The Goodwin model and the periodicity of unemployment and factor shares in the UK

Paul Ormerod (**pormerod@volterra.co.uk**)

Volterra Consulting Ltd

The Old Power Station

121 Mortlake High Street

London SW14 8SN

December 2001

Summary:

Unemployment and the share of wages in national income are key variables in the Goodwin growth cycle model. This paper examines the empirical evidence on the periodicity of these two variables using a long time-series of annual data for the UK.

Two techniques of modern statistical analysis are used. First, spectral analysis applied to the data after it is passed through a kernel filter to remove very low frequency persistence of the data in level form. Second, singular spectrum analysis, based on the eigenstates of delay matrices formed from the original data.

We find definite evidence of the existence of regular periodicity of both unemployment and the labour share. This is predominantly at frequencies usually associated with the business cycle of between 4 and 8 years. However, in both cases the cycle is only determined weakly and the data contain a great deal of noise. The results using the two separate techniques are very similar.

Theoretical specifications of the Goodwin model need to be able to generate results which are consistent with this evidence.

1. Introduction

The classic paper by Atkinson [1] examined the empirical time scale of a range of models in economic theory. One such model investigated was the Goodwin model [2] of the business cycle. Atkinson found that for reasonable values of the parameters in this model, the period of the typical cycle was 'better suited to explaining the 16-22 year "Kuznets" cycle than the post-war trade cycle'.

This paper looks in detail at the empirical evidence for the periodicity of unemployment and the share of wages in national income in the UK, using a variety of modern statistical techniques over a long run of data. We also examine how well the cycles are determined.

2. Empirical evidence from the UK

Figure 1a plots the UK unemployment rate from 1870 through 1993. The 1870-1965 data are taken from Feinstein [3], and later years use Department of Employment figures spliced onto the Feinstein basis.

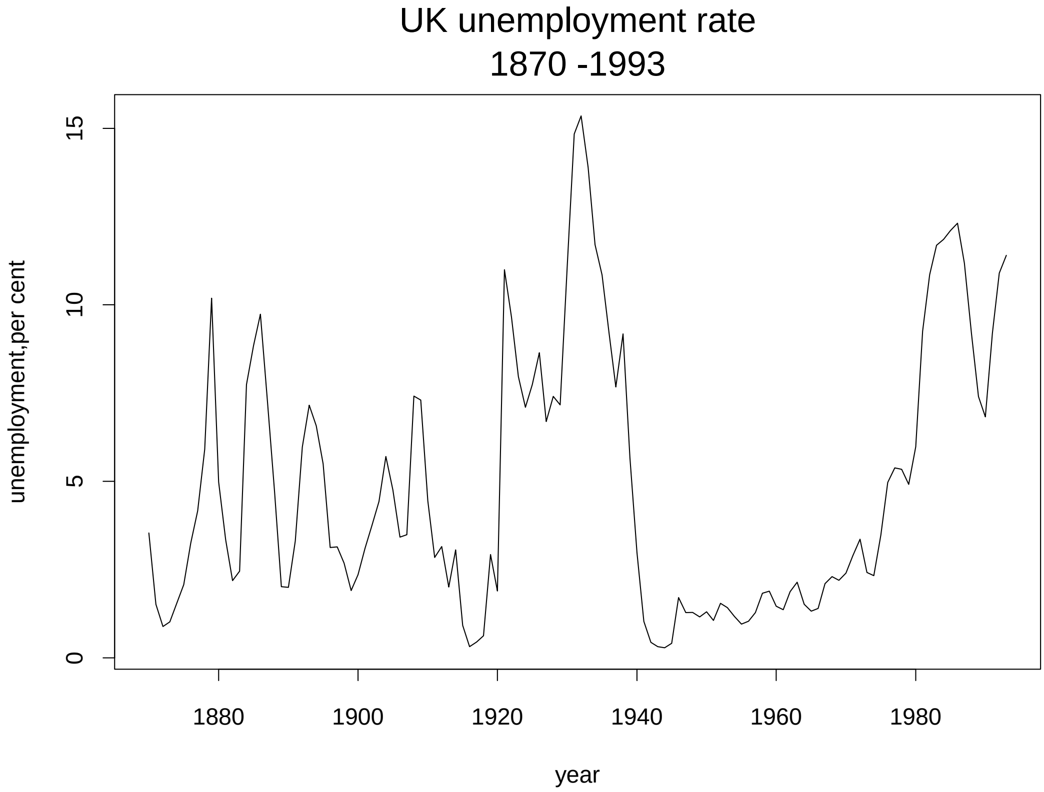

Figure 1a

By inspection, the series exhibits persistence over long periods of time. Unemployment was very low in the war years and in the period 1945-1975. It was very high in the 1920-39 period, and also from 1980 to the early 1990s. This complicates the issue of trying to extract information on frequencies at or around those typically associated with the business cycle (2 - 8 years).

It is arguable that the series as a whole is not really suitable for statistical analysis. Economic history suggests several quite distinct regimes. But, with this qualification in mind, we proceed.

Figure 1b plots labour's share of national income, again taken from Feinstein for 1870-1965, the data for later years being taken from the OECD *National Accounts* and spliced to the Feinstein data. Transparently, this is strongly trended.

Figure 1b

In many ways, the autocorrelation function is the basis for far more sophisticated ways of estimating the periodicity/frequency of time series data. The power spectrum, for example, is the Fourier transform pair of the autocorrelation function. Figures 2a and 2b show at once the problems which are involved in investigating the existence of periodicity corresponding to that of the business cycle. They show, respectively, the autocorrelation functions of the residuals of unemployment and the labour share from the least squares linear fits of the data against time. Persistence in both is very strong, and dominates the data.

