



BIROn - Birkbeck Institutional Research Online

Garratt, Anthony and Lee, K. and Mise, E. and Shields, K. (2008) Real-time representations of the output gap. *Review of Economics & Statistics* 90 (4), pp. 792-904. ISSN 0034-6535.

Downloaded from: <https://eprints.bbk.ac.uk/id/eprint/1939/>

Usage Guidelines:

Please refer to usage guidelines at <https://eprints.bbk.ac.uk/policies.html>
contact lib-eprints@bbk.ac.uk.

or alternatively

**Birkbeck ePrints: an open access repository of the
research output of Birkbeck College**

<http://eprints.bbk.ac.uk>

Garratt, A.; Lee, K.; Mise, E. and Shields, K.
Real-time representations of the output gap
Review of Economics & Statistics - 90(4), pp.792-
904 (2008)

This is an exact copy of an article published in *Review of Economics & Statistics* (ISSN: 0034-6535) made available here with kind permission of: © 2008 Elsevier. All rights reserved.

All articles available through Birkbeck ePrints are protected by intellectual property law, including copyright law. Any use made of the contents should comply with the relevant law.

Citation for this version:

Garratt, A.; Lee, K.; Mise, E. and Shields, K.
Real-time representations of the output gap
London: Birkbeck ePrints. Available at: <http://eprints.bbk.ac.uk/1939>

Citation for publisher's version:

Garratt, A.; Lee, K.; Mise, E. and Shields, K.
Real-time representations of the output gap
Review of Economics & Statistics - 90(4), pp.792-904 (2008)

Real Time Representations of the Output Gap*

by

Anthony Garratt,[†] Kevin Lee,^{††} Emi Mise^{††} and Kalvinder Shields^{†††}

Abstract

Methods are described for the appropriate use of data obtained and analysed in real time to represent the output gap. The methods employ cointegrating VAR techniques to model real time measures and realisations of output series jointly. The model is used to mitigate the impact of data revisions; to generate appropriate forecasts that can deliver economically-meaningful output trends and that can take into account the end-of-sample problems encountered in measuring these trends; and to calculate probability forecasts that convey in a clear way the uncertainties associated with the gap measures. The methods are applied to data for the US 1965q4-2004q4 and the improvements over standard methods are illustrated.

Keywords: Output gap measurement, real time data, data revision, probability forecasts.

JEL Classification: E52, E58.

*[†]Birkbeck College, London, UK, ^{††}University of Leicester, UK, ^{†††}University of Melbourne, Australia. We have received helpful comments from two referees, Simon van Norden, Adrian Pagan and participants at the CIRANO and Bank of Canada's Workshop on 'Macroeconomic Forecasting, Analysis and Policy with Data Revisions'. Corresponding author: Kalvinder K. Shields, Department of Economics, University of Melbourne, Victoria, 3010, Australia. E-mail: k.shields@unimelb.edu.au, tel: 00 613-83443549, fax: 00 613-83446899.

1 Introduction

The measurement of the output gap, i.e. the difference between the economy's actual output and its potential or trend level, is central to much applied macroeconometric work and particularly the analysis of monetary policy. However, it is widely recognised that the output gap is measured with considerable uncertainty, and this is especially true for the measures considered in real-time decision-making.¹ For example, Orphanides and van Norden (2002) [OvN] show, using US data, that the standard measures of this central concept are extremely unreliable, with ex post revisions of the gap in the US of the same order of magnitude as the estimated gap itself. Much of the unreliability arises because the gap measures are based on output data which is subsequently revised and on measures of the trend output level which are subject to estimation error. OvN decompose the revisions observed in their output gap measures into two parts reflecting these two sources of change. They show that, for their data, the effect of changes in the measurement of the trend exceed the effects of changes in the published data but that both effects are significant.²

The OvN analysis highlights the problems involved in real time decision-making by illustrating how their gap measure changes as new information on the actual and trend output levels becomes available with the release of each new vintage of data. However, the OvN decomposition is based on a recursive analysis of each successive vintage of data taken in turn. This ignores the possibility that the sequence of vintages released over time may in itself contain useful information with which to interpret the most recent vintage of data and to anticipate future outcomes (as discussed in Howery, 1978). Hence, for example, there might be systematic patterns in the data revisions that can be used, in

¹It is also acknowledged that the use of ex post revised data can yield misleading descriptions of historical policy and that the use of real-time data generates different real-time policy recommendations to those obtained on the basis of ex post revised data (see, for example, Rotemberg and Woodford (1999), Brunner (2000), Orphanides *et al.* (2000), Orphanides (2001), and Amato and Swanson (2001)).

²These differences are potentially extremely important given the reliance of recent empirical work on the identification of monetary policy shocks and impulse responses on assumptions on the ordering of decisions and the timing of the release of information. See Christiano, Eichenbaum and Evans (1999) or Garratt *et al.* (Chapter 4, 2006) for reviews.

conjunction with the real time data, both to moderate the direct impact of the revisions obtained in successive vintages of data on the perceived current output level and to look forward to offset their impact on the output trend measure.

In this paper, we exploit the information contained in the sequence of vintages more fully than OvN through a cointegrating VAR model which, under reasonable assumptions on the nature of the output series and measurement errors, explains both the changes in the real time data and its revisions. The model is used to generate forecasts of contemporaneous and future values of output. The forecasts improve the accuracy with which the true level of activity is measured and they can be used to supplement the historically-observed series to obtain improved measures of the underlying trends also. For example, as explained in Mise, Kim and Newbold (2005a,b) [denoted MKN], this latter point helps to address the end-of-sample problems associated with the widely-used Hodrick-Prescott (1997) [HP] filter in the measurement of the trend (this being the source of considerable estimation error variance). The model can be estimated recursively, taking into account successive vintages of data. But, because it describes the revision process as well as the underlying output process, the model makes use of all the information available at each point in time, not just the most recent vintage available.

The proposed approach to measuring the output gap has at least three very useful properties. *First*, the output gap is measured relatively precisely because modelling the revision process moderates the effect of changes in published data, while the use of the forecasts mitigates any end-of-sample problems associated with the measure of the trend. *Second*, by linking the trend measure to forecasts of future output levels, it can readily be interpreted in terms of economically-meaningful concepts such as ‘potential output’. And *third*, as well as producing point estimates of the output gap that are measured relatively precisely, the underlying model can be used to describe clearly the uncertainties associated with the measure of the gap. This is extremely useful because, while it is important to recognise the unreliability of the output gap measures, the estimated values of the gap at different horizons are nevertheless an essential requirement in many decision-making contexts. The output gap measures can be used appropriately, taking into account the uncertainties surrounding them, when the model is used to supplement the point forecasts

with forecasts of the probability of the occurrence of particular events involving the gap.

