RESERVE BANK OF NEW ZEALAND

DP2000/04

A multivariate unobserved components model of cyclical activity

Alasdair Scott

January 2000

JEL classification: C32, E32

Discussion Paper Series

Abstract[1, 2]

This paper presents results from the estimation of a multivariate unobserved components model of cyclical activity. The model is motivated by a desire to let the data speak as much as possible, and hence to avoid imposing ad hoc and unjustifiable assumptions about trends and cycles. Estimated over the period 1970:1 to 1999:3 via the Kalman filter and maximum likelihood, the model identifies a common, trend-reverting component to real output, unemployment and capacity utilisation. The structure of the model allows an interesting factor interpretation to be put on the estimate of the output gap. These estimates are consistent with priors, but there is no consistent match to any one simple smoother such as the HP filter.

---

1 I wish to thank participants at the New Zealand Econometric Study Group Meeting in Hamilton, 12-13 February 1999, and in particular my discussant, John Haywood, for comments on an earlier version of this paper. I also wish to thank John McDermott and other colleagues in the Economics Department of the Reserve Bank of New Zealand for helpful discussions.

1 Introduction

One of the most important pieces of information for an economic policymaker is a measure of the extent to which current activity is above or below its long-run trend or "sustainable" level. However, the policymaker immediately faces two problems. First, the concept of "economic activity" may not be perfectly captured by any one standard series from a national statistical office. Second, "potential output" cannot be directly observed, and instead has to be somehow inferred from the data.

Typically the output gap – the difference between the actual and potential levels of output – is derived by running GDP through some kind of smoother, such as the Hodrick-Prescott (HP) filter.[3,4] However, it is now well-known that the HP filter imposes properties on the estimated cycle by construction, when in fact the properties of the cycle may be exactly what is in question.[5] Ideally, we would let the data speak and estimate a structural model for the (permanent) trend and the trend-reverting component. At the same time, we would like to be able to use series other than output which are believed to hold information about the level of economic activity.

Following Clark (1989), I estimate a multivariate unobserved components model designed to provide an estimate of common cyclical activity.[6] The model is motivated by a desire to let the data speak as much as possible, and hence to avoid imposing ad hoc and unjustifiable assumptions about trends and cycles. Estimated over the period 1970:1 to 1998:4 via the Kalman filter and maximum likelihood, the model successfully identifies a common, trend-reverting component to real output, unemployment and capacity utilisation.

The estimated output gap appears reasonable, but is not matched by the results from any one simple time-series smoother. Compared with the popular HP filter with $\lambda = 1600$, the model attributes more 'demand' movement in the last 15 years. At the same time, the structural nature of the model offers an interesting interpretation of the output gap according to capital and labour inputs, and this interpretation is lacking in "naïve" time series estimates of the output gap.

But how much weight should the policymaker put on such a model? There can be considerable uncertainty around the levels of trend and cycle, and the model suffers

---

3 This paper supposes that the estimate of underlying cyclical real activity can be treated as an output gap. Strictly speaking, however, since this implies some notion of a level of output that can be sustained without inflationary or deflationary pressure, the model should connect the nominal and real sides of the economy. Such a model is to be found in Kuttner (1992).

4 See Hodrick and Prescott (1997).

5 See, *inter alia*, Harvey and Jaeger (1993) and Canova (1998).

6 In this paper, the term "business cycle" will refer to the "growth" cycle, rather than the "classical" cycle. The classical cycle is concerned with movements in the *level* of economic activity. The growth cycle concerns fluctuations in the level of *detrended* economic activity. The two may be very different. For a discussion, see Pagan (1997a, 1997b).

heavily from the "end-point" problem.[7] Because it is the current estimate of the output gap that matters for the policymaker, this type of uncertainty is a substantial disadvantage.

The paper proceeds as follows. In section 2, problems with univariate models are described. This discussion motivates the multivariate model outlined in section 3. Results are presented in section 4, and in section 5 the model is evaluated from a policymaker's perspective. Section 6 finishes the paper with some concluding remarks.

2 The univariate problem

Before describing the multivariate model, it is helpful to briefly consider some of the problems with univariate treatments, as these motivate the multivariate approach adopted later in the paper. I start with the well-known HP filter, and show how the estimated gaps will depend on the choice of smoothness parameter.[8] The important link to the next section comes by showing that the HP filter is simply a restricted version of a more general model, the parameters of which can be estimated. The addition of economic structure in section 3 is then a natural extension.

Using the HP filter to infer a level of potential output relies on the assumption that the trend can be well-captured by the smoothed values. The user therefore needs to have some 'degree of smoothness' in mind when applying the filter. Prescott (1986) used a value of $\lambda = 1600$, on the grounds that this resulted in growth cycle of approximately 32 quarters in duration.

However, there is nothing that says that λ must always be set to 1600. A higher value will mean that the trend is smoother, and a lower value will mean that the resulting trend will tend to follow the data more closely. To illustrate, figure 1 plots HP trends for values of $\lambda = 1600$ and $\lambda = 100{,}000$ for New Zealand real output.

The value of the smoothing parameter will correspondingly have implications for the value of the output gap (figure 2). As we can see, the value for the smoothness parameter affects both the magnitude and the duration of the recession centered around 1991-1992. Moreover, one estimate of trend implies that the economy is slightly above potential at the end of the sample, whereas the other implies that it is well below. This illustrates the sensitivity of the "end point" to assumptions made in the use of a smoother. Unfortunately, it is precisely the end point that the policymaker is most reliant on for policy decisions.

---

7 By end-point problem, I refer to the way that the estimate of the level of the trend or the gap will be altered as new data are added to the sample.

8 Some definitions are necessary here for clarity. In the discussion of the model estimated via the Kalman filter, a "filtered" estimate is one-sided – that is, it uses information only up to time *t*. A "smoothed" value is two-sided and uses information from the whole sample, up to time *T*. In this sense, the HP filter is really a smoother, since it can be thought of as a two-sided moving average.

Methods do exist for determining "optimal" values for smoothness parameters for the HP filter and other smoothers.[9] However, this assumes that a smoother like this is capable of accurately capturing the permanent trend. To illustrate this problem, it helps to show that the HP filter is really a restricted version of a more general time series model, the *local linear trend* model. This posits that a series, y_t, can be decomposed into an underlying trend, μ_t, and an irregular component, ε_t. The trend is a random walk process with a drift (or slope) term, β_t, that is itself a random walk. The model takes the form

$$y_t = \mu_t + \varepsilon_t, \qquad \mathrm{E}(\varepsilon_t) = 0, \quad \mathrm{Var}(\varepsilon_t) = \sigma^2_\varepsilon, \tag{1a}$$
$$\mu_t = \mu_{t-1} + \beta_{t-1} + \eta_t, \qquad \mathrm{E}(\eta_t) = 0, \quad \mathrm{Var}(\eta_t) = \sigma^2_\eta, \tag{1b}$$
$$\beta_t = \beta_{t-1} + \zeta_t, \qquad \mathrm{E}(\zeta_t) = 0, \quad \mathrm{Var}(\zeta_t) = \sigma^2_\zeta. \tag{1c}$$

The HP filter imposes smoothness by construction – the HP filter is the local linear trend model with zero variance for the innovation in the level of the trend, σ^2_η. Hence, all innovations must come through changes to the slope. The smoothness parameter in the HP filter, λ, is here equivalent to the ratio of the variances of the irregular term and the slope term, so that $\lambda = \sigma^2_\zeta / \sigma^2_\varepsilon$.

Is the smoothness restriction justified by the data? Harvey and Jaeger (1993) propose a model which augments the local linear trend model by introducing a term for a cyclical component in output. This model takes the form:

$$y_t = \mu_t + \varphi_t + e_t, \qquad \mathrm{E}(e_t) = 0, \quad \mathrm{Var}(e_t) = \sigma^2_e, \tag{2a}$$

where y_t is the observed series, μ_t is the trend, φ_t the cycle, and e_t an irregular term with variance σ^2_e. The trend is a local linear trend defined as before:

$$\mu_t = \mu_{t-1} + \beta_{t-1} + \eta_t, \qquad \mathrm{E}(\eta_t) = 0, \quad \mathrm{Var}(\eta_t) = \sigma^2_\eta, \tag{2b}$$
$$\beta_t = \beta_{t-1} + \zeta_t \qquad \mathrm{E}(\zeta_t) = 0, \quad \mathrm{Var}(\zeta_t) = \sigma^2_\zeta, \tag{2c}$$

The stochastic cycle is made up of sinusoids:

$$\varphi_t = \rho \cos \lambda_c \varphi_{t-1} + \rho \sin \lambda_c \varphi^*_{t-1} + \kappa_t, \qquad \mathrm{E}(\kappa_t) = 0, \quad \mathrm{Var}(\kappa_t) = \sigma^2_\kappa, \tag{2d}$$
$$\varphi^*_t = -\rho \sin \lambda_c \varphi_{t-1} + \rho \cos \lambda_c \varphi^*_{t-1} + \kappa^*_t, \qquad \mathrm{E}(\kappa^*_t) = 0, \quad \mathrm{Var}(\kappa^*_t) = \sigma^2_\kappa, \tag{2e}$$

9 By optimal, I mean that there is some kind of defined criterion for selecting a value for the smoothness parameter. For example an example of this approach applied to the HP filter, see Coe and McDermott (1997).

where ρ is a factor representing the amplitude of the cycle and λ_c is the frequency of the cycle in radians.

Now it can be seen that the HP filter is a restricted version of this model, with $\sigma_\eta^2 = 0$, $\lambda = \sigma_\zeta^2 / \sigma_\varepsilon^2$, and $\rho = 0$ or $\sigma_\kappa^2 = 0$. Estimated over US data, Harvey and Jaeger find that σ_η^2 is zero, justifying the imposition of a smoothly-evolving trend. However, the term is non-zero for other series they test on, such as Austrian GNP. This leads the authors to conclude "the striking coincidence between the estimated business cycle component and the HP cycle suggests that the HP filter is tailor-made for extracting the business cycle component *from US GNP*."[10,11]

What about New Zealand output? To illustrate, I estimate the model (2) over New Zealand real GDP, using a sample of 119 observations dating from 1970:1 to 1999:3. The model is clearly not very successful here at uncovering what would look like a business cycle. The fitted trend is very close to actual output. There is a cyclical component present, but it does not match what we would think of as the "output gap." The cycle appears to be almost completely irregular, while the amplitude is extremely small.