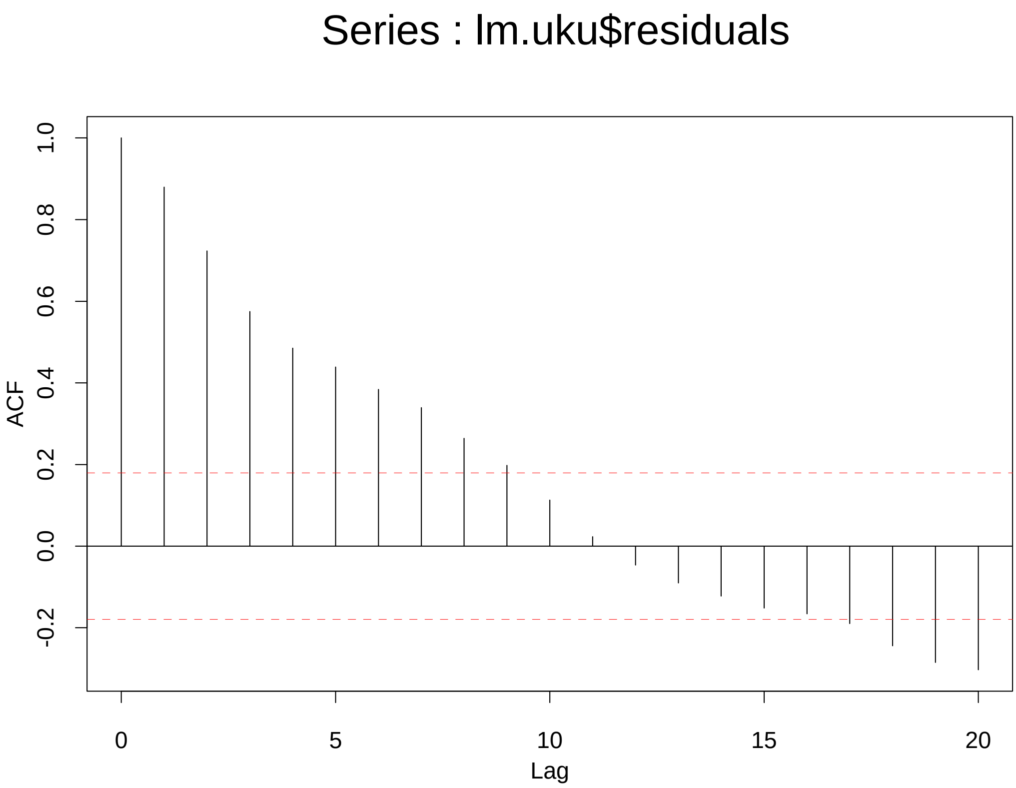

Figure 2a

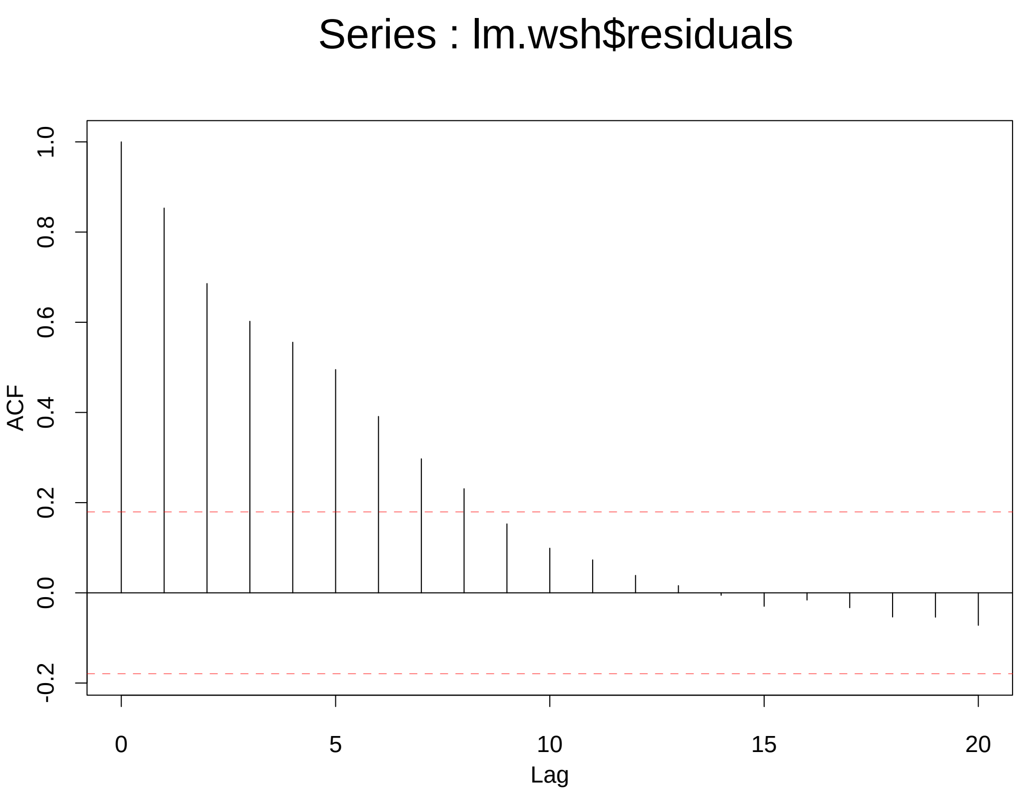

Figure 2b

The underlying periodicity at around the frequency of the business cycle may potentially be revealed more clearly when more sophisticated smoothing than simple linear regression is carried out on the data. The package S-Plus offers a wide variety of smoothing techniques. We examined several nonparametric approaches, each of which rely on the data to specify the form of the model, and which fit a curve to the data points *locally*. In other words, at any point the curve at that point depends only on the observations at that point and some specified neighbouring points.

By appropriate calibration of the smoothing procedure, it is possible to obtain similar results for a range of them, including locally weighted regression smoothers, cubic smoothing splines and a kernel-type smoother. The results we report are with the kernel smoother. Essentially, this is a type of local average smoother that, for each target point x_i in predictor space, calculates a weighted average $\hat{y}_i$ of the observations in a neigbourhood of the target point. A parameter known as the bandwidth parameter determines how large a neighbourhood of the target point is used to calculate the local average. (A range of kernels is available, but Hastie and Tibshirani [4] show that the choice of bandwidth is much more important).

More formally, we used the following procedure. First, calculate the weighted average $\hat{y}_i$ of the observations in a neighbourhood of the target point:

$$\hat{y}_i = \sum_{j=1}^{n} w_{ij} y_j \qquad (1)$$

where
$$w_{ij} = \frac{K\left(\frac{x_i - x_j}{b}\right)}{\sum_{k=1}^{n} K\left(\frac{x_i - x_j}{b}\right)}$$

are weights which sum to one:

$$\sum_{j=1}^{n} w_{ij} = 1$$

The function K used to calculate the weights is called the kernel function, which typically has the following properties:

i) $K(t) \geq 0 \forall t$

ii) $\int_{-\infty}^{\infty} K(t) dt = 1$

iii) $K(-t) = K(t) \forall t$

The parameter b is the bandwidth parameter, which determines how large a neighbourhood of the target point is used to calculate the local average. A large bandwidth generates a smoother curve.

