrum analysis of a data series identifies the existence of cycles or similar structure and suggests the existence of a signal of reasonable strength, the problems of capturing the behaviour of the data in an explicit model may still be formidable.

The information provided by SSA does not tell us, for example, whether an explicit model of a data series should be univariate or multivariate since, as we describe below, it identifies the degree of consistency or structure present in a *single* series. There is some potential for confusion between the sense in which econometricians use the word structure and what we mean by it with respect to SSA. The issues surrounding SSA and the choice of a model can be illustrated by an example.

Suppose that the price of wheat was determined completely by the number of sunspots. The choice of sunspots as the illustrative explanatory variable is made because the sunspots series is clearly, in econometrician's terms, an exogenous variable. The price of wheat could, obviously, only exhibit such regularities of behaviour and signal to noise ratio as were found in the time series data on sunspots. If the sunspots series had a strong and consistent periodicity, it would be possible to build a univariate model of the wheat price that gave very accurate forecasts without knowing that the price depended on the number of sunspots. But if the sunspots were white noise, so, too, would be the price of wheat, and a univariate model would not be able to produce forecasts that were of any use.

A clever econometrician might eventually discover in this latter case that the price of wheat depended upon the number of sunspots and build what econometricians call a 'structural' model. But, by definition in this example, the complete lack of structure in the data would mean that meaningful forecasts of the price of wheat could not be made. More generally, if the price was determined by a number of factors in addition to sunspots, then the properties of the singular spectrum analysis of the wheat price would be similar to those of the underlying factors. SSA would not tell us anything about *which* factors determined the price of wheat, but would tell us about their regularity and consistency and, hence, whether meaningful forecasts could be carried out *regardless* of whether a univariate or multivariate model were used.

SSA identifies characteristics of a series by finding the dominant patterns or structure in smaller sections or windows of the data. For example, for a series X(t) where t runs from 1 to N, we choose a number m which is small by comparison with N and define vectors x_t in m dimensional space by taking the m consecutive observations centred on t, i.e.

$$x_t = (X(t-(m-1)/2), \ldots X(t-1), X(t), X(t+1), \ldots X(t+(m-1)/2).$$

These vectors are called *m-histories* and m is chosen chosen so that the length of time spanned by the observations in the embedded vector is long enough to include any structure in the data.

We will also refer to the series of vectors x_t as the *embedded series* derived from X(t) and call m the *embedding dimension*.

The embedded series x_t forms a cloud of points in Euclidean space. In general any regularity or consistency in the series X(t) will give the embedded series some discernible structure. For example if the series is predominantly periodic, the embedded series will be largely confined to two dimensions rather than extending in all of the dimensions of Euclidean space. In this case the embedded series is essentially the same as a scatter plot of the series against the lagged value of the series. On the other hand if the original series is genuinely random and entirely lacking in structure then the embedded series will be a relatively symmetric ball without any significant pattern.

The regularity or structure present in a data series is measured by taking the embedded series to be approximately an ellipsoid and calculating its principal axes or directions in Euclidean space and calculating the (root mean square) projection of the vectors x_t on to those principal directions. This amounts to finding the eigenvectors and eigenvalues of a matrix of covariance coefficients of the series X(t). The eigenvalues of this matrix, which are typically arranged in decreasing order, are usually referred to as the singular spectrum of the series. This phrase is by no means ideal, for the technique of singular spectrum analysis relies on both the eigenvalues and eigenvectors of this matrix.

The eigenvectors and eigenvalues obtained in this way from the embedded series of m-histories indicate, respectively, the characteristic patterns in the data that exist over a period of time covered by m observations and the relative significance or strength of those patterns.

The ability of SSA to identify structure is not limited to series generated by linear or only mildly nonlinear systems such as a periodic series. Provided the dimension of the embedding space is large enough to cover the time scale of the structure in the series, complex or even chaotic systems will also show some degree of regularity. As is to be expected, SSA finds less evidence of structure the more complex or chaotic the system. However, this should not be seen as a shortcoming of the technique, for it is not clear that there is a distinction between a very complex system with a large number of contributing factors and a genuinely random process either in practice or in theory. (Thinking in terms of economics, the Fisher and Slutsky hypotheses represent, respectively, these two distinct concepts, but if either is true, the implications for predicting the business cycle are the same).