We can see what is going on more closely by looking at the estimated parameters:

$$
\begin{aligned}
y_t &= \mu_t + \varphi_t + e_t, & \mathrm{Var}(e_t) &= 2.08\text{e-}8, \\
\mu_t &= \mu_{t-1} + \beta_{t-1} + \eta_t, & \mathrm{Var}(\eta_t) &= 3.04\text{e-}8, \\
\beta_t &= \beta_{t-1} + \zeta_t, & \mathrm{Var}(\zeta_t) &= 1.00\text{e-}2, \\
\varphi_t &= 2.3\mathrm{e}-7 \cdot \cos \cdot 0.95\varphi_{t-1} + 2.3\mathrm{e}-7 \cdot \sin \cdot 0.95\varphi_{t-1}^* + \kappa_t, & \mathrm{Var}(\kappa_t) &= 2.78\text{e-}3, \\
\varphi_t^* &= -2.3\mathrm{e}-7 \cdot \sin \cdot 0.95\varphi_{t-1} + 2.3\mathrm{e}-7 \cdot \cos \cdot 0.95\varphi_{t-1}^* + \kappa_t^*, & \mathrm{Var}(\kappa_t^*) &= 2.78\text{e-}3.
\end{aligned}
$$

The estimated frequency parameter, $\hat{\lambda}_c$, implies a cycle with a period of approximately 6 quarters, and the amplitude estimated for that cycle is small and insignificant, at 2.27e-07. Hence in trying to avoid some of the pitfalls of imposing a smoother on the series we have hit some new problems. This model attributes most of the "energy" of the series to an underlying trend, and little to business cycle frequencies.

---

10 Harvey and Jaeger (1993) p.236; my italics.

11 This "trend plus cycle" model is perhaps too restrictive. Rather than a trend and a truly cyclical component, we are really only interested in decomposing a series into what is permanent and what is temporary. Watson (1986) and Clark (1989) both employ univariate models where the temporary component is modelled as an AR(2) process. Their results, nonetheless, are very similar to Harvey and Jaeger's and indicate that the HP filter is a good approximation to the permanent component of US output data.

3 A common cycle model

One of the advantages of the unobserved components approach used to estimate model (2) is that the trend and the cyclical component are treated simultaneously. Nonstationary series such as output can be modelled without any need for transformations such as differencing. However, it appears that in the case of the univariate model the estimator has a hard time distinguishing between a permanent trend and a temporary component of business cycle frequencies. It may be that the business cycle is not well-captured in the official data, and that other series may lend more information about underlying cyclical activity.

These problems motivate a multivariate treatment of the problem. From a time series point of view, there may be other series where the cycle at business cycle frequencies is more pronounced, and this may help the estimation of the business cycle component in output. From an economic point of view, it also makes sense to add series that reflect information from different sectors. One obvious candidate is data on unemployment, to reflect the state of pressure and activity in the labour market.[12] A measure of excess demand for or supply of capital is provided by survey measures of capacity utilisation. Hence, following a basic production function intuition, information on factor inputs can be used to condition the estimate of the output gap. A multivariate time series model with these data can be built using the unobserved components framework shown in section 2.

3.1 The model

In order to implement this approach, three things are required from our model. First, we will need an estimate of the unobserved common trend-reverting behaviour in the three series, including estimates of the dynamics of this temporary component. Second, the model will need to be able to provide estimates of the unobserved permanent trends of output, unemployment, and capacity utilisation. Finally, estimates of the coefficients that describe the linkages from cyclical output to unemployment and capacity utilisation are required.

Clark (1989) presents a model where output and unemployment are decomposed into permanent trend components and a common trend-reverting component (which I will refer to as "the cycle"). The dynamics of the cycle are described by a simple autoregressive process, which is assumed to be non-explosive, and the underlying trends are random walks with drift. The advantage with this approach is that detrending, estimation of the cycle, and estimation of the Okun's relationship are done simultaneously. The data are left relatively free to speak about the degree of persistence in the common cycle, the energy attributable to permanent and temporary changes to output and unemployment, and the impact of cyclical movements on unemployment levels.

---

12 Employment could also be used in this framework. However, there is more evidence of cyclical rises and falls in unemployment data than employment.

I utilise Clark's approach, with some necessary extensions. First, observed output is decomposed into three unobserved components:

$$y_t = \mu_t^y + z_t + e_{1t}, \qquad \mathrm{E}(e_{1t}) = 0, \qquad \mathrm{Var}(e_{1t}) = \sigma_{e1}^2, \qquad (4a)$$

where y_t represents output, μ_t^y the permanent trend for output, z_t a temporary trend-reverting component, and e_{1t} an irregular noise component.[13]

Observed unemployment is also decomposed into permanent, temporary and irregular components, with an Okun's relationship:

$$u_t = \mu_t^u + \mathbf{D}(L)\, z_t + e_{2t}, \qquad \mathrm{E}(e_{2t}) = 0, \qquad \mathrm{Var}(e_{2t}) = \sigma_{e2}^2. \qquad (4b)$$

Here u_t represents unemployment, μ_t^u is its permanent trend, $\mathbf{D}(L)$ is a polynomial in the lag operator, and e_{2t} is an irregular component.

Capacity utilisation is also decomposed in the same way:

$$cap_t = \mu_t^{cap} + \mathbf{G}(L)\, z_t + e_{3t}, \qquad \mathrm{E}(e_{3t}) = 0, \qquad \mathrm{Var}(e_{3t}) = \sigma_{e3}^2. \qquad (4c)$$

The dynamics of the cyclical component are simply described by an AR(p) process:

$$z_t = \mathbf{C}(L) z_t + \varepsilon_t, \qquad \mathrm{E}(\varepsilon_t) = 0, \qquad \mathrm{Var}(\varepsilon_t) = \sigma_\varepsilon^2, \qquad (4d)$$

where $\mathbf{C}(L)$ is a polynomial in the lag operator and ε_t is the innovation to the process.[14]

Hence this is a model where the cycle is common to the equations for all three observed variables. The underlying trends for output, unemployment and capacity are independent. They are all modelled as local linear trends, rather than random walks with constant drift.[15] In the case of output

$$\mu_t^y = \mu_{t-1}^y + \beta_{t-1}^y + \eta_{1t}, \ \mathrm{E}(\eta_{1t}) = 0, \qquad \mathrm{Var}(\eta_{1t}) = \sigma_{\eta 1}^2, \qquad (4e)$$

$$\beta_t^y = \beta_{t-1}^y + \zeta_{1t}, \qquad \mathrm{E}(\zeta_{1t}) = 0, \qquad \mathrm{Var}(\zeta_{1t}) = \sigma_{\zeta 1}^2, \qquad (4f)$$

---

13 Irregular terms are omitted in Clark's bivariate model; they are included here to "soak up" noise in the data not associated with trend or trend-reverting components.

14 It is assumed that the roots of the estimated AR specification are within the unit circle; this is not imposed.

15 For series as variable as New Zealand output and unemployment, this flexibility was found to be necessary. The trend components are therefore I(2), which has the unfortunate implication that an innovation to the growth rate of output early in the sample has an increasing effect on the level of output as time goes by. It also has some implications for the use of the Kalman filter during estimation; see section 4 and appendix B.

so that the trend, μ_t^y, depends on the previous lag of μ_t^y, innovations to trend, η_{1t}, and a slope term, β_t^y, which is itself a random walk with innovation ζ_{1t}. The underlying trends for unemployment and capacity are defined in the same way.

The model can be easily re-written in its minimal state space realisation. The basic model to be estimated assumes that the Okun's relationship and the linkage from the cycle to capacity work contemporaneously and with lags one to four, and that the cycle can be parameterised by an AR(4) process.[16] Hence the observation equation can be written as:

$$\begin{bmatrix} y_t \\ u_t \\ capu_t \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ d_0 & d_1 & d_2 & d_3 & d_4 & 0 & 0 & 1 & 0 & 0 & 0 \\ g_0 & g_1 & g_2 & g_3 & g_4 & 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} z_t \\ z_{t-1} \\ z_{t-2} \\ z_{t-3} \\ z_{t-4} \\ \mu_t^y \\ \beta_t^y \\ \mu_t^u \\ \beta_t^u \\ \mu_t^{cap} \\ \beta_t^{cap} \end{bmatrix} + \begin{bmatrix} e_{1t} \\ e_{2t} \\ e_{3t} \end{bmatrix} \qquad (5a)$$

The corresponding state equation takes the form:

$$\begin{bmatrix} z_t \\ z_{t-1} \\ z_{t-2} \\ z_{t-3} \\ z_{t-4} \\ \mu_t^y \\ \beta_t^y \\ \mu_t^u \\ \beta_t^u \\ \mu_t^{cap} \\ \beta_t^{cap} \end{bmatrix} = \begin{bmatrix} \begin{matrix} c_1 & c_2 & c_3 & c_4 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \end{matrix} & \mathbf{0} \\ \mathbf{0} & \begin{matrix} 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{matrix} \end{bmatrix} \begin{bmatrix} z_{t-1} \\ z_{t-2} \\ z_{t-3} \\ z_{t-4} \\ z_{t-5} \\ \mu_{t-1}^y \\ \beta_{t-1}^y \\ \mu_{t-1}^u \\ \beta_{t-1}^u \\ \mu_{t-1}^{cap} \\ \beta_{t-1}^{cap} \end{bmatrix} + \begin{bmatrix} \varepsilon_t \\ 0 \\ 0 \\ 0 \\ 0 \\ \eta_{1t} \\ \zeta_{1t} \\ \eta_{2t} \\ \zeta_{2t} \\ \eta_{3t} \\ \zeta_{3t} \end{bmatrix} \qquad (5b)$$

[16] The specification of lags in the observation equations and in the cyclical component is arbitrary. As is common, we face a trade-off between tractability and the wish to capture salient features of the data. Since quarterly data are used, the specification begins with four lags in order to allow for annual patterns.