In this context, a trade-off must be achieved between a smoother which removes the very low frequency persistence of the original data, and at the same time does not smooth so closely to the data itself as to remove all structure. By concentrating the local neighbourhood sufficiently tightly, for example, near-perfect fits to the original data are readily achieved.

Figure 3 plots unemployment along with the data generated by the kernel smoothing procedure with bandwidths of 4 and 8. With bandwidths below 4, the fit to the data rapidly becomes very tight, and so these were not used.

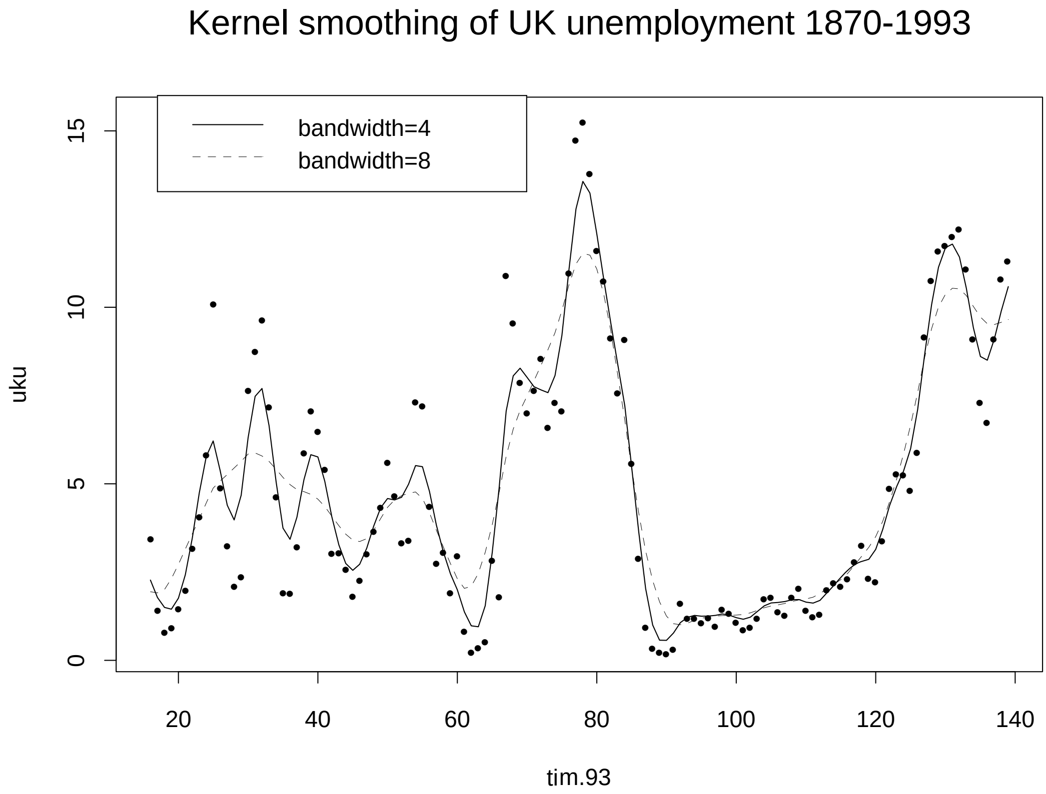

Figure 3

Figure 4 plots the residuals from the kernel smoothing technique (ie: the actual data minus the smoothed values) against the actual data. With bandwidth equal to 8, there is fairly clear heteroskedasticity, with large values of the unemployment rate being associated with large positive residuals. Although there is still some sign of this with the lower bandwidth, it is distinctly less marked. A bandwidth of 4 therefore seems reasonable: it does not fit the data too tightly, yet appears to eliminate from the residuals much of the longer term persistence in the raw data. To check the robustness of the results, however, we use residuals from both bandwidths (4 and 8) in subsequent analysis.

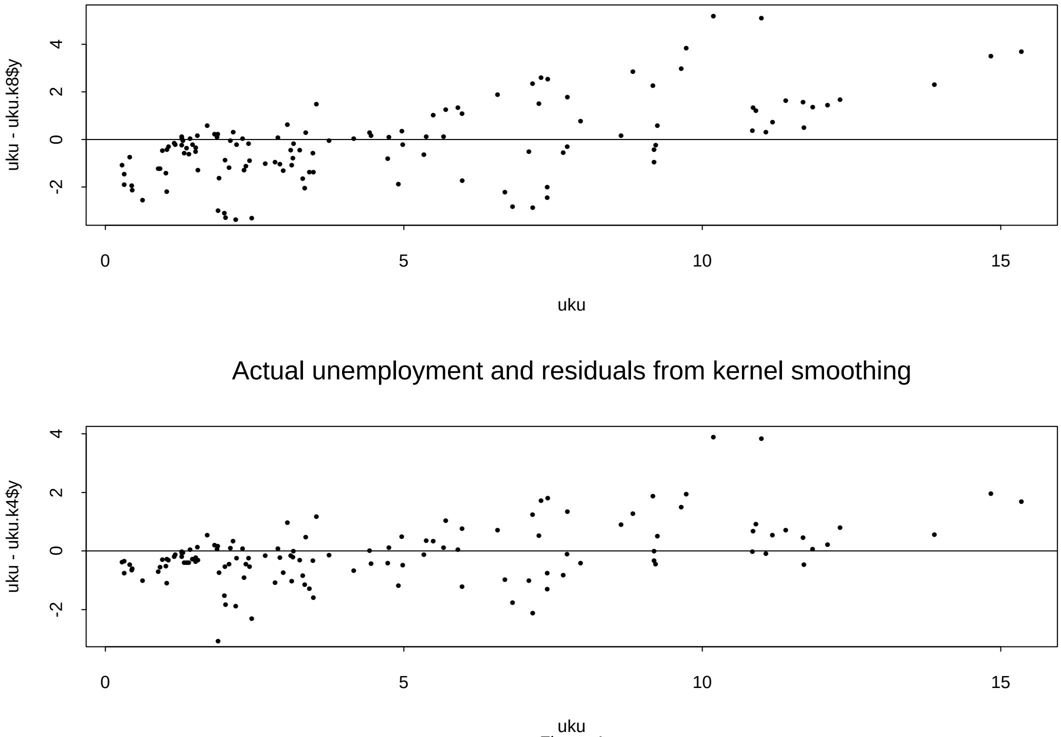

Figure 4

Figures 5 and 6 plot, respectively, the labour share data and kernel smoothing with bandwidth of 5, and the residuals from this kernel smoothing against the actual data. Again, this seems to be a reasonable smoothing in this context, but we also used bandwidths of 3 and 7 in subsequent analysis as a check on the robustness of results.