The remainder of the paper is organised as follows. In Section 2, the proposed method for measuring the output gap is elaborated through a description of the cointegrating VAR model, through a discussion of the end-of-sample problems encountered when measuring trends in real time and through a comment on the calculation of probability forecasts relating to the output gap. Section 3 describes the application of the proposed methods to obtain output gap measures for the US and compares these with measures obtained following alternative procedures. Section 4 presents some probability forecasts obtained using our modelling framework, and Section 5 concludes.

2 Measuring the Output Gap with Real Time Data

To describe our proposed method of measuring the output gap, we need to introduce some notation and terminology. We write (the logarithm of) the output level at time $t - j$ by y_{t-j} , and denote the measure of output at time $t - j$ that is released in time t by ${}^t y_{t-j}$, $j = 0, 1, 2, \dots$. Throughout the paper, the “vintage- t ” dataset is defined by $Y_t = \{{}^t y_{t-1}, {}^t y_{t-2}, {}^t y_{t-3}, \dots\}$ so that it includes the time- t measure of output at time $t - 1$ and before. Note that it is assumed that the first release of output data for any period takes place after a one-period delay; this corresponds to practice in the US, for example. The full information set available at time t , denoted Ω_t , contains the datasets of all vintages dated at t and earlier; i.e. $\Omega_t = \{Y_t, Y_{t-1}, Y_{t-2}, \dots\}$. It is worth noting that the time- $(t + 1)$ measure of a variable is simply the time- t measure plus the revision; i.e. ${}^{t+1}y_{t-1} = {}^t y_{t-1} + ({}^{t+1}y_{t-1} - {}^t y_{t-1})$. Hence, the full information set grows with the addition of successive vintages of datasets by including the news on the output level in the previous period (the ‘first release’ of information on the output level in that period) and the revisions on the output series in previous periods; i.e. $\Omega_{t+1} = \Omega_t \cup \{{}^{t+1}y_t, ({}^{t+1}y_{t-1} - {}^t y_{t-1}), ({}^{t+1}y_{t-2} - {}^t y_{t-2}), \dots\}$. Finally, turning to the output trend, we note that there are a variety of methods employed in the literature to obtain measures of the output trend at time t . Some of these make use of data that becomes available both before and after time t , so that care also needs to be exercised in describing the information set on which the trend measure is based. Specifically, writing the trend output level at time $t - j$ by \tilde{y}_{t-j} , we

denote the measure of trend output at time $t - j$ that is calculated using method k on the basis of an information set available at time t , say Ω_t , by $\tilde{y}_{t-j}^k|\Omega_t$.

In OvN, attention is focused on the differences between ‘real time’ measures of the output gap based on successive vintages of output data and ‘final’ measures obtained from the last available vintage of data. Hence the comparison is between the real time measure of the gap $x_t^{ro} = {}_{t+1}y_t - \tilde{y}_t^o|Y_{t+1}$ and the final measure $x_t^{fo} = {}_T y_t - \tilde{y}_t^o|Y_T$, where $t = 1, \dots, T - 1$, and the ‘o’ superscript denotes the HP filter method used by OvN.³ OvN also consider a ‘quasi-real’ estimate, $x_t^{qo} = {}_T y_t - \tilde{y}_t^o|Y_{T,t}$, in which the time- T measure of output at time t is compared to a trend measure obtained on the basis of a subset of Y_T ; namely $Y_{T,t} = \{{}_T y_t, {}_T y_{t-1}, {}_T y_{t-2}, \dots\}$, $t < T$. Evaluating the differences between the quasi-real measure of the output gap x_t^{qo} and the real time measure x_t^{ro} isolates the changes in the gap arising from the revision of the trend in the light of subsequent data. OvN find this element to be significant but relatively small, and it is argued that it is the addition of new points to the sample, which causes $\tilde{y}_t^o|Y_T$ to deviate from $\tilde{y}_t^o|Y_{T,t}$, that explains much of the difference between the real time and final measures of the output gap.

OvN’s three measures of the output gap highlight the different effects of revisions in published data and of differences in the use of information. But their decomposition is potentially misleading. For example, focusing on vintage- t data without reference to the revisions that have taken place in previous periods’ data potentially overstates the effects of changes in the published data in time t , since these revisions might have been anticipated. Indeed, even if only vintage- t data is used, predictions of future output levels will be helpful in measuring the trend at the end of the sample whenever the time- t value of the trend is related to its value in adjacent periods. The conclusion, then, is that all information available at time t should be employed in constructing an output gap measure in real time, with particular attention paid to forecasts of future values of the output series. The appropriate modelling framework for accommodating all information

³For expositional purposes, we initially focus on the HP filter, but OvN also illustrate the uncertainty in the gap measures arising from the choice of detrending technique: deterministic trends, HP filter, unobserved components, and so on. See also Canova (1998).

is described in the section below, and this is then used to explain how forecasts can be used to eliminate the end-of-sample problems associated with the measures of the trend.

2.1 A Joint Model of Actual and Revised Output Series

In order to make use of the full information available, the real time measures of output should be modelled alongside the “actual”, realised value of output, taking into account the revision process as well as the underlying output process.⁴ In most of this section, we assume for illustrative purposes that data is revised just once after its initial release, so that we can model the two processes jointly in a bivariate VAR. However, we note also that if revisions continue up to q periods after the first release of data, then a VAR of size $q + 1$ would be required to model the processes adequately, and we illustrate this more general case too.

Our modelling approach assumes first that actual output is first-difference stationary. This means that, if data on output is released with a one period delay and the actual output is observed with the revision after one further period, $({}_t y_{t-2} - {}_{t-1} y_{t-3})$ is stationary. The approach also assumes that measurement errors (i.e. revisions) are also stationary. The first of these assumptions is supported by considerable empirical evidence,⁵ and the latter is eminently reasonable. Under these assumptions, any linear combination of these two series can be modelled in a bivariate VAR.⁶ Hence, the output growth measure $({}_t y_{t-1} - {}_{t-1} y_{t-2})$ and the data revision series have the following joint fundamental Wold representation:

$$\begin{bmatrix} {}_t y_{t-1} - {}_{t-1} y_{t-2} \\ {}_t y_{t-2} - {}_{t-1} y_{t-2} \end{bmatrix} = \begin{bmatrix} \alpha_1 \\ \alpha_2 \end{bmatrix} + \mathbf{A}(\mathbf{L}) \begin{bmatrix} \epsilon_t \\ \xi_t \end{bmatrix} \quad (2.1)$$

Here, α_1 is mean output growth (measured by ‘first-release’ data), α_2 is the mean value of the revisions, $\mathbf{A}(\mathbf{L}) = \sum_{j=0}^{\infty} \mathbf{A}_j(L)$, where the $\{\mathbf{A}_j\}$ are 2×2 matrices of parameters,

⁴See also Howery (1978) and Diebold and Rudebusch (1991).

⁵See, for example, Pappell and Prodan (2004).