Following this overview, we now give a more precise mathematical description of SSA. A more detailed account is given in Vautard and Ghil (*op.cit.*). We assume that the series X(t) where t = 1 … N has sample mean 0 and sample variance σ^2. The estimate C_k of the k'th covariance coefficient of the series X(t) for positive and negative integer values of k is defined by

$$C_k = 1/N \sum_{t = 1, N - |k|} X(t) X(t + k). \quad (1)$$

The coefficient C_0 equals the variance σ^2. Let m be the embedding dimension and denote by Γ the m x m matrix of covariance coefficients with (i, j) entry C_{i-j}, i..e.

$$\Gamma_{i, j} = C_{i-j}$$

and in accordance with the definition of the m-histories x_t, we take the indices i and j to run from $-(m - 1)/2$ to $(m - 1)/2$. We introduce a scaling factor and define the matrix A by

$$A = 1/m\, \Gamma. \quad (2)$$

The development of singular spectrum analysis is based on this matrix.

It is well known that the covariance matrix Γ given by the estimate in equation (1) is non-negative definite. Therefore since it is also symmetric, the eigenvalues λ_k of the matrix A are non-negative and can be arranged in decreasing order,

$$\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_m \geq 0.$$

We call the eigenvalues of the matrix A the *singular spectrum* of the series X.

Since the sum of the eigenvalues equals the trace of the matrix A, we have

$$\sum \lambda_k = C_0 = \sigma^2$$

where σ^2 is the variance of the series X. By normalising series to have variance 1, this scaling provides a basis for comparing the singular spectrums for different series.

Let ρ_k for k = 1 ... m be the eigenvectors of A. Like the m-histories x_t, the eigenvectors are indexed from $-(m-1)/2$ to $(m-1)/2$. The eigenvectors are orthogonal and we assume that they have norm 1. Therefore they form an orthonormal basis for the m-dimensional embedding space. The symmetry of the covariance coefficients about 0, i.e. $C_k = C_{-k}$ for $k \neq 0$, can be shown to give the eigenvectors a similar property. The eigenvectors are either odd or even functions of their index, i.e. $\rho_k(i) = -\rho_k(-i)$ in the case of the odd eigenvectors or $\rho_k(i) = \rho_k(-i)$ for even eigenvectors. It follows that for odd eigenvectors $\rho_k(0) = 0$.

We can measure the components of the original series X(t) corresponding to each of the eigenvectors of A. For each k define a derived series $X_k(t)$ by

$$X_k(t) = \sum_i X(t+i)\, \rho_k(i) \qquad (3)$$

where the sum over the index i runs from $-(m-1)/2$ to $(m-1)/2$. This could be written more simply as (x_t, ρ_k), the usual inner product of the vectors x_t and ρ_k. In other words $X_k(t)$ is a projection of $X(t)$ onto the eigenvector ρ_k.

The variance of $X_k(t)$ indicates the size of the component ρ_k in the series $X(t)$, and can be calculated easily. The matrix with columns $X_k(t)$ can be written as Ξ P where Ξ is the matrix or row vectors x_t and P is the matrix of column vectors ρ_k. The covariance matrix C is approximately equal to $1/N\ \Xi^{tr}\ \Xi$, the difference between this estimate and equation (1) is due to 'edge effects'. Therefore, since $A = 1/m\ C$, for the variance-covariance matrix of the derived series we have,

$$1/N\ (\Xi\ P)^{tr}\ (\Xi\ P) = 1/N\ P^{tr}\ \Xi^{tr}\ \Xi\ P \approx m\ P^{tr}\ A\ P = m\ \mathrm{diag}(\lambda_k)$$

where $\mathrm{diag}(\lambda_k)$ is the diagonal matrix with the eigenvectors λ_k on the main diagonal. It follows that the derived series are uncorrelated and that $X_k(t)$ which measures the component of $X(t)$ corresponding to ρ_k has variance λ_k.

The original series can be reconstructed from the projections $X_k(t)$ by inverting equation (3). Each embedded vector x_t can be written as the sum $\sum_k X_k(t)\ \rho_k$ over k from 1 to m. The series $X(t)$ is just the 0-th component of the vectors x_t, i.e.

$$X(t) = \sum_k X_k(t)\ \rho_k(0).$$

This also leads to a way of filtering the series. Taking the first r terms in this sum rather than all the terms from 1 to m, amounts to extracting those components of the series corresponding to the first r eigenvectors. This would be appropriate if the eigenvectors r + 1 to m had been identified as representing noise. Mullin (*op.cit.*) and Medio (*op.cit.*) give applications of this means of filtering.