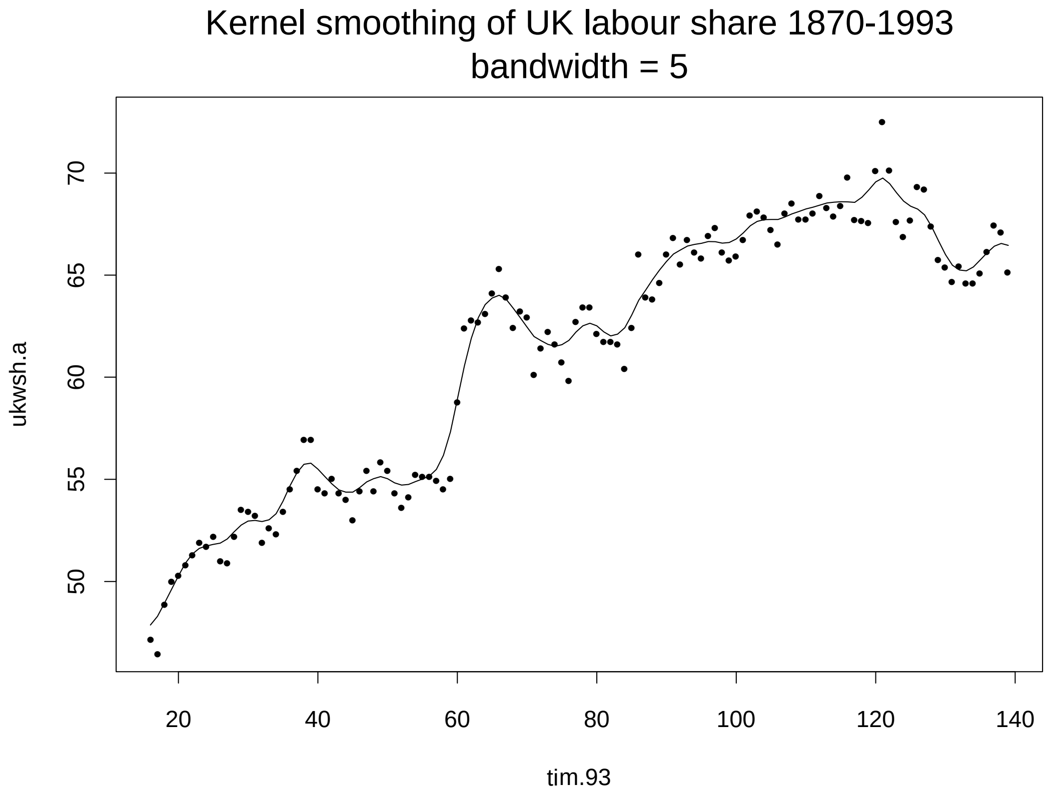

Figure 5

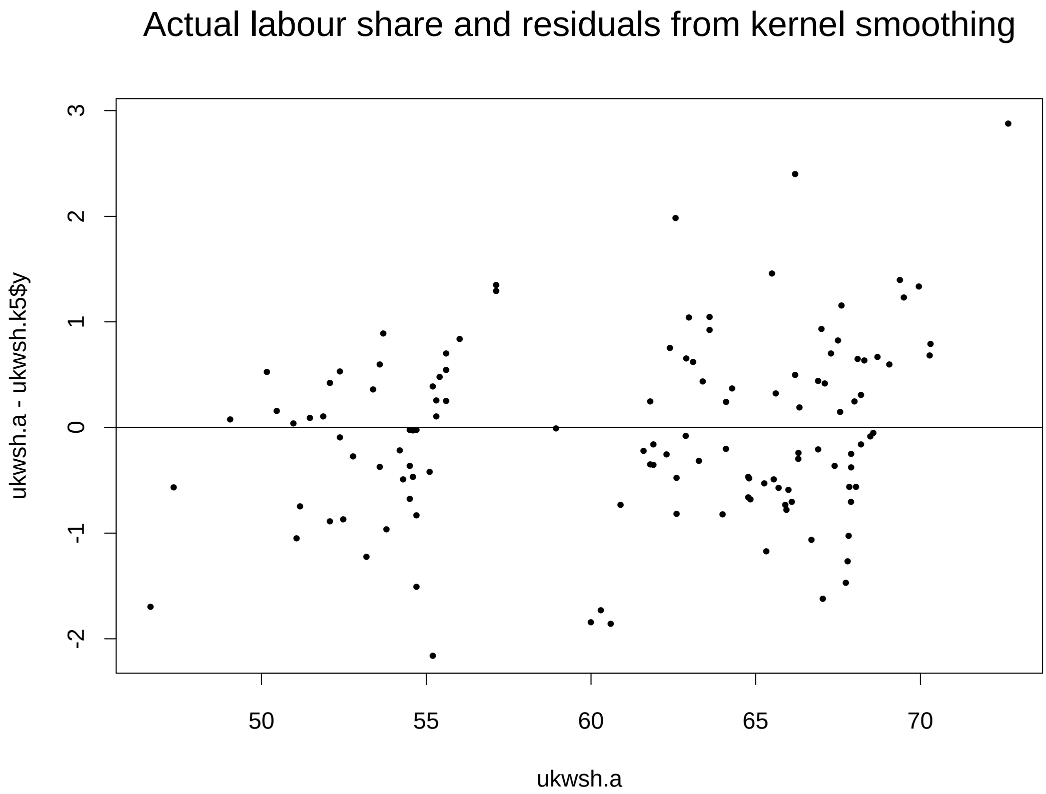

Figure 6

Figure 7 plots the autocorrelation functions of the residuals from the smoothings of the unemployment data, with bandwidths of 4 and 8. The dotted lines show the (approximate) 95 per cent confidence intervals around zero.