In this form, the model can be estimated via the Kalman filter and maximum likelihood. (appendices 1 and 2 briefly explain the Kalman filter and estimation, respectively.)

3.2 Data

The sample period over which the model is estimated is 1970:1 to 1999:3. The output data series used is log quarterly real seasonally-adjusted production-measure GDP from Statistics New Zealand. Unemployment data are from the Household Labour Force Survey of unemployment in the workforce over the age of 18.[17] Capacity utilisation is derived by the New Zealand Institute for Economic Research from a survey questionnaire of manufacturers and builders, which asks them by how much they can raise production from existing plant and machinery without raising costs.[18] The median response is used to compile the series, and is represented as a ratio between zero and one.[19]

The series for output is clearly nonstationary, owing to the simple fact of economic growth. However, it is not clear that there is any simple trend; the growth rate has varied considerably over the years. Nor is there much evidence of a cycle of regular periodicity, as seen in figure 4. We might expect data for unemployment to be stationary, if only because the data are notionally bounded. However, looking at figure 5, there does appear to be a rising trend in the unemployment series as well. That said, there do seem to be regular peaks and troughs at business cycle frequencies, indicating that this series will help condition the estimate of the output gap. The series for capacity utilisation (figure 6) also shows some clearly cyclical behaviour. The periodicity of the cycle seems to be becoming longer as time goes by. At the same time, there is some appearance that there has been an evolving trend in the series, first down, then flat, then up from the early 1990s.

---

17 While output and capacity data are non-seasonal, the official seasonally-adjusted unemployment series is only available from 1986:1. The unadjusted series was therefore deseasonalised using the Basic Structural Model with trigonometric seasonality advocated by Harvey (1989, p.47) and implemented in the package STAMP (Koopman *et al*, 1995).

18 It is not absolutely clear whether respondents interpret the question in a narrow sense, relating only to their existing levels of capital, or in wider sense that also takes into account labour resources and intermediate inputs. Silverstone (1993, p.68) reports a high correlation (0.72) between the capacity utilisation responses of firms and their net responses with respect to investment intentions in plant and machinery. I take this as some corroborative evidence that the capacity utilisation data can be used in a narrow sense as an indication of excess demand for and supply of capital.

19 Respondents select one of five percentage increases in production, and the median is used as a measure of excess capacity relative to full (100 percent) capacity. Capacity utilisation as a percentage is therefore $100/(1+\text{excess})$. For example, if the median response was 10 percent, then capacity utilisation would be recorded at 90.91 percent.

4 Results

The model was estimated using the Kalman filter and exact maximum likelihood.[20] In practice, careful attention was required for the initialisation of both the Kalman filter and the numerical optimisation; these points are described in further detail in appendix B. It was found that the unemployment equation and the capacity utilisation could be made more parsimonious without affecting the resulting trends and common cycle.[21] The final, preferred model took the form:

$$y_t = \mu_t^y + z_t + e_{1t}, \qquad \text{Var}(e_{1t}) = 2.71\text{e-}16,$$

$$\underset{}{u_t = \mu_t^u - \underset{(3.55)}{21.75}\, z_t - \underset{(4.98)}{15.18}\, z_{t-1} - \underset{(1.30)}{7.68} z_{t-2} + e_{2t}}, \qquad \text{Var}(e_{2t}) = 7.01\text{e-}04,$$

$$cap_t = \mu_t^{cap} + \underset{(6.17)}{2.45} z_{t-1} - \underset{(5.06)}{2.15}\, z_{t-2} + e_{3t}, \qquad \text{Var}(e_{3t}) = 2.04\text{e-}16,$$

$$z_t = \underset{(24.54)}{1.78} z_{t-1} - \underset{(12.95)}{0.82} z_{t-2} + \varepsilon_t, \qquad \text{Var}(\varepsilon_t) = 1.45\text{e-}05,$$

$$\mu_t^y = \mu_{t-1}^y + \beta_{t-1}^y + \eta_{1t}, \qquad \text{Var}(\eta_{1t}) = 8.92\text{e-}05,$$

$$\beta_t^y = \beta_{t-1}^y + \zeta_{1t}, \qquad \text{Var}(\zeta_{1t}) = 7.66\text{e-}07,$$

$$\mu_t^u = \mu_{t-1}^u + \beta_{t-1}^u + \eta_{2t}, \qquad \text{Var}(\eta_{2t}) = 6.03\text{e-}02,$$

$$\beta_t^u = \beta_{t-1}^u + \zeta_{2t}, \qquad \text{Var}(\zeta_{2t}) = 4.58\text{e-}04,$$

$$\mu_t^{cap} = \mu_{t-1}^{cap} + \beta_{t-1}^{cap} + \eta_{3t}, \qquad \text{Var}(\eta_{3t}) = 1.22\text{e-}05,$$

$$\beta_t^{cap} = \beta_{t-1}^{cap} + \zeta_{3t}, \qquad \text{Var}(\zeta_{3t}) = 1.40\text{e-}19.$$

The long-run unemployment equation provides an estimate of the Okun's relationship. This is slightly below 2:1, which means that a one unit increase in cyclical output is associated with a fall in cyclical unemployment of one half.[22]

---

20 The Kalman filter, Kalman smoother and likelihood evaluation procedures for estimating the model were written in GAUSS. Numerical optimisation was carried out using the BFGS algorithm available in the GAUSS module OPTIMUM.

21 This appears to be a grey area. I selected the model specification using the value of the likelihood function as a measure of model fit. By contrast, removing some parameters that were insignificant lowered model fit and produced estimates of the gap that could be discounted.

22 This ratio is a long-run value derived by summing the coefficients and dividing into one hundred to convert from a log difference to percentage units.

The coefficients in the cycle equation provide an indication of the characteristics of the estimated common cycle. Note that these alternate sign. Solving for the roots of the equation, we see that there are the complex conjugates $0.89 \pm 0.18i$. The existence of complex conjugates means that the trend-reverting component is truly cyclical – that is, the impulse response from a single shock would oscillate. The periodicity of this cycle is 32 quarters. All of the roots are within the unit circle, so that the oscillation is damped and the cyclical component is indeed temporary. At 0.91, the dominant root is quite close to the unit circle, implying that the cycles are quite persistent.

Because of the normalisation of the cycle in the model (equation 4a) and the small estimated variance for the irregular term, the estimates for the common cycle and the output gap are almost identical. Looking at the output gap in figure 7, the estimate clearly shows longer periodicities that look much more like business cycles than the estimates from the univariate case of section 2 (figure 3). Compared with the results from the HP filter, these estimates bear more resemblance to the results using $\lambda = 1600$ in the first half of the sample and $\lambda = 100{,}000$ in the second half. To the extent that the value of λ in the HP filter reflects a view on the ratio of demand (temporary) to supply (permanent) shocks, this implies very approximately that the gap from the model is more 'demand dominated' in the last 15 years.

The estimated trends for the three series are plotted in figures 8 to 10.[23] For the trends in output and unemployment, the impact of the large variances in the innovations to their levels is immediately obvious. The trends appear to evolve in discreet jumps; results are clearly not consistent with the smoothness restriction imposed by the HP filter.[24] This is a result of the structure of the model, which effectively maps a smooth cycle into output space.

The trend in unemployment rises in large steps through the sample period. It is noticeably persistent: it cuts across the peaks and troughs of the late 1970s and early 1980s and through the cycles after that. A similarly steady fall and rise is seen in the trend for capacity utilisation. These two series appear, therefore, to have been quite successful at conditioning the estimate of the output gap. Moreover, these two series have their interpretations as measures of the demand pressure for the factors of production. In this respect, the recession of 1991-1992 was the product of both excess supply of labour and excess supply of capital. However, the 'unemployment gap' is longer-lived and typically lags the pressure on capital. This interpretation needs to be

23 Note that the trends shown are the final estimates from the Kalman smoother, which revises the estimates from the Kalman filter for all information over the sample, rather than up to time *t*. See appendix B for further details.

24 The permanent trend components are not very smooth, even though the model includes irregular terms in all of the observation equations. To see whether this was the result of relatively large energy in the levels of the underlying trends, the variance terms η_1, η_2, and η_3 were restricted to zero. This did not change the results substantially – the algorithm forced enough variation into the slopes so as to leave the appearances of the estimated trends virtually unchanged. The "discrete jump" behaviour of the trends therefore appears to be an important and robust regularity over the sample. However, this is not to say that variation in the slopes is unimportant: the data firmly rejected attempts to fit trends as random walks with constant drift terms.

done with some care, but the ability of this kind of model to offer an economic story behind the estimate of the output gap would appear to be a major advantage.

Two extensions to the model were experimented with. First, in its present form, the model is over-identified, as all off-diagonal elements of the covariance matrix for the observation equation, **Q**, are constrained to zero (see appendix A). Some of these elements could, however, be expected to be non-zero.[25] However, experimentation with relaxing this over-identification did not advance the model.[26] Second, in this framework it is possible to allow time-varying variances.[27] Indeed, the local linear trend model of section 2, augmented with a simple break in variances, when re-estimated over the output series showed a clear reduction in the levels variance from 1985:2 onwards. This is consistent with the story of a shift to a larger component of temporary shocks later in the cycle. However, it proved impossible to introduce the same approach to the larger model.[28]

5 Policy and uncertainty

While the model has been successful in identifying a common trend at business cycle frequencies, the natural question raised is how precise these estimates are. For example, Staiger, Stock and Watson (1996) find that estimates of the natural rate of unemployment for US data are highly imprecise. In this section, I examine the uncertainty associated with the estimated trends and common cycle.