Estimates of correlation coefficients play a central role in SSA. There is no satisfactory sampling theory for estimates of such coefficients, although approximations can be made (see, for example, Priestley (1981)). It is therefore even more difficult to try to estimate the statistical significance of the singular spectrum. In practice, it is a matter of judgment to infer from the estimated spectrum and the number of eigenvectors which exhibit regular patterns of behaviour, the degree of structure in the data. In section 3 below we also describe what is effectively a bootstrapping procedure which provide further information.

By way of illustration of the SSA technique, we examine the series generated by the differential equation investigated by Medio (1992)

$$(D/n + 1)^n x = r x (1 - x) \qquad (4)$$

where D is the differential operator d/dt and n and r are parameters. This is clearly related to the widely studied iterated quadratic map. We choose values n = 10 and r = 5, generating a series which Medio demonstrates to be mildly chaotic. Equation (4) was integrated numerically and the value recorded every 0.1 time units to give a series of 250 observations. The resulting data series M_1 is plotted in figure 1, along with a second series M_2 obtained from the first by adding noise generated from the normal distribution. The ratio of the variance of the noise to the variance of the original series is approximately 0.3.

The series M_1 is quasi-periodic with period of about 2 time units. Therefore we should choose the embedding dimension m so that it covers this interval. Figure 2 shows the singular spectrum of M_1 for embedding dimensions of 30 and 50 which correspond to time intervals of 3 and 5 units of time respectively, since we record values of the series every 0.1 time units. The eigenvalues are plotted on a $\log_{10}$-scale. Two points should be noted. First, the values range over a full three orders of magnitude. Second, the singular spectrums fall away rapidly, suggesting a considerable degree of structure since the embedded series is largely restricted to the directions corresponding to the dominant eigenvalues. This is also the pattern in embedding dimension as low as 10.

Figure 3 compares the singular spectrum of M_2 with that of M_1 for embedding dimension m = 30. The first part of the spectrums are virtually identical but the eigenvalues associated with M_2 soon reach a floor that is much larger than the smallest eigenvalues of M_1. This floor corresponds to the noisy component of M_2. The embedded series derived from M_2 is still dominated by the structure in the series M_1 but it also extends in every other direction as a result of the added noise. The qualitative nature of this result is not affected by the choice of m.

Although this is an artificial example, it illustrates that even in the presence of considerable noise, SSA can identify underlying structure. It is important to note that the results of the SSA do not in any way point us to the choice of any particular model as a way of representing the data. They simply show that a component of sufficient structure exists within the data to enable a potentially reasonable model to be built. Of course, without knowledge of (4) and being confronted with data for M_2, the practical search for such a model may be a formidable undertaking.

3. Applications of SSA to US and UK quarterly GDP data

For both the US and the UK, we examined the quarterly growth in real national output, on a seasonally adjusted basis over the post-war period. Growth was defined in the conventional way in applied economics, namely as the first difference of the log of real national output. For the US, the sample period was from 1947Q2 through 1990Q4, the same as that used in the Potter and Tiao and Tsay articles quoted above. For the UK, the sample period was 1955Q2 through 1995Q3. The two series are plotted in Figures 4a and 4b respectively. (Gross National Product is the data series for the US and Gross Domestic Product for the UK, in accordance with the usual conventions of applied economics. There is a small difference between these two definitions of national output which is more important for the US than it is for European economies, and which need not concern us here.)

Visual inspection of the two data series does not immediately suggest any obvious periodicity. Mullin (*op.cit.*) suggests that information on the choice of m, the embedding dimension of the data, can potentially be provided by the autocorrelation function (ACF). But in these examples, as in a range of other series we have investigated, it is not terribly helpful. The use of formal criteria such as the AIC to choose the embedding dimension leads to values of m which are very low - three, for example, in the case of the US data. Lags 1 and 2 are significantly different from zero at the conventional 5 per cent level with the US data, but they are the only ones. The values at lags 5 and 14 are larger than the others in absolute value, but are not significant. For the UK, no single value of the ACF is significantly different from zero, the largest absolute values being at lags 6, 8, and 12. As a much better guide, it has become the convention in economics to regard the business cycle as lasting anything between two or three to seven or eight years. Accordingly, by choosing m = 41, in other words spanning a window of ten years of data, in the first instance, we make adequate allowance for any such potential structure.