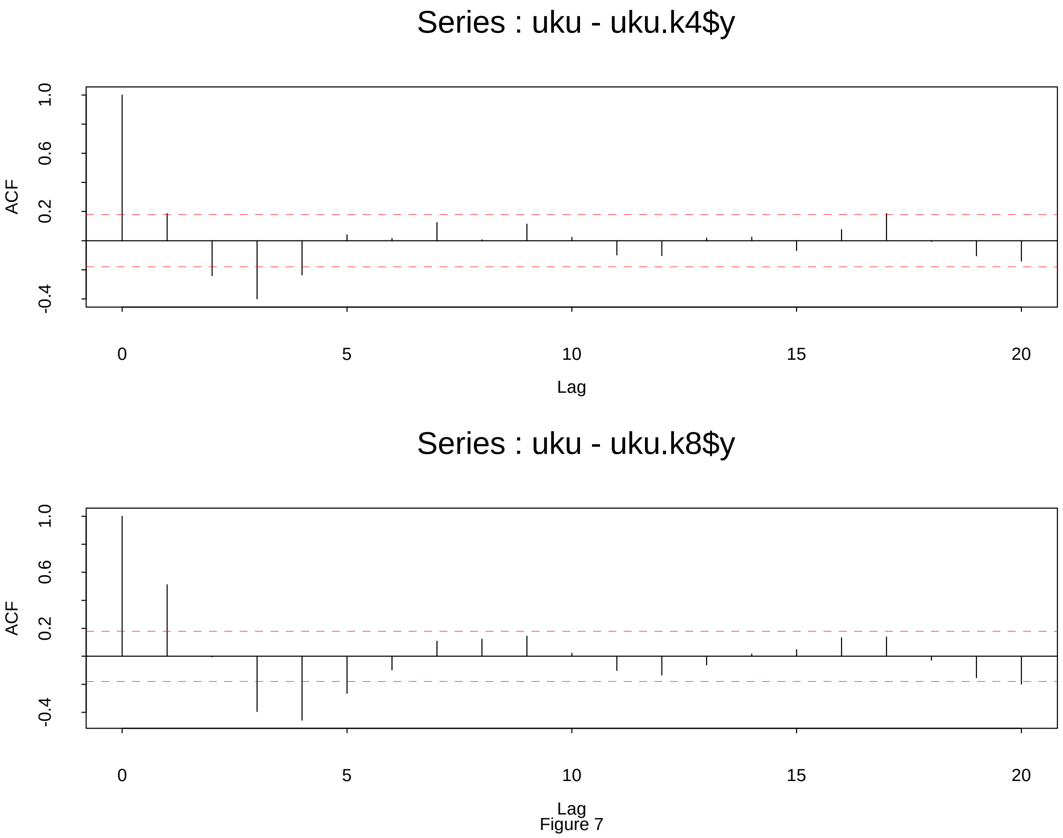

Figure 7

Most of the long-term persistence of the raw data is removed by these filters. Both, in slightly differing ways, display significant values over the lags 1 through 5, but of a non-persistent form. There is little clear evidence of significance at higher order lags, except a hint at around lags 16 - 20.

The ACFs of the labour share residuals plotted in Figure 8 show a similar pattern.

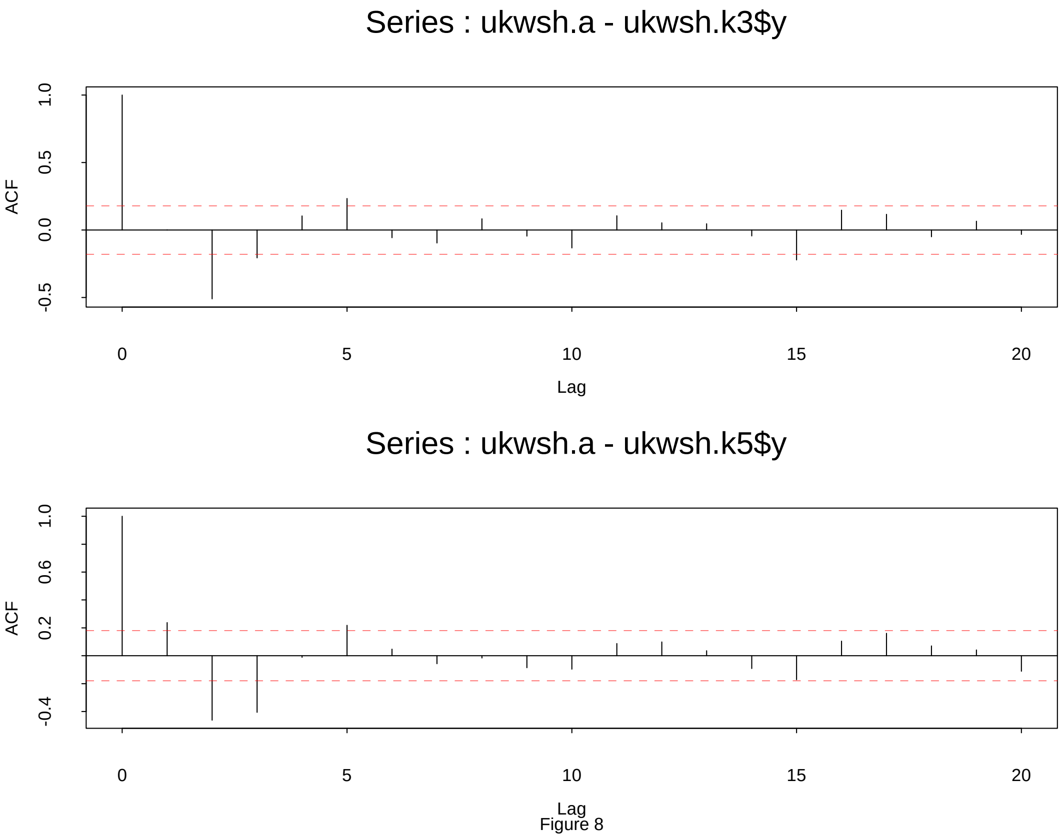

Figure 8

Classical spectral analysis of the unemployment series itself reveals very little of value. We need to ensure that the mean of the data is zero, and this is accomplished by subtracting a linear trend from the original series ie: the analysis is carried out on the series $(u(t) - \alpha + \beta t)$, where $\alpha + \beta t$ is the least squares linear fit to the data.