⁶For example, output growth measured by the change in the ‘first-release’ output level, $({}_t y_{t-1} - {}_{t-1} y_{t-2})$, can be written in terms of actual growth and the relevant revisions and so is itself stationary; i.e. $({}_t y_{t-1} - {}_{t-1} y_{t-2}) = ({}_{t+1} y_{t-1} - {}_t y_{t-2}) + ({}_t y_{t-1} - {}_{t+1} y_{t-1}) - ({}_{t-1} y_{t-2} - {}_t y_{t-2})$.

assumed to be absolutely summable, and L is the lag-operator. Also, ϵ_t and ξ_t are mean zero, stationary innovations, with non-singular covariance matrix $\Psi = \psi_{jk}$, $j, k = 1, 2$. The model in (2.1) emphasises the point that the chosen measure of output growth at time $t - 1$ and the revision of the measure of output at time $t - 2$ between $t - 1$ and t are both revealed at time t . For notational convenience, in what follows we write $\alpha = (\alpha_1, \alpha_2)'$, where $\alpha_2 = 0$ if there is no bias in the measurement error.

The general model in (2.1) can be expressed in various different ways. For example, assume that $\mathbf{A}^{-1}(\mathbf{L})$ can be approximated by the lag polynomial $\mathbf{A}^{-1}(\mathbf{L}) = \mathbf{B}_0 + \mathbf{B}_1\mathbf{L} + \dots + \mathbf{B}_{p-1}\mathbf{L}^{p-1}$, where $\mathbf{B}_0 = \mathbf{I}_2$ without loss of generality. In this case, (2.1) can be rewritten to obtain the AR representation

$$\begin{bmatrix} {}_t y_{t-1} & {}_{-t-1} y_{t-2} \\ {}_t y_{t-2} & {}_{-t-1} y_{t-2} \end{bmatrix} = \mathbf{a} - \mathbf{B}_1 \begin{bmatrix} {}_{t-1} y_{t-2} & {}_{-t-2} y_{t-3} \\ {}_{t-1} y_{t-3} & {}_{-t-2} y_{t-3} \end{bmatrix} - \dots - \mathbf{B}_{p-1} \begin{bmatrix} {}_{t-p+1} y_{t-p} & {}_{-t-p} y_{t-p-1} \\ {}_{t-p+1} y_{t-p-1} & {}_{-t-p} y_{t-p-1} \end{bmatrix} + \begin{bmatrix} \epsilon_t \\ \xi_t \end{bmatrix} \quad (2.2)$$

and hence

$$\begin{bmatrix} {}_t y_{t-1} \\ {}_t y_{t-2} \end{bmatrix} = \mathbf{a} + \Phi_1 \begin{bmatrix} {}_{t-1} y_{t-2} \\ {}_{t-1} y_{t-3} \end{bmatrix} + \Phi_2 \begin{bmatrix} {}_{t-2} y_{t-3} \\ {}_{t-2} y_{t-4} \end{bmatrix} + \dots + \Phi_p \begin{bmatrix} {}_{t-p} y_{t-p-1} \\ {}_{t-p} y_{t-p-2} \end{bmatrix} + \begin{bmatrix} \epsilon_t \\ \xi_t \end{bmatrix} \quad (2.3)$$

where $\mathbf{a} = \mathbf{A}^{-1}(\mathbf{1})\alpha$,

$$\Phi_j = \mathbf{B}_{j-1} \begin{bmatrix} 1 & 0 \\ 1 & 0 \end{bmatrix} - \mathbf{B}_j \text{ for } j = 1, \dots, p-1, \quad \text{and} \quad \Phi_p = \mathbf{B}_{p-1} \begin{bmatrix} 1 & 0 \\ 1 & 0 \end{bmatrix}.$$

Seen in the context of (2.3), the vector of errors $(\epsilon_t, \xi_t)'$ has a clear interpretation: ϵ_t is the “news on output level in time $t - 1$ contained in the first-release data becoming available at time t ”; and ξ_t is the “news on the level of output in time $t - 2$ contained in the revised data becoming available at time t ”.

Alternatively, manipulation of (2.3) also provides the VECM representation explaining the changes in the first release measures and the change in output realisations, $[\Delta_t y_{t-1}, \Delta_t y_{t-2}]$ where $\Delta = (1 - L)$ is the difference operator. As shown in the Appendix, the VECM representation includes the lagged value of $({}_t y_{t-1} - {}_t y_{t-2})$ as a regressor since these two series are cointegrated, with cointegrating vector $\beta' = [1, -1]$. This property holds because revisions are taken to be stationary in this model, so that first-release and

actual output levels are cointegrated by assumption.⁷ Note that the model at (2.1), and its equivalent forms, are quite general and have no implications for the nature of the measurement error other than it is stationary. However, the assumption that real time measures are unbiased (in the sense that measurement errors have no systematic content) can be accommodated in the model through the imposition of restrictions. If first-release measures are unbiased, we would have ${}^t y_{t-2} = {}_{t-1} y_{t-2} + \xi_t$ so that, in (2.3), the second row of $\Phi_1 = \begin{pmatrix} 1 & 0 \end{pmatrix}$, and the second row of $\Phi_j = \begin{pmatrix} 0 & 0 \end{pmatrix}$, $j = 2, \dots, p$.

Finally here, we note that the above models can be readily extended when the revision process extends beyond just one period. Hence, for example, if quarterly data continues to be revised for up to a year, then the data requires a four-variable VAR to capture the joint determination of the first-release output series and the three successive revisions. Hence, the model that will accommodate the news on output levels contained in the first-release data (ϵ_t) and in all the revised data becoming available at time t on the previous periods ($\xi_{1t}, \xi_{2t}, \xi_{3t}$) can be written in a form corresponding to (2.2),

$$\begin{bmatrix} {}^t y_{t-1} & {}_{t-1} y_{t-2} & {}^t y_{t-2} \\ {}^t y_{t-2} & {}_{t-1} y_{t-2} & {}^t y_{t-2} \\ {}^t y_{t-3} & {}_{t-1} y_{t-3} & {}^t y_{t-3} \\ {}^t y_{t-4} & {}_{t-1} y_{t-4} & {}^t y_{t-4} \end{bmatrix} = \mathbf{a} - \mathbf{B}_1 \begin{bmatrix} {}_{t-1} y_{t-2} & {}_{t-2} y_{t-3} \\ {}_{t-1} y_{t-3} & {}_{t-2} y_{t-3} \\ {}_{t-1} y_{t-4} & {}_{t-2} y_{t-4} \\ {}_{t-1} y_{t-5} & {}_{t-2} y_{t-5} \end{bmatrix} - \dots \\ - \mathbf{B}_{p-1} \begin{bmatrix} {}_{t-p+1} y_{t-p} & {}_{t-p} y_{t-p-1} \\ {}_{t-p+1} y_{t-p-1} & {}_{t-p} y_{t-p-1} \\ {}_{t-p+1} y_{t-p-2} & {}_{t-p} y_{t-p-2} \\ {}_{t-p+1} y_{t-p-3} & {}_{t-p} y_{t-p-3} \end{bmatrix} + \begin{bmatrix} \epsilon_t \\ \xi_{1t} \\ \xi_{2t} \\ \xi_{3t} \end{bmatrix}. \quad (2.4)$$

This can be rewritten in levels form, in VECM form and in MA form exactly as in (2.3) and the models of the Appendix.