Figure 5 plots the resulting singular spectrums of US and UK data. The results are drastically different for both series than the examples given in section 2. In neither case does the spectrum show evidence of much structure. The largest eigenvalue is relatively small and the eigenvalues do not fall away rapidly. The range of the eigenvalues is only one order of magnitude. There may be somewhat more structure to the US data since the larger eigenvalues are greater relative to the rest than those of the UK, but the difference is only small. Experimenting with different values of the embedding dimension, m, does not alter these results.

The eigenvectors associated with the six largest eigenvalues are plotted in figures 6a and 6b for the US and the UK respectively. Here there is a marked difference between the two countries. As discussed in section 2, the vectors represent principal directions in the cloud of points formed by the embedded series. They therefore indicate the pattern or regularity that exists in the data on the time scale covered by the choice of embedding dimension. For the US data, the first four eigenvectors have a strong pattern of regularity. By contrast, for the UK only the first two show any regularity at all, and the remainder are highly irregular.

The pattern of the eigenvalues does in fact look similar to that of a purely random series and, therefore, to try to assess the significance of these results, we carried out a bootstrapping exercise. From each of the original series, a random draw with replacement was carried out to generate a new data series of the same length. By construction, there is only a small probability that this series has any discernible structure. The singular spectrum of this artificial data was calculated, and the procedure repeated 500 times.

This exercise confirms the impression gained from inspection of the singular spectrums and eigenvectors for each of the data series.

The singular spectrum of the UK data seems to be indistinguishable from the spectrum of the artifically generated random series. For example, in terms of the largest eigenvalue of each spectrum, no less than 256 of the random series have one which is larger than that of the UK data itself. The proportion of trials for which the eigenvalue is larger than the corresponding value of the original series is plotted for each successive eigenvalue in figure 7a. The proportion shows some pattern of fluctuation but this is not at all significant. The variation is typical of what is observed when any one of the artificially generated series is compared with the remaining trials.

The results for the US data are quite different. In this case, the spectrum of the original series is distinct from those of the artifically generated random series. Figure 7b repeats the exercise of figure 7a for the US data. For the first ten values, the proportion of trials in which the artifical data have eigenvalues larger than the actual data is very small, and there is then a very clear crossover point. It is important to remember, however, that the range of the singular spectrum of the US data is small, so that the larger values do not dominate the rest. The degree of structure is therefore not large.

Overall, taking into account the evidence from the eigenvalues, the eigenvectors and the results of the bootstrapping exercise, the clear interpretation of this analysis is that quarterly changes in real GDP in the UK possess no underlying structure. To all intents and purposes, it is indistinguishable from a random series. This offers strong support for the Fisher-Slutsky hypothesis that the business cycle is inherently unpredictable. Although the US data does show some small sign of structure or regularity, it is small and is by no means so clear as to contradict the hypothesis.

4. Conclusion

The programme of research in macro-economic modelling and forecasting now spans a period of some fifty years. Particularly during the most recent decades, a great deal of work has been carried out. The content of the models follows closely the fashions of economic theory. Perhaps more importantly, the methodology of time series econometrics which is used to build empirical equations for forecasting has become very much more sophisticated.

But despite these apparent advances, the forecasting record of the models remains poor. Indeed, from a scientific standpoint it is hard to avoid the phrase ‘very poor’. For forecasts of macro-economic variables such as the growth in national output (GNP or GDP) and inflation, the track record of conventional forecasts is really no better than the simple rule that next year’s value will be the same as this year’s.

The use of statistical techniques for nonlinear modelling is relatively new in macro-economics. A small number of studies suggest that such techniques can improve upon linear models in terms of the power of the in-sample fit. But even in terms of just one-quarter ahead forecasts for GDP growth, say, the variance of the forecast errors of such models is still around 90 per cent of the variance of the actual data.

In this chapter, we apply the technique of singular spectrum analysis to post-war quarterly series of GNP/GDP growth in the United States and the United Kingdom. The technique has been developed over the past decade or so, and is well established in the literature on the analysis of dynamic systems. It is designed to characterise the general properties of a time-series such as the degree to which the series has some discernible structure or the extent of the noise present in the data. In other words, singular spectrum analysis can be used to identify consistent or regular features of data, for example periodicity or more general systematic behaviour, and also provides information about the nature of the random or noisy component of a series.

The results suggest very clearly that the inability to make satisfactory forecasts of GDP growth, in the sense that the variance of the forecast errors over time is less than the variance of the data, arises from fundamental characteristics of the data. This is particularly the case for the UK, where singular spectrum analysis suggests that the series is effectively indistinguishable from a purely random series. The data for the US appear to possess a certain amount of structure, but the series is heavily dominated by noise.