Figure 9 plots the periodogram of this series, both in its raw form and with a mild amount of smoothing applied[1]. The top chart plots the unsmoothed periodogram. The peak is very clearly at the far right hand side of both charts, at around the frequencies 0.025 - 0.05. Given that the data is annual, this implies that the principal periodicity of the data is between 25 and 40 years (ie: 1/0.05 and 1/0.025). A very similar result obtains for the labour share.

[1] The smoothing in this instance is applied to the periodogram. It is distinct from the kernel smoothing of the actual data discussed above.

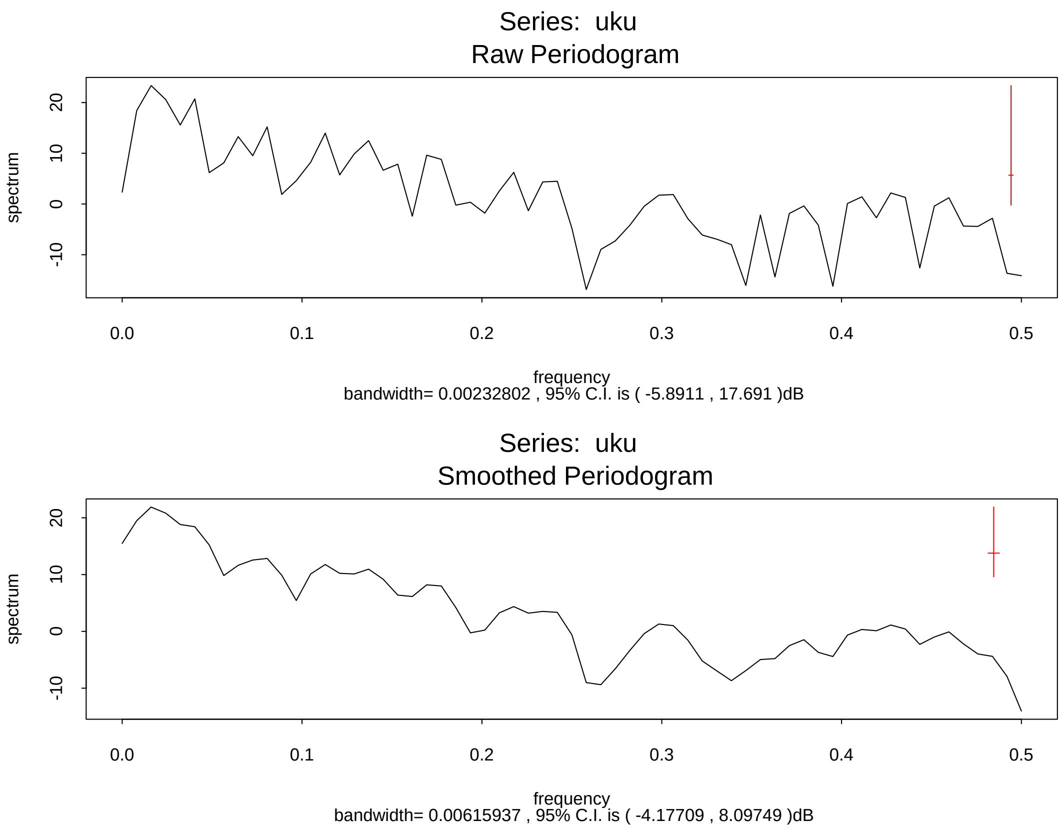

Figure 10 shows the smoothed periodogram of the unemployment residuals from the kernel smoothing with bandwidth of 4. The peak is around the frequencies of 0.12 - 0.18, indicating a cycle of some 5.5 - 8 years. But the degree of concentration is not dramatic. The 95 per cent confidence interval is printed at the bottom of the chart, and is marked by the vertical line at the top right hand corner. There is some evidence for secondary cycles at frequencies of just under and slightly over 0.4, suggesting cycles of around 3.5 and 2.5 years respectively.

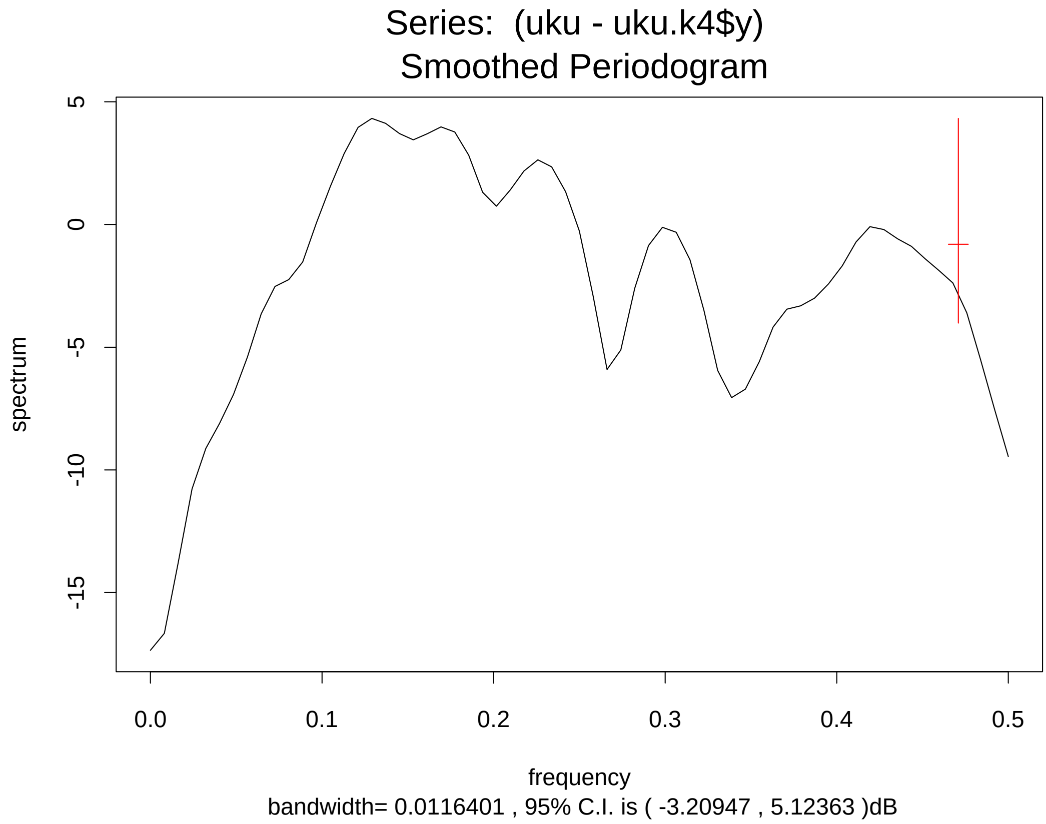

The use of a bandwidth of 8 in the kernel smoothing gives a very similar result, with slightly more concentration at a cycle of some 6 - 8 years. In neither case is there any evidence pointing to the existence of a longer cycle of 15 - 20 years.