2.2 Measuring Trend Output and the Output Gap

Estimates of the bivariate or multivariate models derived above can be used to generate forecasts of the output series infinitely into the future and, in this section, we argue that

⁷The VECM representation also has implications for the corresponding MA representation in first differences; see Appendix for details.

these can be usefully applied in the measurement of the output trend whenever this is related to the trend in adjacent periods (i.e. both backwards and forwards in time). To motivate this procedure, we focus on the HP filter which is an additive decomposition $y_t = \tilde{y}_t + x_t$ where \tilde{y}_t is identified as a growth (trend) component and x_t as a cyclical component. The HP filter is an exponentially weighted moving average filter, and is two-sided symmetric in the sense that it uses both past and future observations with equal importance in order to decompose any one observation in a series. The HP filter has the desirable property that it is optimal, in the expected squared error sense, for data generating processes of the form

$$\begin{aligned} (1 - L)^2 \tilde{y}_t &= A(L) \varepsilon_t & ; & & x_t = A(L) u_t & \quad (2.5) \\ A(L) &= \sum_{j=0}^{\infty} a_j L^j & ; & & \sum_{j=0}^{\infty} a_j^2 < \infty \end{aligned}$$

where ε_t and u_t are mutually stochastically uncorrelated white noise processes (i.e. $E(\varepsilon_t u_s) = 0 \forall t, s$), and where their variance ratio is $\lambda = \left[\frac{\sigma_u}{\sigma_\varepsilon} \right]^2$, with λ being the value of the ‘smoothness’ parameter.⁸ Moreover, although the optimality conditions are expressed in terms of unobserved components, MKN show that all ARIMA($p, 2, q$) models that can be fitted to the observed series y_t can be expressed in this framework. In particular, this holds true for all possible ARIMA($p, 1, q$) models, with $A(L)$ in (2.5) involving a unit moving average root, so that the series and its trend component are $I(1)$. Here, if y_t is an ARIMA($p, 1, q$), then \tilde{y}_t is ARIMA($p + 2, 1, q$) and x_t is ARMA($p + 2, q + 1$).

However, an important feature of the HP filter is that, when we have a finite series, the optimality properties only hold for the mid-point of the series. As we move towards the end of the series, the HP filter becomes increasingly one-sided, and for the last observation of the series, the filter is completely one-sided. MKN note that the filter continues to provide an unbiased estimate of the quantity x_t at the endpoints of a finite series but that the estimates are inefficient. They illustrate the extent of the inefficiency by comparing the estimated HP trend measures with the actual trends present in a variety of simulated series obtained using different trend and cycle specifications, finding that the estimation variance

⁸This parameter is conventionally set to 1600 for quarterly data, following a suggestion by Hodrick and Prescott (1997).

of the trend is up to 40 times that of the error inherent in the series in some circumstances (see also Baxter and King, 1999, and St-Amant and van Norden, 1998). To address the inefficiency issue, MKN note Burman’s (1980) suggestion to augment the observed series with optimal linear forecasts and demonstrate, through their simulation exercises, that the application of the HP filter to the augmented series provides an estimate of the end-of-sample observation which is optimal. Indeed, by augmenting a series by its forecast, the standard deviation of the estimation error for the cyclical component is reduced by up to half (relative to the standard application of the HP filter) in their various simulations.⁹

The clear implication of these results is that the output gap should be calculated using a trend obtained by applying the filter to the forecast-augmented output series. For the series described in the previous section, the model at (2.1), or its equivalent forms in (2.2) or (2.3), can provide the vehicle for generating these forecasts. Forecasts of the output series $t+1y_t, t+2y_{t+1}, t+3y_{t+2}, \dots$ could be generated using a univariate model of the vintage- t data, but this will generally be less efficient than that provided by the bivariate model of (2.1) which uses all the information available. We shall denote the end-of-sample trend measure obtained by applying the HP filter to the output series augmented by forecasts from the univariate model obtained using vintage- t data by $\tilde{y}_{t-1}^{uh}|Y_t$ and the corresponding measure obtained using the bivariate model of (2.1) by $\tilde{y}_{t-1}^{mh}|\Omega_t$. In the empirical section, we shall also consider gap measures obtained by applying an exponential smoothing filter and Watson’s (1986) unobserved components model to the forecast-augmented data for the purpose of comparison; these are denoted with a ‘e’ and ‘w’ superscript so that the multivariate versions of the series are $\tilde{y}_{t-1}^{me}|\Omega_t$ and $\tilde{y}_{t-1}^{mw}|\Omega_t$ respectively.

The application of the HP filter to the forecast-augmented series not only improves the statistical properties of the derived series but it also justifies an economically-meaningful interpretation of the trend. Specifically, forecasts of future output levels show the expected

⁹MKN also note that the HP filter is often used in contexts where there is no assumed underlying ‘true’ trend and cycle measures of the form (2.5) or indeed any other form. They comment that the reliability of a trend measure can be assessed in these circumstances if a measure based on a sample of data $1, \dots, T$ is revised as little as possible in the light of subsequent observations; this matches the discussion of OvN on the comparison of their ‘quasi real’ and ‘final’ trend estimates. MKN confirm through their simulations that these revisions are indeed minimised when the HP filter is applied to the forecast-augmented series.

evolution of the series in the absence of further shocks, so that the infinite-horizon outcome can be readily interpreted as the economy's "potential output" level.¹⁰ A trend measure based on a forecast-augmented series will coincide with this potential output series at long horizons by construction. As discussed above, the optimality of a particular filter in identifying the trend at shorter horizons depends on the underlying data generating process. Unless economic theory can provide sufficient detail on the nature of the short run dynamics, an investigator might want to consider a number of alternative trend measures. But focusing attention on trends using forecast-augmented series ensures the trend is consistent with expected future output levels and matches the potential output concept in the long run.

2.3 Conveying the Uncertainty Surrounding the Output Gap Measures

In practice, decision-makers faced with the complete set of vintages of data up to and including that at time T are concerned with obtaining a measure of the output gap for the end-of-sample period (and possibly into the future). In some cases, attention focuses simply on whether the gap is positive or negative, but in any case it is the time- T (and future) magnitudes that matter in real time decision-making. Here, assuming again that data is released with a one period delay and there is a single revision made, this means decision-makers are interested in forecasts of $x_T^{fk} =_{T+2} y_T - \tilde{y}_T^k | \Omega_{T+N}$ for a trend measure k and for large N . Hence, the relevant output level to be forecast is $_{T+2}y_T$, the time- T output level that will be observed in $T+2$, taking into account the one period delay in the release of data and after any revisions in the data have been fully taken into account. And the relevant trend measure to be forecast is that obtained on the basis of an information set that is available at some forecast horizon well into the future (at $T+N$) so that there are no end-of sample problems for the measure at T .