The idea that movements in GDP over the course of the business cycle are inherently unpredictable is not new in economics, and some of the early quantitative thinking about the cycle, by for example Fisher (*op.cit.*) and Slutsky (*op.cit.*) in the 1920s and 1930s, advanced this as a hypothesis. The use of singular spectrum analysis confirms the validity of the hypothesis.

If this evidence existed purely in isolation we would suggest that the chapter offered strong but not conclusive support for the hypothesis that meaningful short-term forecasts of real output growth could not be carried out. However, it does not by any means exist in isolation.

The actual ex ante forecasting record of the past thirty or more years supports it. The conclusion we draw is that the forecasting record is poor, not necessarily because of bad economic theory or bad statistical technique on the part of the forecasters, but because of a deep lack of structure in the data. In the current state of scientific knowledge, short-term economic forecasting is a fairly hopeless task.

A great deal of economic policy in the West since the war has been and still is conducted on the basis of short-term forecasts of the economy. Governments try to anticipate the movements of the economy over the course of the business cycle and to take action to correct any adverse consequences which would arise if such predictions were correct. Again, criticisms of such a strategy are by no means new, with Milton Friedman, for example, in the 1950s arguing that government intervention was as likely to increase the overall variance of GDP growth over the course of the cycle as it was to reduce it. Keynes was in many ways a rather delphic economist, and it is certainly possible to interpret his *General Theory,* written in the 1930s, in the same way. Singular spectrum analysis confirms this view.

It must be emphasised, however, that the results of this chapter cannot be interpreted to mean that governments should never intervene in the economy. Keynes, after all, was writing at a time, not when it was predicted that there would be a major depression, but when the economies of the West were actually in one and he was addressing the question of what governments could do about it. The results of this chapter are not concerned with this latter point. Rather they show that the attempts of governments to anticipate and correct short-term fluctuations over the course of the business cycle are in general fruitless, given the inherent properties of the data series involved.

References

Broomhead,D.S. and King,G.P. (1986) ‘Extracting qualitative dynamics from experimental data’ *Physica*, **20D**, pp.217-236

Clements,M.P. (1995) 'Rationality and the role of judgement in macroeconomic forecasting' *Economic Journal,* vol.105, no.429, pp.410-420

Chatfield,C. (1995) 'Model uncertainty, data mining and statistical inference', *J.R.Statist.Soc,* **A,** vol.158, part 3

Fisher,I (1925) 'Our unstable dollar and the so-called business cycle', *Journal of the American Statistical Association*, June

Gershenfeld,N.A. and Weigend,A.S. (1993) 'The future of time series: learning and understanding', in Gershenfeld and Weigend, eds., *Time Series Prediction: Forecasting the Future and Understanding the Past*, Addison-Wesley, Reading MA

Klein,L.R. (1947) 'The use of econometric models as a guide to the policy process', *Econometrica*, vol.15, pp.111-152

McNees,S.K. (1988) 'How accurate are macroeconomic forecasts?', *New England Economic Review*, July/August, pp.15-36

Medio,A. (1992) *Chaotic dynamics: theory and applications to economics*, CUP

Mullin,T. (1993) 'A dynamical systems approach to time series analysis', in T.Mullin, ed., *The Nature of Chaos,* Oxford Scientific Publications

OECD (1993) *Economic Outlook*, June, Paris

Ormerod,P. (1979) ed., *Economic Modelling*, Heinemann

Potter,S.M. (1995) 'A nonlinear approach to US GNP', *Journal of Applied Econometrics*, vol.10, no.2, pp.109-126

Priestley,M.B. (1981) *Spectral Analysis and Time Series*, 2 vols. London, Academic Press

Slutsky,E. (1937) 'The summation of random causes as the source of cyclical processes', *Econometrica*, vol.5, no.2, pp.105-46

Tiao,G.C. and Tsay,R.S. (1994) 'Some advances in non-linear and adaptive modelling in time-series analysis', *Journal of Forecasting*, vol.13, no.2, pp. 109-132

Vautard,R. and Ghil,M. (1989) ‘Singular spectrum analysis in nonlinear dynamics, with applications to paleoclimatic time series’ *Physica,* **35D**, pp.395-424

Wallis,K.F. (1989) 'Macroeconomic forecasting: a survey', *Economic Journal,* vol.99, no.394, pp.28-61