The smoothed periodogram of the labour share residuals from the kernel with bandwidth of 5 is plotted in Figure 11.

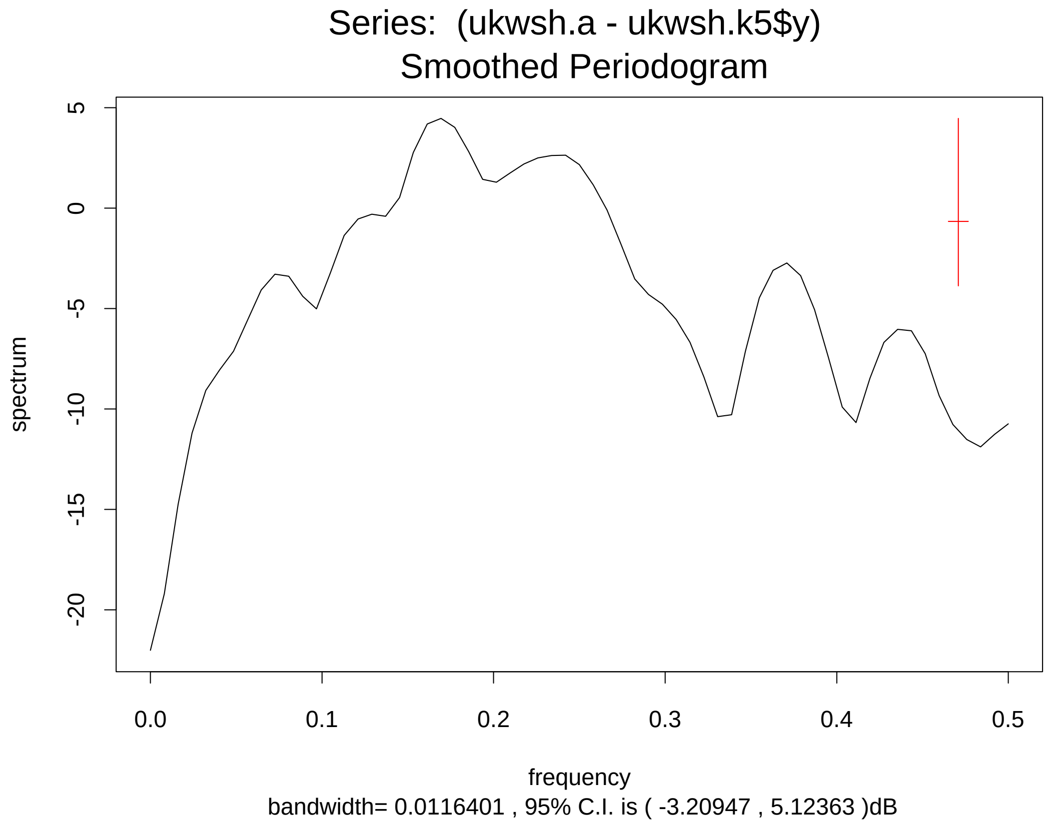

This is similar, but not identical, to that of unemployment. The main concentration is at slightly higher frequencies, in the range of around 0.15 - 0.22, suggesting a cycle of 4.5 - 6.5 years. The secondary frequencies also exist, around 2.75 and 2.25 years. The results with a bandwidth of 7 in the kernel smoother are very similar. A bandwidth of 3 fits the raw data more closely than one of 5, and as a result there is less concentration of the periodogram of the residuals in any particular range of frequencies, but qualitatively the results are again very similar.

An alternative way of investigating the periodicity of time series data is provided by a modern technique with the rather unhelpful name of singular spectrum analysis (SSA). Two seminal references are Broomhead and King [3] and Vautard and Ghil [4]. A more accessible recent account is given by Mullin [5], and an illustration of its application to macro-economic time series data in Ormerod and Campbell [6].

SSA is a powerful technique which can provide information on several important but distinct aspects of the fundamental properties of a time series. The one of interest in this context arises from the phase portrait reconstruction of any attractor which might exist in the data. The description below is highly condensed, but might help to illustrate its relevance.

A very simple method of trying to identify an attractor - should one exist - in a time series is the so-called method of delays. In three dimensions, for example, from a single time series, z(t) say, the points (z(t), z(t-1), z(t-2)) can be plotted and connected in sequence. This primitive approach is able to reveal the underlying structure of, for example, each of the individual time series generated by the Lorenz equations.

In practice, however, and particularly with noisy series this approach has very definite limitations. These mainly arise from the fact that the original time series is projected onto an arbitrary basis - in the example above, onto itself at time t, time (t-1) and time (t-2).

SSA can be thought of as finding the vectors onto which the data can be projected optimally. Initially, a so-called delay matrix is formed from the original time series. Given a series z(t), where t runs from 1 to n, choose a maximum delay d, and the first row of the matrix is z_1, z_2,..., z_d. The second row consists of z_2, z_3,..., z_{d+1}, and so on. The eigenvectors associated with the principal eigenvalues of the moment matrix of this delay matrix form the appropriate co-ordinate system.

These eigenvectors give us information on the periodicity of the original data. Figure 12a plots them for the unemployment residuals of the bandwidth 4 kernel smoother. The first eigenvector has a cycle of around 8 years, and the second around 7. There is less regularity in the vectors from 3 onwards. In terms of classical spectral analysis, the periodogram is concentrated to a certain extent around this periodicity. The relative sizes of the eigenvalues also give us information on the degree of concentration, as it were, which is not high. Only the first two stand out from the rest, and even then not to any great extent.