We can obtain point forecasts of this magnitude relatively easily: the point forecast of $_{T+2}y_T$ is obtained straightforwardly from the bivariate model of (2.2) based on Ω_T ;

¹⁰The Beveridge-Nelson (1981) trend highlights precisely this infinite-horizon outcome, abstracting from the dynamic path that will be involved in reaching the potential output level.

and the forecast of $\tilde{y}_T^k|\Omega_{T+N}$, based on Ω_T , is simply the period- T observation of $\tilde{y}_T^k|\Omega_T$.¹¹ But the point forecast of the gap obviously does not convey the uncertainty associated with the output gap measure, and this is potentially significant here given that forecasts of the revised and unrevised series are used in various different ways in the construction of the measure. So, using the information set Ω_T for example, there will be uncertainty associated with the output gap measure at time $T - 2$ because of the need to forecast the values of output beyond T and the consequent imprecision in the measure of the trend. (Of course, the estimation variance due to the end-of-sample problem is reduced by the forecast augmentation but not eliminated). This uncertainty is compounded in the measure dated at $T - 1$ by the forecast revisions that will be made to the first-release data on ${}_T y_{T-1}$ and then further compounded at T and beyond as the unrevised output series and revisions are subsequently forecasted.

It is important, therefore, that any output gap measure is supplemented with information on the uncertainties associated with the measure. Indeed, it is sometimes argued that decision-makers' objective functions are concerned with 'booms' and 'recessions' (i.e. whether the output gap is positive or negative, irrespective of its size) and that these episodes are not valued symmetrically so that the costs incurred during a recession might outweigh the benefits experienced in boom, say. (See Cukierman and Gerlach, 2003, and references therein, for example). Similarly, there is an argument that policy-makers are concerned with whether conditions are improving or deteriorating, with the gap rising or falling (see Walsh 2003, for example). In these circumstances, the decision-maker requires the entire probability density function (pdf) of the estimated output gap measure rather than its point forecast or, at least, explicit forecasts of the probability of the event of interest (i.e. the probability that the output gap will exceed or fall below zero, or the probability of a turning point).¹²

¹¹This follows because the measure $\tilde{y}_T^k|\Omega_T$ is itself based on forecast values of the future unrevised and revised series and in the absence of any additional information, the value of the updated series expected to be observed in $T + N$ is unchanged from that measured in period T (cf. the Law of Iterated Expectations).

¹²Point forecasts will provide sufficient information for decisions only in the special case of the "LQ problem" involving a single decision variable (where the objective function is quadratic and constraints, if they exist, are linear); see Pesaran and Skouras (2002). For output gap measures, it is widely recognised

The calculation of probability forecasts and pdf's of this sort is relatively unusual in economics (where uncertainty is typically conveyed, if at all, by the reporting of confidence intervals). But the methods are relatively straightforward to implement and are described in Garratt *et al.* (2003). For example, abstracting from parameter uncertainty for the time being, to calculate the pdf associated with the forecast of $x_T^{fk} = {}_{T+2}y_T - \hat{y}_T^k | \Omega_{T+N}$, one would use the estimated model of (2.2) to generate R replications of the future vintages of data, denoted $\hat{Y}_{T+n}^{(r)}$ for $n = 1, \dots, N$ and $r = 1, \dots, R$. These include values of ${}_{T+2+n}\hat{y}_{T+n}^{(r)}$, $n = 0, 1, 2, \dots, N - 2$, on which the trend measure $\hat{y}_T^{k(r)} | \Omega_{T+N}^{(r)}$ can be based. The simulated distribution of $\hat{x}_T^{fk(r)} = {}_{T+2}\hat{y}_T^{(r)} - \hat{y}_T^{k(r)} | \Omega_{T+N}^{(r)}$ obtained in this way provides the pdf of the output gap measure directly. Equally, counting the number of times an event occurs in these simulations provides a forecast of the probability that the event will occur; the fraction of the simulations in which $\hat{x}_T^{fk(r)} > 0$ provides an estimate of the forecast probability that the time- T output gap is positive, for example. Extending the simulation exercise to accommodate parameter uncertainty is relatively straightforward (see Garratt *et al.* (2003) for more details) and the methods can also readily accommodate the use of alternative trend measures.¹³ Hence, a complete characterisation of the uncertainty surrounding the output gap measure can be obtained, accommodating stochastic uncertainty, parameter uncertainty, and the uncertainties associated with the appropriate measure of trend.

3 Output Gaps in the US

The methods described above are applied to the vintages of US output data provided by the Federal Reserve Bank of Philadelphia at <http://www.phil.frb.org/econ/forecast/index.html>. This dataset includes 157 vintages of data; the first vintage is dated 1965q4 and the final vintage is dated 2004q4. All vintages of data run from 1947q1 up to one period prior to the

that the design of optimal monetary policy requires a more sophisticated treatment of uncertainty than the LQ framework; see, for example, Svensson (2001, 2002).

¹³Specifically, the alternative measures of the trend can be calculated in each of the simulation exercises to provide alternative gap measures. Assigning appropriate weights to the alternative trend measures, the simulations for each trend can then be pooled to provide density functions for the gap measures and associated event probability forecasts.

release date; i.e. $Y_t = \{ {}_t y_{1947q1}, \dots, {}_t y_{t-1} \}$, $t = 1965q4, \dots, 2004q4$. The US National Income and Product Account (NIPA) figures that include an observation of output in quarter t for the first time are released at the end of the first month of quarter $t + 1$. This is the vintage that is identified in the Philadelphia database as being the data that exists at the mid-point of the quarter $(t + 1)$ and which we term Y_{t+1} . The effects of two subsequent revisions to the NIPA data, taking place at the end of the second and third months of quarter $(t + 1)$, are captured in the Philadelphia database when it reports the available data at the mid-point of the following quarter $(t + 2)$, termed Y_{t+2} in this paper.¹⁴

The first exercise undertaken on this data aims to investigate the gains from using the forecast augmented approach to defining the trend, focusing on the case where the trend is obtained using the HP filter. In the first instance, we follow OvN and consider the successive vintages of data, applying the HP filter, to derive the ‘real-time measure’ $\tilde{y}_t^o | Y_{t+1}$, $t = 1965q4, \dots, 2004q3$ as the end-of-sample observation of the trend in each recursion. We compare this with the ‘quasi real’ measure $\tilde{y}_t^o | Y_{T,t}$, also derived recursively, and the ‘final’ measure $\tilde{y}_t^o | Y_T$. We also derive the corresponding trends based on data augmented by forecasts. The forecasts are based on eighth-order univariate autoregressions explaining $({}_t y_{t-1} - {}_t y_{t-2})$; an eighth-order autoregression is applied to ensure there is no serial correlation in the residuals.¹⁵

Table 1 reports statistics relating to the output gaps considered by OvN, namely $x_t^{ro} = ({}_{t+1} y_t - \tilde{y}_t^o | Y_{t+1})$, $x_t^{qo} = ({}_T y_t - \tilde{y}_t^o | Y_{T,t})$ and $x_t^{fo} = ({}_T y_t - \tilde{y}_t^o | Y_T)$, for $t = 1965q3, \dots, 2004q3$, and $T = 2004q4$, and illustrates the considerable differences arising out of data revisions and the end-of-sample effects on the underlying trends. The Table shows that the correlation between the real time and final measures of OvN is just 0.526, and the two measures agree on whether output growth is above or below trend in only 63% of the sample period.