5.1 Confidence intervals around the estimates of trend and cycle

The Kalman recursion automatically returns estimates of the error covariance matrix associated with the estimates of the state variables, $\mathbf{P}_t$. (See appendix B for further details.) The diagonal elements of these matrices may be used to construct confidence intervals for the corresponding elements of the state vector. Figure 11 shows the plot of the 1 standard deviation level of uncertainty around the output gap (as a percentage of the gap) that comes from $\mathbf{P}_t$. As is quickly seen, the level of uncertainty is time varying, ranging in level from 1 or 2 percent to nearly 100 percent. However, comparing this plot to the cycle in figure 9, the level of uncertainty appears not to be associated with the turning points in the cycle, as one might expect. Instead, it

---

25 For example, structural changes to the economy over the sample period could mean that shocks to trend capacity are not orthogonal to shocks to trend unemployment. Alternatively, hysteresis models could allow correlation between shocks to the common cycle and shocks to the trends.

26 Introducing correlations between the shocks to the trends consistently produced severe convergence problems. Alternatively, a very wide range of values, negative and positive, were consistent with values for covariances between cycle and trend innovations.

27 The model is still Gaussian, and the construction of a likelihood function still proceeds via the prediction error decomposition, the only difference being that there are now more parameters to estimate. See Harvey (1989), pp.346-348.

28 Convergence problems were experienced.

appears to be more associated with the "cross-over" points where output is close to the trend.

As Ansley and Kohn (1986) and Hamilton (1986) show, estimates of uncertainty from $\mathbf{P}_t$ systematically underestimate the true standard errors around the state vector.[29] Accordingly, I estimated confidence intervals via the Monte Carlo integration method used in Hamilton (1985). (The method is discussed in detail in appendix C.) The 95 percent confidence intervals around the trend are huge (so large as to be uninteresting), and imply that the policymaker could never be sure that the economy had been in a boom or a recession.

5.2 End-point revisions

Because we are particularly interested in the values for the state variables of the model, the Kalman smoother has been used to derive final values. This allows all the observations in the sample to be used to infer the likely values. Like all smoothers, however, the Kalman smoother becomes more and more one-sided at the end of the sample. At this point, the smoothed value and the filtered value are the same. Revisions are therefore continually made for the smoothed estimates of the gaps as more and more observations were added to the sample. The key question, therefore, is what is the nature of these revisions, and, in particular, how big are they?

To give some idea of the nature of the revisions, figure 12 plots the revisions to the output gap (that is, the difference between the filtered or one-sided value and the final smoothed or two-sided value at each point in time.) The revisions for all of the estimated gaps are very similar: they are highly serially correlated; moreover, the scale of revisions is very large, the differences often being as much as a recession or boom itself.

7 Concluding remarks

One of the main motivations for the multivariate unobserved components approach described in this paper was a simple question: if we let the data speak, what will they tell us? By introducing some economic structure to the identification of a common cyclical component to observed data, a further advantage of the model was that an interesting factor input interpretation could be placed on the results.

Such an approach asks a great deal from the data. As it turns out, the data take a while to tell us their story. In this respect, the model does not escape from the end-point problem which affects other methods of estimating potential. The extent and persistence of the revisions from filtered estimates to smoothed estimates is large and, from a policymaker's perspective, disconcerting. The levels of the cycle and trends are estimated imprecisely.

---

[29] The reason is that the estimate $\mathbf{P}_t$ is conditional on the assumption that the parameters of the model are known. Of course, here these parameters are also being estimated, and this brings its own uncertainty to the estimate of the state vector.

Nonetheless, the model proved quite effective in estimating plausible coefficients for the linkage of the common cycle to unemployment and capacity utilisation measures, while successfully conditioning the estimate of the output gap to business cycle frequencies. It was noticeable that the resulting estimates are not matched by any one particular simple time-series smoother, and, compared to different versions of HP-filtered gaps, indicated some shift towards a more ‘demand-led’ cycle. The structure of the model allows for some interesting interpretations of the composition of the gap according to factor inputs. Hence, while some considerable care in the interpretation and use of the estimates is required, the results indicate that there are important benefits from pursuing structural, data-driven approaches such as this.

References

Ansley, Craig and R Kohn (1986), "Prediction mean square error for state space models with estimated parameters" *Biometrika* 73, 467-74.

Canova, Fabio (1998), "Detrending and business cycle facts" *Journal of Monetary Economics* 41, 475-512.

Clark, Peter (1989), "Trend reversion in real output and unemployment" *Journal of Econometrics* 40, 15-32.

Coe, David T and C John McDermott (1997), "Does the gap model work in Asia?" International Monetary Fund *Staff Papers* 44(1), 59-80.

Davidson, R and J G MacKinnon (1993), *Estimation and Inference in Econometrics*, Oxford University Press.

Hamilton, James (1985), "Uncovering financial market expectations of inflation" *Journal of Political Economy* 93, 1224-41.

Hamilton, James (1986), "A standard error for the estimated state vector of a state space model" *Journal of Econometrics* 33, 387-97.

Hamilton, James (1992), "Was the deflation during the great depression anticipated? Evidence from the commodity futures market" *American Economic Review* 82(1), 157-78.

Hamilton, James (1994), *Time Series Analysis*, Princeton: Princeton University Press.

Harvey, Andrew C (1989), *Forecasting, Structural Time Series Models and the Kalman Filter*, Cambridge University Press.

Harvey, Andrew C and Albert Jaeger (1993), "Detrending, stylised facts and the business cycle" *Journal of Applied Econometrics* 8, 231-247.

Hodrick, R and E Prescott (1997), "Post-war US business cycles: an empirical investigation" *Journal of Money, Credit and Banking* 29(1), 1-16.

Koopman, Siem J, Andrew C. Harvey, Jurgen A. Doornik, and Neil Shephard (1995), *STAMP 5.0 Structural Time Series Analyser, Modeller and Predictor* London: Chapman and Hall.

Kuttner, Kenneth (1992), "Monetary policy with uncertain estimates of potential output" Federal Reserve Bank of Chicago *Economic Perspectives* Vol.16(1) 2-15.

Meinhold, R J and N D Singpurwalla (1983), "Understanding the Kalman filter" *American Statistician* 37, 123-127.

Okun, Arthur (1962), "Potential GNP: its measurement and significance" in *Proceedings of the Business and Economics Section*, American Statistical Association, Washington, DC, 98-104.

Pagan, Adrian (1997a), "Towards an understanding of some business cycle characteristics" *Australian Economic Review* 30, 1-15.

Pagan, Adrian (1997a), "Policy, theory and the cycle" *Oxford Review of Economic Policy* 13, 19-23.

Silverstone, Brian (1993), "Some practical aspects of the NZIER measure of capacity utilisation" New Zealand Institute of Economic Research *Quarterly Predictions* 113.

Staiger, Douglas, Stock, James, and Mark Watson (1996), "How precise are estimates of the natural rate of unemployment?" *NBER Working Paper No*. 5477.

Watson, Mark (1986), "Univariate detrending methods with stochastic trends" *Journal of Monetary Economics* 18, 49-75.

Appendix A:
State space representations of unobserved components models

This appendix briefly describes state space representations of unobserved components models. More comprehensive descriptions are available in Harvey (1989) and Hamilton (1994).[30]

The first part of the state space representation is an *observation* or *measurement equation*, which describes a linear relationship between a vector of observed variables and a vector of variables that describe the *state of nature* of the system. This takes the form:

$$x_t = \mathbf{Z}_t \alpha_t + \mathbf{d}_t + \varepsilon_t , \qquad \mathrm{E}(\varepsilon_t) = 0 , \qquad \mathrm{Var}(\varepsilon_t) = \mathbf{H}_t , \qquad \text{(A.1)}$$

where x_t is the $(n \times 1)$ vector of observed variables, α_t is the $(p \times 1)$ state vector of unobserved variables, and $\mathbf{d}_t$ is a vector of predetermined variables. The matrix of coefficients $\mathbf{Z}_t$ - the *observation matrix* - is correspondingly $(n \times p)$. The observation error is assumed to be distributed with conditional expectation zero and a covariance matrix $\mathbf{H}_t$.

The essential difference between state space models and conventional linear models is that in the former the state of nature is not assumed to be constant but can change with time. The state space representation is completed by adding another equation, the *state* or *transition* equation. This describes the law of motion for the unobserved components:

$$\alpha_t = \mathbf{T}_t \alpha_{t-1} + \mathbf{c}_t + \mathbf{R}_t \eta_t \qquad \mathrm{E}(\eta_t) = 0 , \qquad \mathrm{Var}(\eta_t) = \mathbf{Q}_t \qquad \text{(A.2)}$$

As can be seen, this is assumed to be a first order autoregressive process. Exogenous components can be added in $\mathbf{c}_t$. The *transition* matrix $\mathbf{T}_t$ is $(p \times p)$. $\mathbf{R}_t$ is a $(g \times p)$ matrix factorising the innovations in the state vector. As denoted by the time subscripts, these matrices can be time-varying, though techniques to handle this are generally complicated and expensive.

The dimension of α_t is independent of x_t - for example, x_t could be a scalar and α_t a vector, implying that x_t was a function of several unobserved components; or alternatively x_t could be a vector and α_t a scalar, implying that the elements of x_t were a function of a single state of nature or a common factor.

---

[30] The description here and in appendix B uses the notation in Harvey (1989).

Appendix B:
Estimation of the model via the Kalman filter and exact maximum likelihood in the time domain

The obvious problem is how to estimate a model in the form described in appendix A. There are two sets of unknowns: the parameters of the model in $\mathbf{T}_t$, $\mathbf{Z}_t$, $\mathbf{d}_t$, $\mathbf{c}_t$, $\mathbf{H}_t$, $\mathbf{R}_t$ and $\mathbf{Q}_t$; and the elements of the state vector α_t. However, once a model is cast into its minimal state space representation, the addition of certain assumptions allows the model to be estimated using maximum likelihood and the Kalman filter.