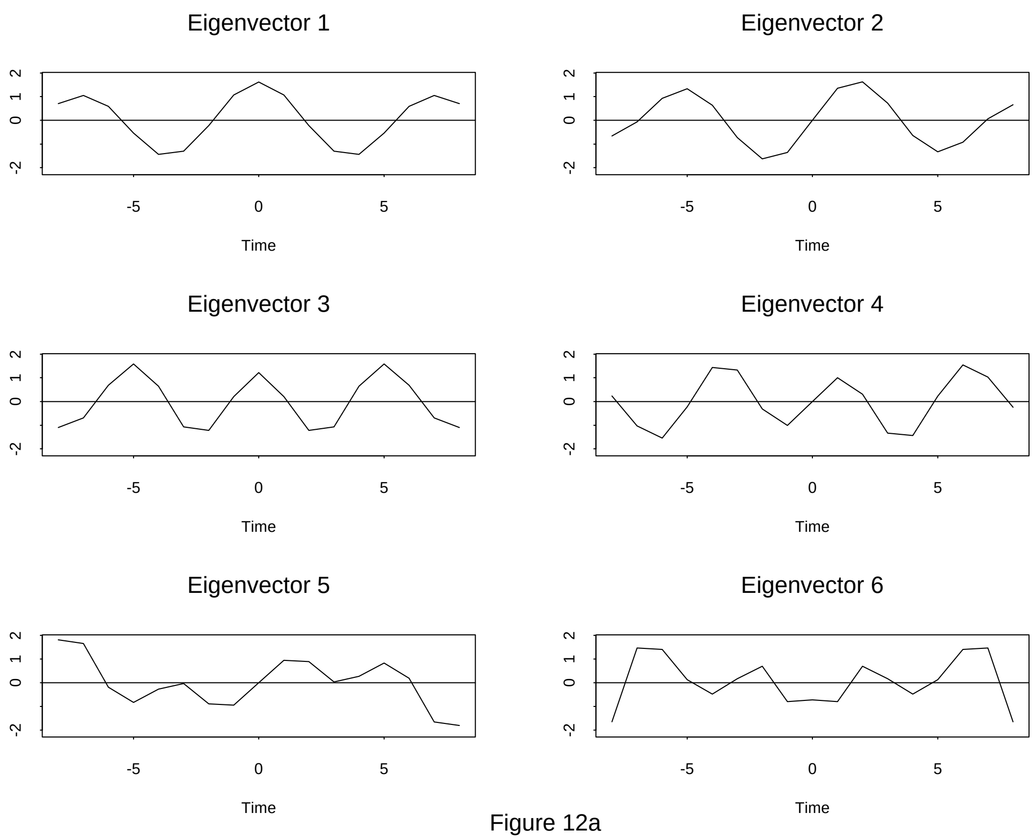

Figure 12a

The results with the labour share are plotted in Figure 12b, using bandwidth 5. As with classical spectral analysis, the underlying cycles appear to be somewhat shorter than with unemployment, with the first eigenvector exhibiting a cycle of around 5 years.

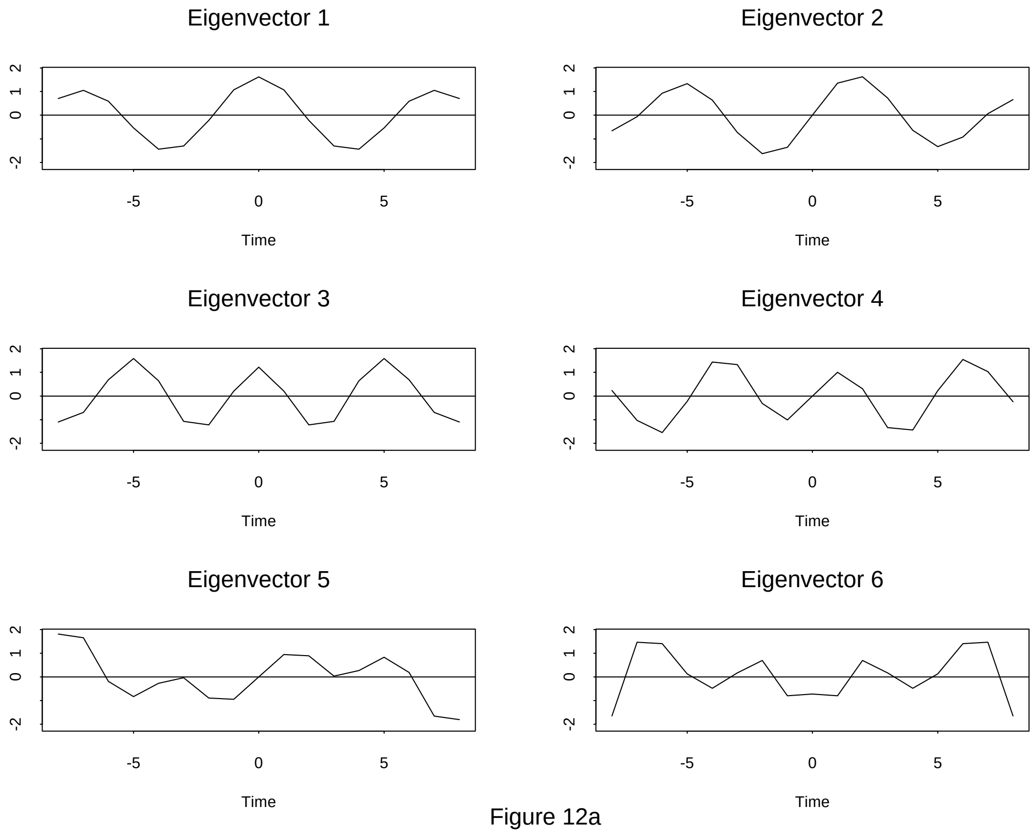

Figure 12a

The results of both are robust with respect to different values of the bandwidth of the kernel.

Overall, using modern statistical techniques, there is clear evidence on the following:

i) there is regular periodicity in both unemployment and the labour share of national income in the UK

ii) this periodicity is at the frequency of the business cycle

iii) neither variable exhibits very strong cyclical behaviour

iv) unemployment has an underlying cycle of some 5.5 - 8 years

v) the labour share cycle is somewhat shorter, at 4.5 - 6.5 years.

vi) ‘best’ estimates are probably 7 - 8 years for unemployment and around 5 years for the labour share

References

1. C.H.Feinstein, *National Income, Expenditure and Output of the UK 1855-1965*, CUP, 1972

2. T.Hastie and R.Tibshirani, *Generalized Additive Models*, Chapman and Hall, London, 1990

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

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

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

6. P.Ormerod and M.Campbell, 'Predictability and Economic Time-Series' in C.Heij, J.M.Schumacher, B.Hanzon and C.Praagman, eds., *System Dynamics in Economic and Financial Models,* John Wiley, 1997