¹⁴The NIPA data is also usually revised each July for the prior three years. These July revisions are different in nature to those occurring at other times and this could mean that their differential impact, and the consequent seasonal effects, should be taken into account in the model. However, extensions made to accommodate these July effects in the multivariate models discussed below failed to add significantly to the fit of the models. (Details are available from the authors on request.)

¹⁵Details of regressions, and diagnostic tests relating to the order of integration of the output and revision series, are not presented for space considerations but are available from the authors on request.

These figures rise to 0.553 and 69% respectively when the comparison is between the quasi real time and final measures (abstracting from effects of data revision) but the figures are clearly still not high. Taking the final measure x_t^{fo} as the best indicator of the true output gap available to OvN, it is the poor performance of the x_t^{ro} and x_t^{qo} measures in reflecting the true output gap that is the basis of OvN's conclusion that real time measures of the gap are unreliable.

Table 1 also describes the effect of employing the forecast-augmentation approach to calculating the trends on the three gap measures, x_t^{ruh} , x_t^{quh} and x_t^{fuh} where the 'uh' superscript indicates that the underlying trend is based on the HP filter applied to a forecast-augmented series obtained using the univariate model. This has a substantial impact on the variability of the output gap series, cutting the standard deviation and range of values for the real time measure by around 30% and by nearer 40% in the case of the quasi-real measure. This illustrates that the forecast-augmentation is having a considerable impact on the trend measure as the estimation error variance associated with the application of the HP filter at the end-of-sample is reduced. The effect is to raise the correlation between the final measure x_t^{fuh} and the real and quasi real measures to 0.77 and 0.78 respectively, and agreement on the occurrence of booms and recessions rise to 83% and 81% respectively also. The improvement in reliability using the forecast-augmentation method is pronounced and shows the importance of the augmentation in calculating output gap measures.

Next, we turn to the multivariate analysis of the output growth and revision processes together, considering whether there are systematic patterns in the data revisions that underlie the successive vintages of data and the extent to which a model of the output growth data is enhanced by modelling the measured output growth and revisions data jointly. To do this, we need to choose the lag length p in the multivariate model in (2.4) and the length of the "revision horizon" (after which revisions are unsystematic and insignificant). The maximum lag length we consider is $p = 4$ and the maximum length of the revision horizon we consider is 3, as in (2.4). It turns out that the data is described adequately if we allow for a revision horizon of two quarters and lags in the VAR of order 2. To demonstrate this, Table 2 provides estimates of (2.4) obtained using

the entire data up to and including Y_{2004q4} .¹⁶ The Table shows that a revision horizon of 2 is sufficient to capture systematic elements in the revision process, since none of the variables in the fourth column, explaining time- t revisions of data at $t - 4$ are individually or jointly statistically significant. And the Table also provides variable exclusion tests, denoted $\chi^2_{LM}(10)$, showing that the third and fourth lags of the first three variables in our system and all four lags of the fourth can be safely dropped from the regressions without violating the data. Table 2 therefore confirms that the joint modelling of the growth series and the revisions is a useful approach: both the lagged growth series and the lagged revisions contribute significantly to the explanation of the time- t growth $({}_t y_{t-1} - {}_{t-1} y_{t-2})$, meaning that the univariate model is misspecified, and there are very significant systematic elements in the revisions $({}_t y_{t-2} - {}_{t-1} y_{t-2})$ and $({}_t y_{t-3} - {}_{t-1} y_{t-3})$.¹⁷

The regression analysis shows that only the first two revisions $({}_t y_{t-2} - {}_{t-1} y_{t-2})$ and $({}_t y_{t-3} - {}_{t-1} y_{t-3})$ contain systematic content, and this suggests that it might be reasonable to work with an adjusted dataset in which the subsequent revisions are assumed to be precisely zero (so that ${}_{t-k+s} y_{t-k} = {}_t y_{t-k}$, $k = 4, 5, \dots$ and for $s = 3, 4, \dots, k - 1$). The treatment of the unsystematic revisions in the regression analysis, and the choice between using the adjusted or unadjusted dataset, determines the way in which measurement error enters the system and could potentially introduce biases in the estimated parameters. The choice between the two datasets depends on the nature of the (unobservable) data generating process for the revisions and output data. The use of the adjusted data is appropriate if the revision process is a function of the "true" output whose historical values are accurately measured by the most recent vintage of data. The use of the unadjusted data would be more reasonable if the revisions are functions of growth as measured at the time (cf. Koenig et al., 2003). In the event, the correlation between the real time and final vintage measures of the gaps based on the adjusted and unadjusted datasets are 0.95 and 0.93 respectively (with agreement on booms and slumps in 91% and 93%

¹⁶In fact, we conducted this exercise recursively on the full information set, Ω_t , consisting of all of the vintages of data upto and including Y_t , for $t = 1965q4, \dots, 2004q4$. Although we only report the results of the 2004q4 analysis, qualitatively similar results were obtained throughout.

¹⁷Similar systematic elements are found in Swanson et al. (1999).

of the sample). This confirms that the adjustment and the choice of the dataset has a relatively minor impact in this case. Further, employing the adjusted dataset ensures that the most up-to-date information on historical output levels is used in constructing the output gap measures at any time. Hence, our suggested measure of the output gap based on the HP filter is that obtained by applying the forecast-augmented technique based on the multivariate model estimated using the adjusted dataset; this is denoted $x_t^{fmh} = {}_T y_t - \tilde{y}_t^{mh} | \Omega_T$, $t = 1, \dots, T - 3$.¹⁸

Table 3 provides summary statistics relating to this series, and the corresponding real time measure obtained applying the procedure recursively over time, x_t^{rmh} , for our data up to 2004q1 (i.e. for $T = 2004q4$). These figures show that the advantages of the forecast-augmentation remain, with a correlation between the real time measure and the final measure of 0.75 and agreement on booms and recessions in 84% of the sample. Table 3 also provides statistics relating to the output gap measures obtained using two alternative methods for measuring the trend in place of the HP filter. The measure denoted x_t^{rme} refers to the gap obtained in real time and based on an exponential smoothing (ES) filter. The filter is applied to the post-revision series augmented with forecasts based on our multivariate model.¹⁹ The measure x_t^{rmw} applies Watson's (1986) unobserved components (UC) model to the same series.²⁰ As discussed in King and Rebelo (1993), the ES smoothing can be considered as a restricted version of the HP filter and can be motivated as providing the filter that minimises trend growth (as opposed to the HP filter which minimises the change in trend growth). The UC model permits more complex dynamics and is consistent with a more volatile trend measure, than HP characterising the trend as a random walk with drift.²¹

¹⁸Restricting attention to $t = 1, \dots, T - 3$ implies that only post-revision measures of actual output are involved and forecasts are used only in measuring trends.