The Kalman filter is simply a recursion which we can use to make inference about the state of nature α_t. The Kalman filter can be derived using classical statistical techniques, but is probably most easily interpreted as an application of Bayes' theorem:[31]

$$\begin{aligned}&\text{Prob}\left\{\text{ state of nature} \mid \text{data }\right\}\\ &\qquad \propto \text{Prob}\left\{\text{ data} \mid \text{state of nature }\right\} \times \text{Prob}\left\{\text{ state of nature }\right\}\end{aligned}$$

or

$$P\left(\alpha_t \middle| x_t\right) \propto P\left(x_t \middle| \alpha_t, x_{t-1}\right) \times P\left(\alpha_t \middle| x_{t-1}\right)$$

The recursion can probably be best explained if we focus on a certain time point. We assume for the moment that the parameters of the model in $\mathbf{T}_t$, $\mathbf{Z}_t$, $\mathbf{d}_t$, $\mathbf{c}_t$, $\mathbf{H}_t$, $\mathbf{R}_t$ and $\mathbf{Q}_t$ are known and that at time *t* and we have information up to and including the observations x_{t-1}.[32] Let $\mathbf{a}_{t-1}$ denote the optimal estimator of α_{t-1} based on this information, and $\mathbf{P}_{t-1}$ denote the covariance matrix of the estimation error.

Because this is a recursion, we need some starting values for $\mathbf{a}_0$ and $\mathbf{P}_0$. We supply best guesses for these matrices.[33] We look forward to time *t*, but in two stages: prior to observing x_t, and then after observing x_t.

---

31 For a more complete development of this intuition, see Meinhold and Singpurwhalla (1983).

32 We will also assume that the errors are Gaussian.

33 One approach is to assume that, where applicable, the underlying state at time 0 is equal to the observed value at time 1. This was used in the estimation of the models in this paper. Hence the underlying trend at time 0 was assumed to be equal to the observed value of the series at time 1, whereas the cyclical elements were assumed to be zero. A common approach to the specification of $\mathbf{P}_0$ is to assume that it is $\mathbf{I}_p \times k$, where *k* denotes a measure of imprecision about the value, say ten or twenty. Stationary elements of the state vector, however, should have corresponding elements in $\mathbf{P}_0$ that are small. In model (4), a value of 0.5 was used. Small values were also used for the non-stationary trend components. This was used to tie the estimated trend down to a path that would run through the data. This seems to be necessary due to the high volatility of the data in the first observations of the samples in conjunction with the I(2) trends used – small changes in the growth path early in the recursion could easily cause the estimates of trend to diverge away from the observed data.

The first step is to make forecasts one period ahead of the state vector and the estimation error covariance. The best forecast of the state vector is easy to see, since the model is linear and the conditional expectation of the innovations is zero:

$$\mathbf{a}_{t|t-1} = \mathbf{T}_t \mathbf{a}_{t-1} + \mathbf{c}_t . \tag{B.1}$$

The covariance matrix of the innovation error updates by the equation

$$\mathbf{P}_{t|t-1} = \mathbf{T}_t \mathbf{P}_{t-1} \mathbf{T}_t' + \mathbf{R}_t \mathbf{Q}_t \mathbf{R}_t' . \tag{B.2}$$

Equations (B.1) and (B.2) are equivalent to the probability statement

$$\left(\alpha_t \middle| x_{t-1}\right) \sim N\left(\mathbf{T}_t \hat{\alpha}_t + \mathbf{d}_t , \mathbf{T}_t \mathbf{P}_{t-1} \mathbf{T}_t' + \mathbf{R}_t \mathbf{Q}_t \mathbf{R}_t'\right),$$

which can be thought of as the prior, $P\left(\alpha_t \middle| x_{t-1}\right)$, for the next stage. On observing x_t, the goal is to compute the posterior of α_t. To do this, however, we now have to evaluate the likelihood $L\left(\alpha_t \middle| x_t\right)$ or $P\left(x_t \middle| \alpha_t , x_{t-1}\right)$.

Since $\mathbf{T}_t$, $\mathbf{Z}_t$, $\mathbf{d}_t$, $\mathbf{c}_t$, $\mathbf{H}_t$, $\mathbf{R}_t$, $\mathbf{Q}_t$, and $\mathbf{a}_{t|t-1}$ are all known, observing x_t is equivalent to observing $e_t = x_t - \hat{x}_t$. Hence when the new observation becomes available, the estimates at time *t*-1 – the priors – can be updated:

$$\mathbf{a}_t = \mathbf{a}_{t|t-1} + \mathbf{P}_{t|t-1} \mathbf{Z}_t' \mathbf{F}_t^{-1} \left(x_t - \mathbf{Z}_t \mathbf{a}_{t|t-1} - \mathbf{d}_t \right), \tag{B.3}$$

where

$$\mathbf{F}_t = \mathbf{Z}_t \mathbf{P}_{t|t-1} \mathbf{Z}_t' + \mathbf{H}_t \tag{B.4}$$

Here a certain fraction of the difference between the observable and the predicted state is added to the previous prediction.[34] Finally, the estimated covariance matrix of the innovation error needs to be updated:

$$\mathbf{P}_t = \mathbf{P}_{t|t-1} - \mathbf{P}_{t|t-1} \mathbf{Z}_t' \mathbf{F}_t^{-1} \mathbf{Z}_t \mathbf{P}_{t|t-1} \tag{B.5}$$

Once equations (B.3-B.5) are computed, we can go back to the next cycle of the recursion.

Equations (B.1-B.2) and (B.3-B.5) are termed the *prediction equations* and the *updating equations* respectively. Taken together and used recursively, they make up

---

[34] This further helps to explain why tight priors, $\mathbf{P}_0$, for the Kalman filter were required. The correction factor – the Kalman gain – was insufficiently large in the early stages of the recursion to prevent the estimate of trend being fatally "deflected" by large deviations of new observations from previous observations.

the Kalman filter. Hence if $\mathbf{T}_t$, $\mathbf{Z}_t$, $\mathbf{d}_t$, $\mathbf{c}_t$, $\mathbf{H}_t$, $\mathbf{R}_t$ and $\mathbf{Q}_t$ were known, the Kalman filter would yield a series, $\{a_t\}$, that was the MMSE estimate of the state vector. Thus one way of thinking about the Kalman filter is to view it as a sequential updating procedure that consists of forming a prior guess about the state of nature and then adding a correction to this guess, the correction being determined by how well the guess has performed in predicting the next observation.

Of course, we usually do not know the parameters, so we are only half way to estimating a state space model. The Kalman filter yields an extra piece of information that allows us to go forward. The *prediction error* is simply $e_t = x_t - \hat{x}_t$, or the difference between the best estimate of the state given information up to t-1 and the outturn at t:

$$v_t = x_t - \mathbf{Z}_t \mathbf{a}_{t|t-1} - \mathbf{d}_t . \qquad \text{(B.6)}$$

Equation (B.6) can be used to form a *prediction error decomposition* of the likelihood function, so that for a Gaussian model the likelihood function has the form

$$\log L = -\frac{NT}{2}\log 2\pi - \frac{1}{2}\sum_{t=d+1}^{T}\log\left|\mathbf{F}_t\right| - \frac{1}{2}\sum_{t=d+1}^{T} v_t' \mathbf{F}_t^{-1} v_t . \qquad \text{(B.7)}$$

Note that the likelihood is evaluated from $t = d+1$, rather than 1, where d equals the number of nonstationary elements in the state vector. This is to allow the estimated error covariance, $\mathbf{P}_t$, to settle to its steady-state value, $\overline{\mathbf{P}}$, in the recursion before parameter inference begins.

With the likelihood function defined, the parameters may be estimated using conventional numerical recipes. In this way the Kalman filter can be used with exact maximum likelihood in the time domain to estimate both the parameters of the model and the state vector, α_t.[35] As for all numerical procedures, attention has to be paid to starting values in order to avoid local minima.[36]

One final note: the estimates of the state vector that are yielded from the equations (B.1-B.5) are *filtered* estimates. They are one-sided in that only current and lagged values of x_t are used to inform the estimate of α_t. A further step, often used, is to pass the estimates through a Kalman *smoother*. This yields two-sided estimates of α_t.[37]

---

35 Estimation can also usually be carried out in the frequency domain if the model can be re-written into its spectral generating function. Frequency domain estimation is usually faster than time domain estimation, though time domain estimation is usually regarded as more robust when the numerical search is difficult.

36 Model (4) proved to be quite sensitive to the choice of starting values. However, the alternative estimates from different starting points could always be disregarded as economically nonsensical.

37 Applications in economics can usually use the fixed interval smoother (see Harvey, section 3.6.2), which has the advantage that it can be run after the final estimates of the parameters are derived.

Appendix C:
Estimation of confidence intervals around elements of the state vector

This appendix describes the process of evaluating signal-noise and parameter uncertainty proposed by Hamilton (1986) and utilised in Hamilton (1985), Hamilton (1992), and Kuttner (1992).

Step 1:
Obtain the variance-covariance matrix, $\hat{\Sigma}$, of the maximum likelihood estimator of the parameter vector, $\hat{\theta}$.

Step 2:
Form *n* draws of parameter vectors, θ_i, drawn from a normal distribution with a mean of $\hat{\theta}$ and variances from the diagonal of $\hat{\Sigma}$.[38] This yields a sample, $\left(\theta_1, \theta_2, \ldots, \theta_n\right)$. In practice, even with a large number of draws, the mean of this sample will be quite different from $\hat{\theta}$. This is because the mean of the random numbers generated will be sufficiently different from zero as to cause problems. Re-centering techniques such as de-meaning are to be avoided, since they alter the second moment properties of the distribution, which we also have to replicate in this procedure. Antithetic variate techniques are useful in these situations.[39] In this case, this is simply matter of drawing sequentially from a matrix of error terms of dimension $nobs \times n/2$ to form $\left(\theta_1, \theta_2, \ldots, \theta_{n/2}\right)$, and then repeating the process from the same matrix of error terms with the sign reversed to form $\left(\theta_{n/2+1}, \theta_{n/2+2}, \ldots, \theta_n\right)$.

Step 3:
For each θ_i, run through the observations, $\{y_t\}$, with the Kalman filter and the Kalman smoother to get $\hat{\mathbf{a}}_{t|t}\left(y_T, \theta_i\right)$, $\hat{\mathbf{a}}_{t|T}\left(y_T, \theta_i\right)$, $\hat{\mathbf{P}}_{t|t}\left(y_T, \theta_i\right)$, and $\hat{\mathbf{P}}_{t|T}\left(y_T, \theta_i\right)$.

Step 4:
Average $\hat{\mathbf{P}}_t\left(y_T, \theta_i\right)$ across draws for each *t*. This yields the signal-noise uncertainty at each point of time.

Step 5:
Evaluate the average variance of the deviation of the (new) value of the state from its original estimate, $\frac{1}{n}\sum_{i=1}^{n}\left\{\left[\hat{\mathbf{a}}_t\left(y_T, \hat{\theta}\right) - \hat{\mathbf{a}}_t\left(y_T, \theta_i\right)\right] \cdot \left[\hat{\mathbf{a}}_t\left(y_T, \hat{\theta}\right) - \hat{\mathbf{a}}_t\left(y_T, \theta_i\right)\right]'\right\}$.

Step 6:
Form confidence intervals from the square root of the sum of the two types of uncertainty.

---

38 A further development would be to utilise the covariances in forming the random numbers. In practice, the off-diagonal elements of $\hat{\Sigma}$ are usually very small.