¹⁹The 'smoothing' parameter was set equal to 10. This means that 85% of the weight is on observations one year either side of the observation of interest and 95% on two years either side.

²⁰The forecast augmentation here refers to the forecast of the post-revision output data for the duration of the revision process only. The sample period on which the UC measure is based runs from 1975q1, using the earlier observations to obtain initial values for the Kalman filter estimation.

²¹The variability of growth in the ES and UC trends is 70% and 67% of that of output growth compared to 15% for the HP trend.

The results in Table 3 show that the (relatively) reassuring results obtained for the HP are also found with the other two smoothers. Hence, the correlation between the real time and final vintage gap measures are relatively high, at 0.89 and 0.78 for the UC and the ES model respectively (and agreement on booms and recessions are also high at 88% and 84%). Perhaps more surprisingly, the table also shows reasonably high correlations between the gap measures obtained using the three alternative trends. The (pairwise) correlation between the three final vintage measures are in the range $[0.86, 0.96]$ and agreement on booms and recessions is in the range $[0.80, 0.93]$ despite the differences in the form and motivation of the alternative trend measures. Table 3 therefore not only confirms that the advantages of applying our modelling approach carries over to other methods of detrending, but also shows that the alternative gaps obtained in real time provide a reasonably consensual picture of the macroeconomy, at least as far as the size of the gap is concerned.

Before discussing the treatment of uncertainty in these measures, it is worth commenting on the contribution of our modelling framework, and its use of revisions data, to these results. This contribution can be judged by comparing the forecasting performance of the multivariate model with that of the univariate model and by comparing the in-sample fit of the associated gap measures. In terms of forecasting, the univariate and multivariate models can be estimated recursively over the period 1970q1 - 2004q1 and the models' forecasts of output can be compared with either of two output outcomes: namely, the first release observation of output at time t , ${}_{t+1}y_t$, or the final vintage measure ${}_T y_t$. Using the first release series as the appropriate measure of the output outcome, the root mean square forecast error (RMSFE) defined by $\sqrt{\frac{1}{T-3} \sum_{t=1}^{T-3} ({}_{t+1}y_t - \widehat{{}_{t+1}y_t})^2}$ takes the value 0.0194 for the univariate model and 0.0095 for the multivariate model, representing a 51% improvement over the univariate model. Using the final vintage data as the measure of output outcome, calculating $\sqrt{\frac{1}{T-3} \sum_{t=1}^{T-3} ({}_T y_t - \widehat{{}_T y_t})^2}$, also supports the use of the multivariate model: the RMSFE for the univariate model, for which $\widehat{{}_T y_t} = E[\widehat{{}_{t+1}y_t} | Y_t]$, takes the value 0.0455, while the RMSFE for the multivariate model is equal to 0.0389 where $\widehat{{}_T y_t}$ is measured by $\widehat{{}_T y_t} = E[\widehat{{}_{t+1}y_t} | \Omega_t]$ and equal to 0.0125 where it is measured by $E[\widehat{{}_{t+3}y_t} | \Omega_t]$; i.e. depending on whether the multivariate model is used to forecast the

first-release series ${}_{t+1}y_t$ or the post-revision series ${}_{t+3}y_t$.²² In either case, the multivariate model represents an improvement over the univariate model, of around 15% for the first release series and 73% for the full post-revision case.

The advantages of using the multivariate model are confirmed also by the in-sample root mean squared error (RMSE), defined as the gap between the real time gap measures and the final vintage gap measures, obtained using the two models. Values of the RMSE for x_t^{rmh} , x_t^{rme} and x_t^{rmw} are 0.0110, 0.0692 and 0.0962, respectively, representing improvements of 7%, 26% and 57% over their univariate counterparts. The statistical significance of the explanatory variables in the model of the regressions explaining the revisions, the gains in forecasting of output levels and the gains in the fit of the real time gap measures all confirm that the multivariate model of growth and revisions is appropriate and will provide a firm basis on which to calculate trends and output gap measures.

4 Representing the Output Gap under Uncertainty

The analysis above shows that the uncertainty surrounding the output gap measure can be reduced through the appropriate use of forecast-augmented data, and that the forecasts are best calculated using a multivariate model that describes the measured output growth series and data revisions jointly. Nonetheless, some of the unreliability of the measures highlighted by OvN remains and so it is important that the uncertainties surrounding the measures are properly represented for decision-making purposes.

Figure 1 illustrates the order of magnitude of the uncertainties involved based on x_t^{fmh} using the information available at 2002q2 (leaving ten periods, to 2004q4, for the purpose of “out-of-sample” forecast evaluation). According to the analysis based in real time, the plot shows a period of expansion, in which the economy moves from ‘recession’ (where $x_t^{fmh} < 0$) to ‘boom’ (where $x_t^{fmh} > 0$) up to 2000q2, followed by a contraction that

²²For the univariate model, no revisions are expected to take place after the first-release data and the forecast of ${}_{t+1}y_t$ is the measure of the interest. If it is known that two systematic revisions will occur (as shown in the regression analysis), the forecast of the “true” post-revision series is the expected value of ${}_{t+3}y_t$ and it is this measure, which fully takes into account the role of predicted revisions, that is most appropriate for the multivariate model.

ends in 2001q4. These measures are conveyed with a good degree of precision at first (with 95% confidence intervals no more than $\pm 0.7\%$ prior to 2000q2) when uncertainty is derived solely from the estimation error in the underlying trend.²³ But the uncertainty rises considerably when the data uncertainties are accommodated towards the end of the sample and when forecasting out-of-sample. During the latter stages, it is difficult to interpret the information content of the gap measures either in terms of the size of the gap, the likelihood of turning points or the occurrence of boom or recession when judged simply according to the size of the confidence interval.