39 See Davidson and MacKinnon (1993), pp.744-47.

Figure 1
Output and HP-filtered trends

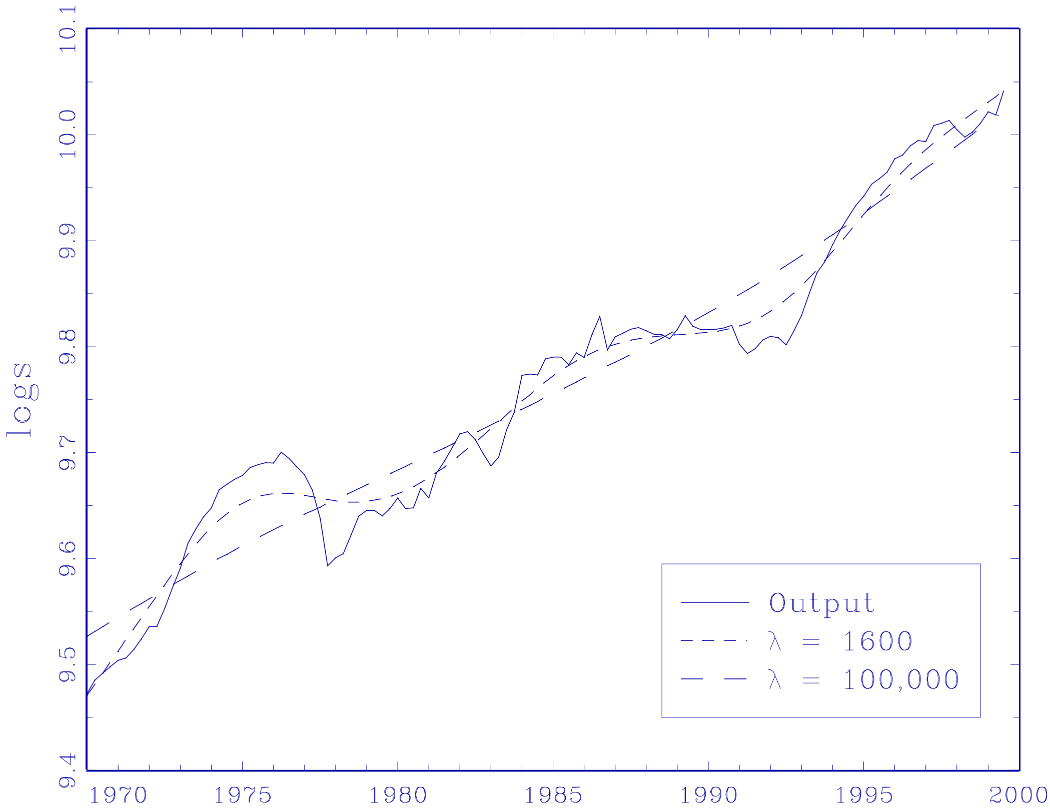

Figure 2
HP-filtered output gaps

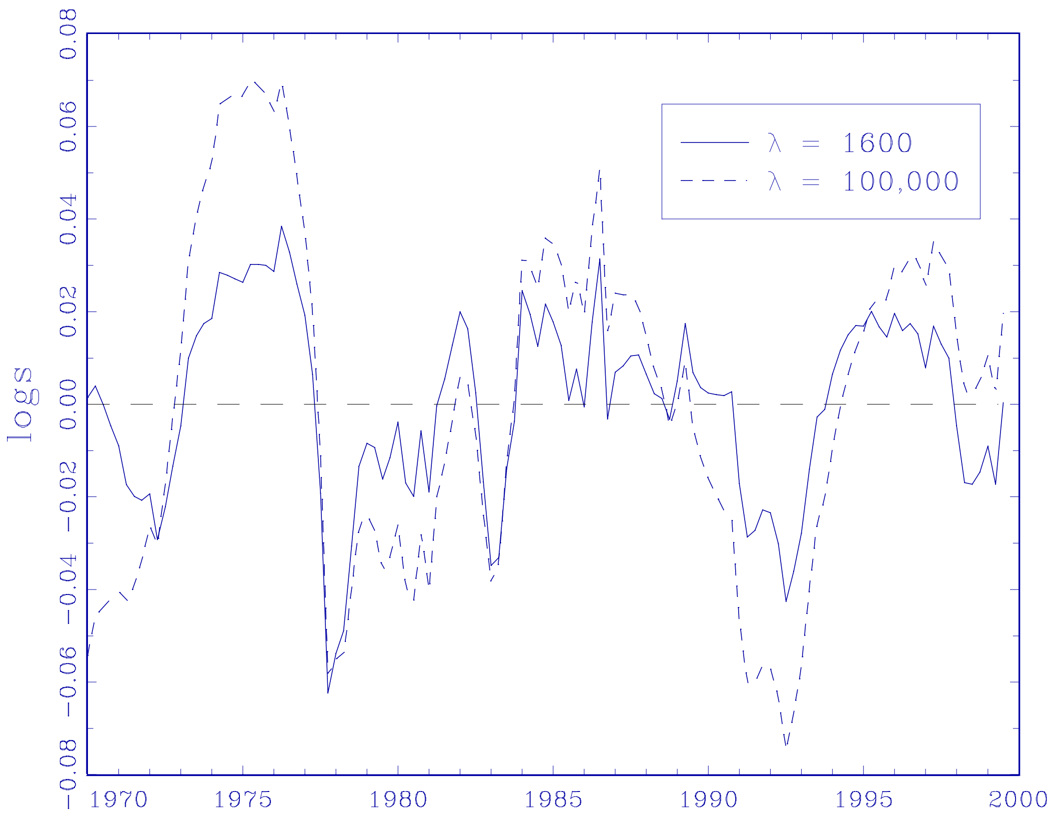

Figure 3
Output gap from trend plus cycle model

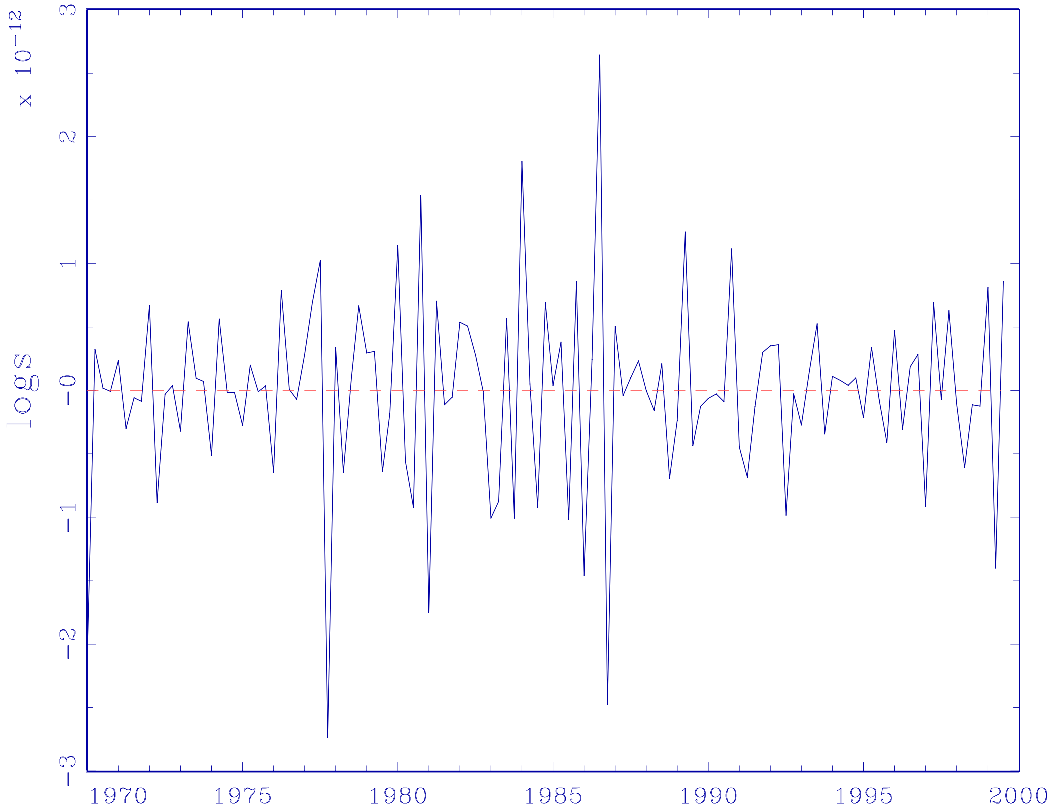

Figure 4
Output

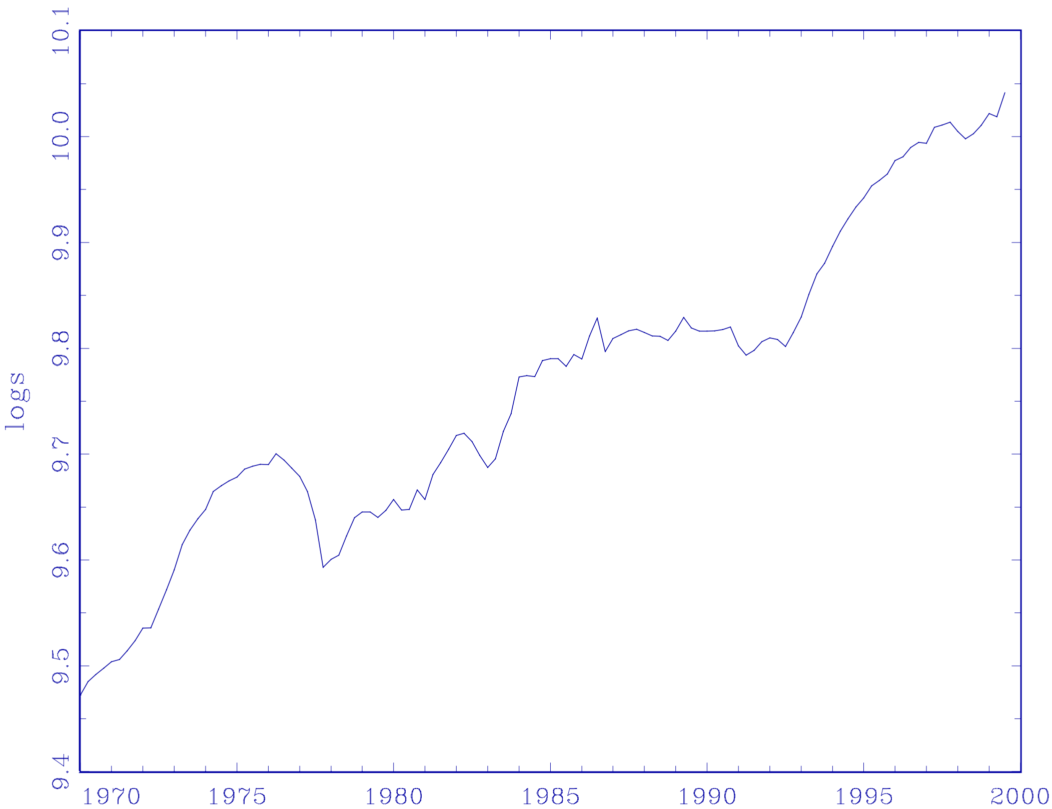

Figure 5
Unemployment

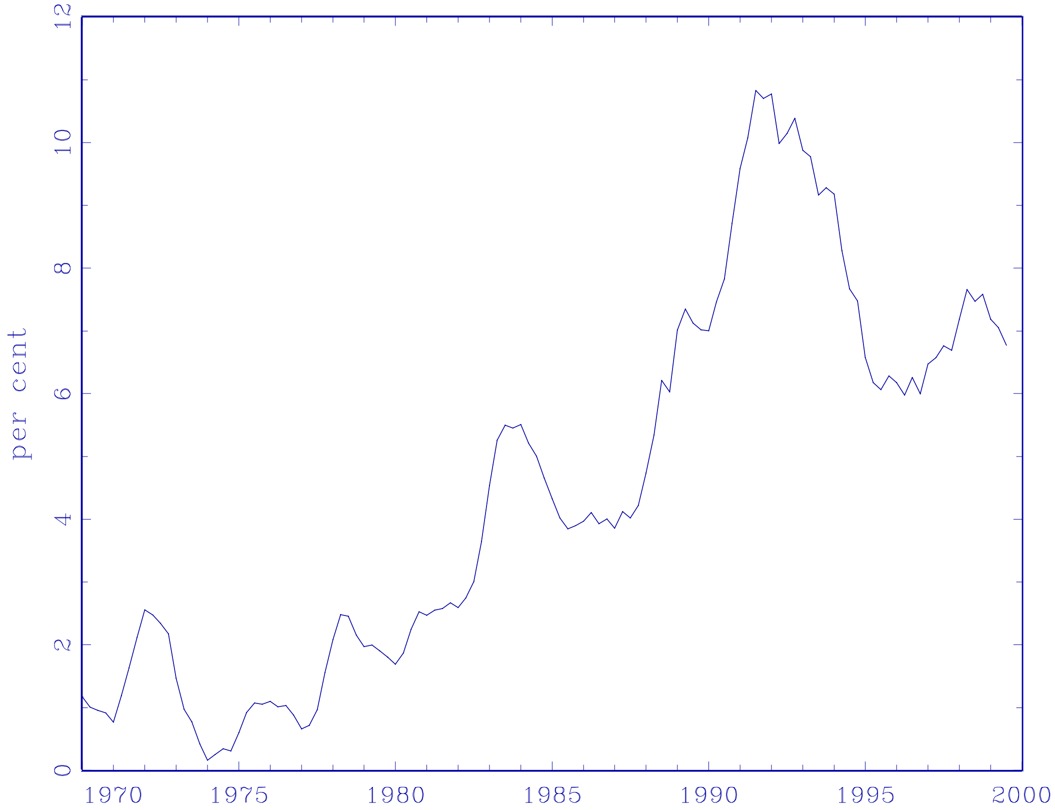

Figure 6
Capacity utilisation

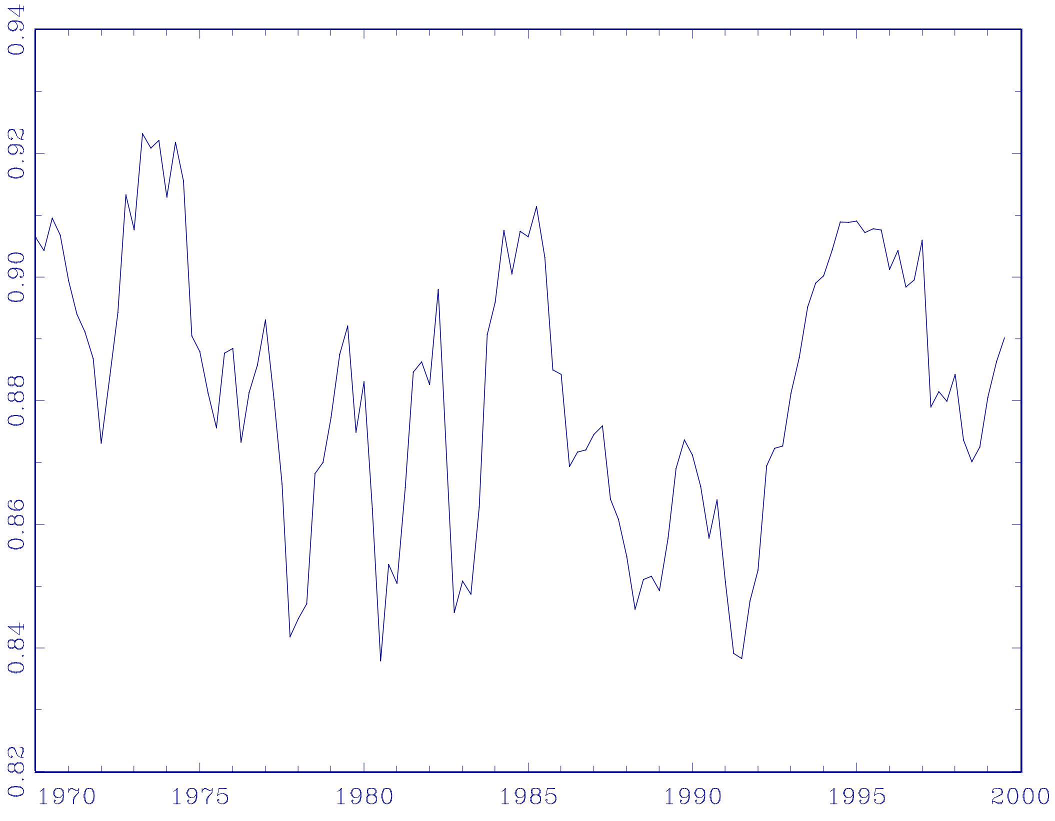

Figure 7
Output gaps compared

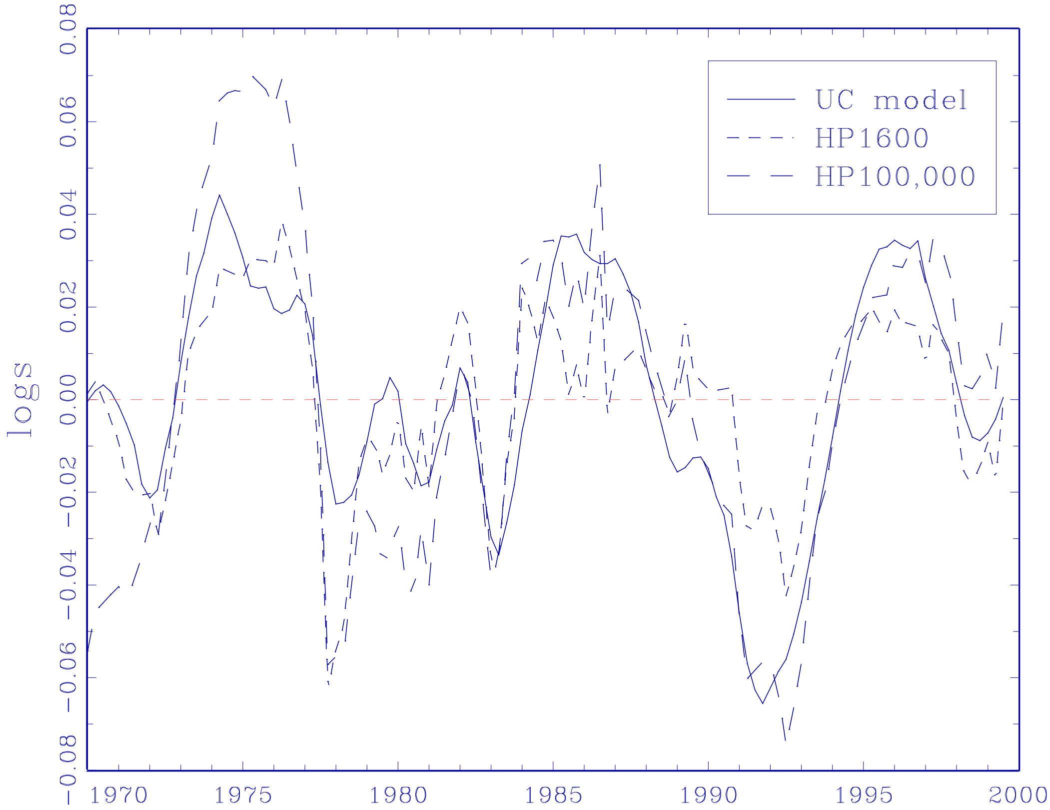

Figure 8
Output and smoothed trend

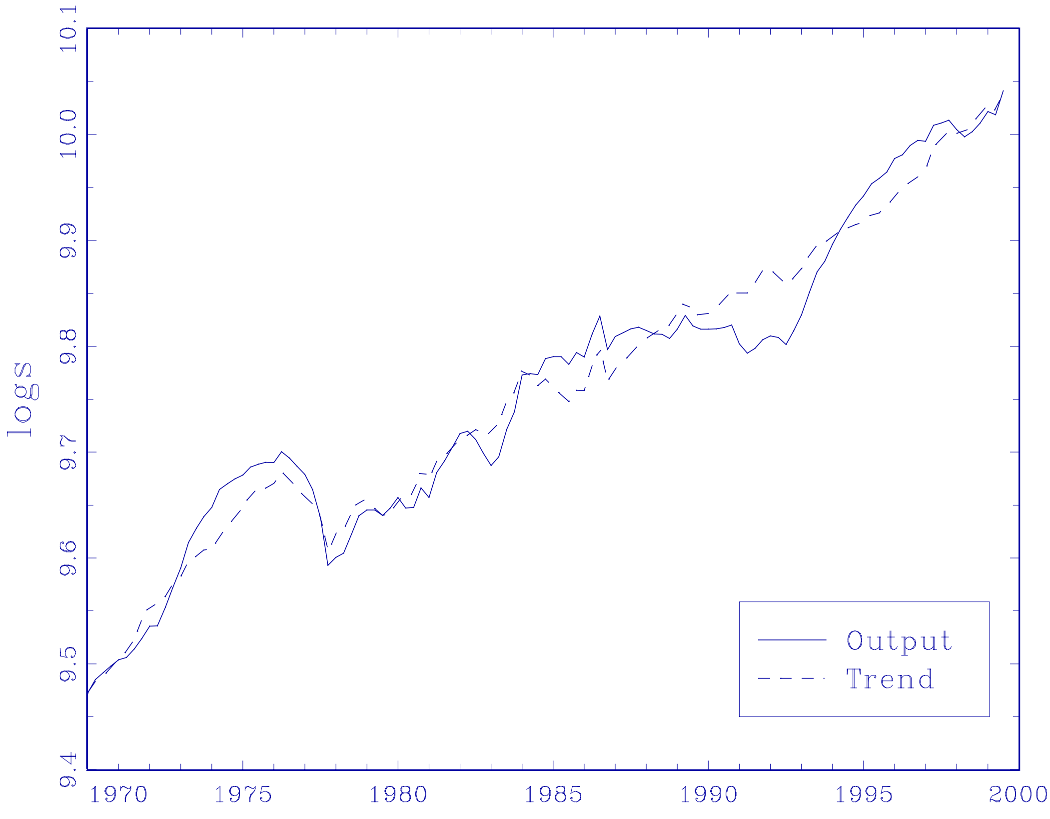

Figure 9
Unemployment and smoothed trend

Figure 10
Capacity utilisation and smoothed trend

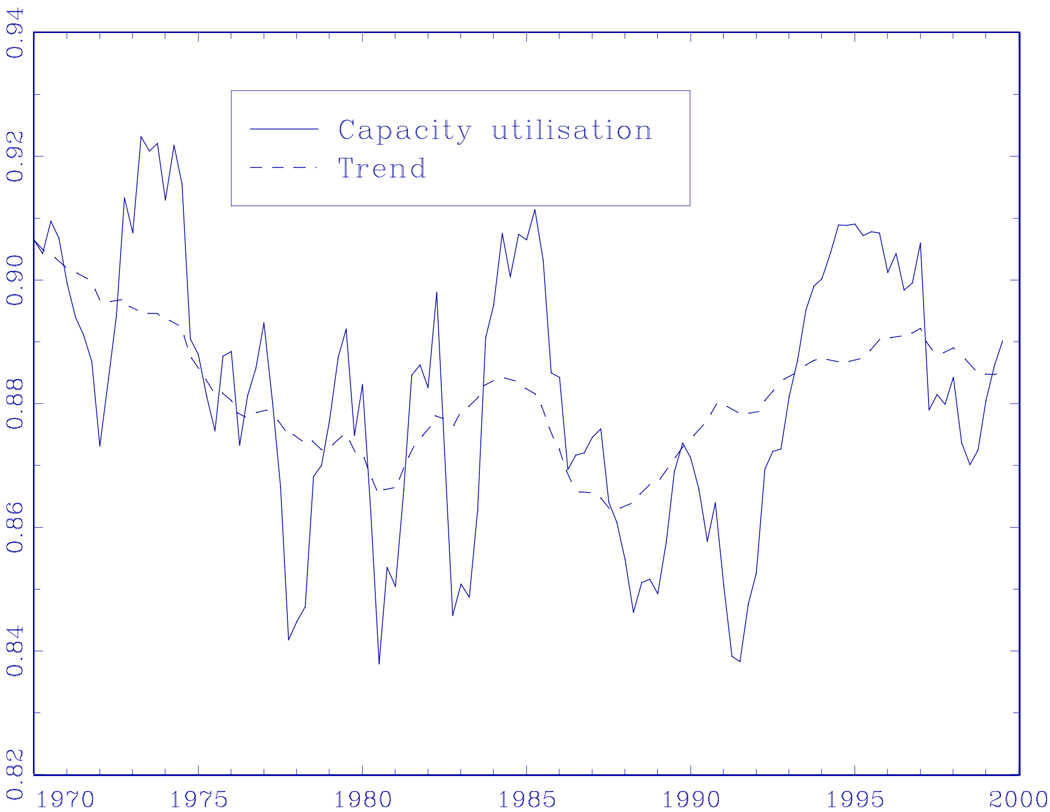

Figure 11
Level of uncertainty around output gap

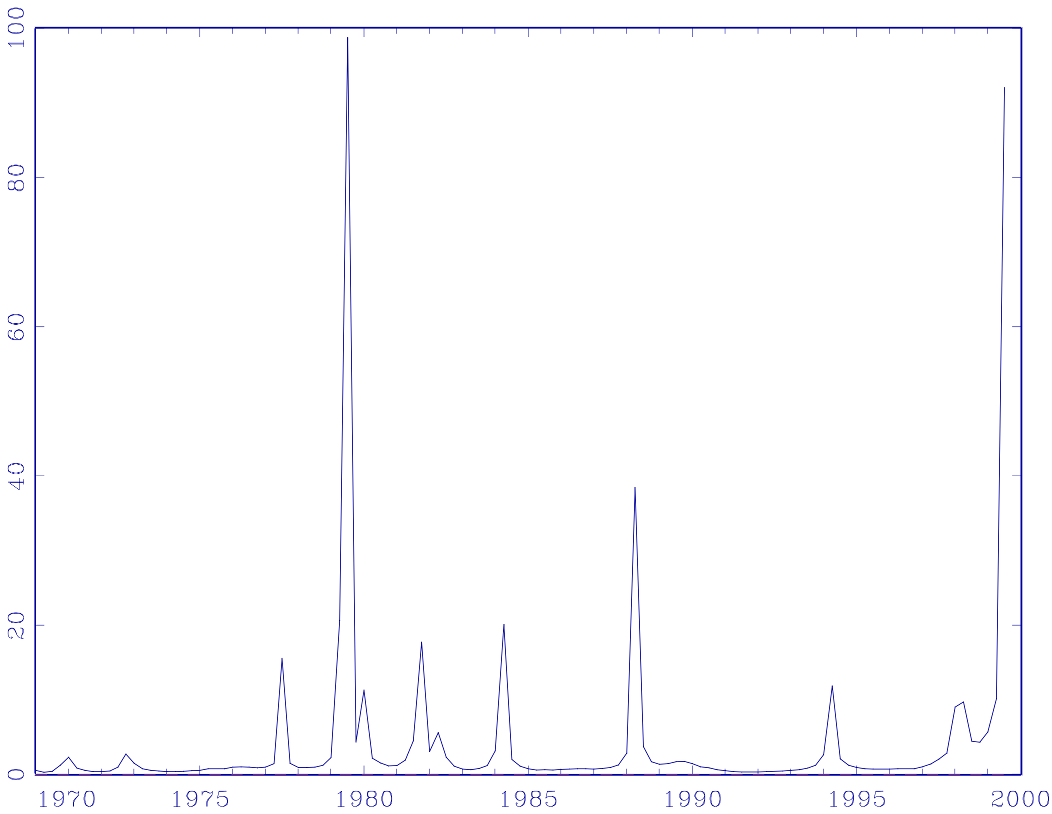

Figure 12
Smoother revision in output gap