The information on the size and the precision of the gap can be conveyed more usefully and more directly through the corresponding probability density functions showing $prob(x_{2002q2+n}^{fm} | \Omega_{2002q2} < c)$ for a range of critical values c at various estimated horizons, n . Figure 2 shows such density functions for $n = -3, 0$ and 4 , generated using the simulation methods of Section 2.3 and again taking into account stochastic and parameter uncertainty. The functions shift to the right over time, reflecting the rising value of the point forecast and become progressively flatter reflecting the accumulating uncertainty at the end-of-sample and into the forecasting horizons. This sequence of densities illustrates well the form in which the output gap can be usefully presented for the purpose of decision-making. Further, the analysis underlying the densities can also convey insight into particular events involving the gap. Hence, the probability of a negative output gap can be seen directly from the densities, falling from 0.93 in 2001q3 to 0.72 in 2002q2 and 0.55 in 2003q2. But the underlying simulations also show, for example, that the probability of a turning point rises from 0.18 in 2001q3 to 0.64 in 2001q4 (where the point forecast turns) and to 0.80 by 2003q1 (when the upturn actually occurred).²⁴ These probabilities convey far more precisely the strength of conviction with which these events are perceived to take place. Given the uncertainties associated with the gaps discussed in

²³These intervals are generated by the simulation methods discussed in Section 2.3 and relate to the stochastic and parameter uncertainty surrounding the measures. Abstracting from parameter uncertainty, by undertaking simulations taking into account stochastic uncertainty only, generates slightly tighter but very similar confidence intervals.

²⁴A turning point is defined here as two consecutive periods of positive output growth following two periods of negative growth.

the previous section, it is clear that the density functions of Figure 2 and the associated probability forecasts provide a more useful form for representing the output gap than the point forecasts and confidence intervals given in Figure 1.

A final illustration of the usefulness of probability forecasts in conveying the uncertainties associated with the gap measures is provided in Figure 3. The figure plots the probability of a positive gap occurring one-step ahead, as measured in real time through the sample 1978q1 – 2004q4. In the figure, x_{t+1}^{rmh} and x_{t+1}^{ruh} represent the real time measures of the gap obtained by applying the HP filter to data augmented by data from the multivariate and the univariate models. The measure \bar{x}_{t+1}^m is the gap in time $t + 1$ based on the density derived by aggregating the three densities obtained for the alternative trend measures HP, ES and UC (using the multivariate model in each case and with equal weight given to each trend measure). The aggregated density accommodates the uncertainties associated with the choice of the trend measure as well as the stochastic and parameter uncertainty underlying the individual gaps in a straightforward way. In Figure 3, the information in the aggregated density is translated into a form that is directly usable by a decision-maker who is concerned with booms and recessions. As it turns out, the probability series based on the aggregated density rarely differs from that based on x_{t+1}^{rmh} by more than 10%, while there are some substantial differences between the gaps based on x_{t+1}^{rmh} and x_{t+1}^{ruh} for example. Hence, in this case at least, the ‘trend uncertainty’ appears less important than the choice of the model used to implement the forecast-augmented approach.

5 Conclusions

The analysis of this paper starts from the point that output gap measures are an essential element of many decisions but that they are measured with considerable uncertainty. This is because of the imprecision of the output data available at the time decisions have to be made and because of the difficulties in establishing a precise measure of trend output. We have shown that these uncertainties can be mitigated by modelling the output process alongside the revision process, making use of forecasts of current and future post-revision output levels, to obtain more precisely estimated measures of the gap for use

in real time decision-making. But the uncertainties surrounding the measures, correctly identified as important by OvN and others, remain and are substantial. We have also shown, therefore, that the production of forecasts of probabilities of events involving the gap convey the information on the level of the gap and the uncertainties associated with this measure more precisely than the point forecasts and confidence intervals typically delivered by analysts. The cumulative density functions that we have discussed, along with the estimated probabilities of particular events of interest, provide a very informative and helpful means of representing the output gap data for use by decision-makers.

6 Appendix

Manipulation of (2.3) in the text provides the VECM representation

$$\begin{bmatrix} \Delta_t y_{t-1} \\ \Delta_t y_{t-2} \end{bmatrix} = \mathbf{a} + \mathbf{\Gamma}_0 \begin{bmatrix} {}_{t-1}y_{t-2} \\ {}_{t-1}y_{t-3} \end{bmatrix} + \sum_{j=1}^p \mathbf{\Gamma}_j \begin{bmatrix} \Delta_{t-j+1} y_{t-j} \\ \Delta_{t-j+1} y_{t-j-1} \end{bmatrix} + \begin{bmatrix} \epsilon_t \\ \xi_t \end{bmatrix} \quad (6.6)$$

where

$$\begin{aligned} \mathbf{\Gamma}_0 &= -(\mathbf{I}_2 - \mathbf{\Phi}_1 - \mathbf{\Phi}_2 - \dots - \mathbf{\Phi}_p) = -\mathbf{\Phi}(\mathbf{1}) \\ \mathbf{\Gamma}_j &= -(\mathbf{\Phi}_{j+1} + \mathbf{\Phi}_{j+2} + \dots + \mathbf{\Phi}_p) \quad j = 1, 2, \dots, p-1 \end{aligned}$$

Given the form of the $\mathbf{\Phi}_i$ described in (2.3), it is easily shown that $\mathbf{\Gamma}_0$ takes the form

$$\mathbf{\Gamma}_0 = \begin{bmatrix} -k_1 & k_1 \\ -k_2 & k_2 \end{bmatrix} = \begin{bmatrix} -k_1 \\ -k_2 \end{bmatrix} \begin{bmatrix} 1 & -1 \end{bmatrix} \quad (6.7)$$

where k_1 and k_2 are functions of the elements of the \mathbf{B}_j , $j = 0, 1, \dots, p-1$. Hence, the model at (2.1) can be written in a VECM form where $\mathbf{\Gamma}_0 = \boldsymbol{\alpha}\boldsymbol{\beta}'$ and $\boldsymbol{\alpha}' = [-k_1, -k_2]$ contains the parameters determining the speed of adjustment to equilibrium and $\boldsymbol{\beta}' = [1, -1]$ is the cointegrating vector. The form of the cointegrating vector captures the assumption that revision errors are stationary through the inclusion of the error correction term $\boldsymbol{\beta}' [{}_{t-1}y_{t-2}, {}_{t-1}y_{t-3}]' = {}_{t-1}y_{t-2} - {}_{t-1}y_{t-3}$.

A final alternative for describing the model is the MA representation obtained through recursive substitution of (2.3):

$$\begin{bmatrix} \Delta_t y_{t-1} \\ \Delta_t y_{t-2} \end{bmatrix} = \mathbf{b} + \mathbf{C}(L) \begin{bmatrix} \epsilon_t \\ \xi_t \end{bmatrix} \quad (6.8)$$

where $\mathbf{b} = \mathbf{C}(1)\mathbf{a}$, $\mathbf{C}(L) = \sum_{j=0}^{\infty} \mathbf{C}_j(L)$, $\mathbf{C}_0 = \mathbf{I}_2$, $\mathbf{C}_1 = \mathbf{\Phi}_1 - \mathbf{I}_2$ and $\mathbf{C}_i = \sum_{j=0}^p \mathbf{C}_{i-j}\mathbf{\Phi}_j$, $i > 1$, $\mathbf{C}_i = 0$, $i < 0$. As is well known, following Engle and Granger (1987), the presence of a cointegrating relationship between the ${}_t y_{t-1}$ and ${}_{t-1} y_{t-2}$ imposes restrictions on the parameters of $\mathbf{C}(L)$; namely, $\boldsymbol{\beta}'\mathbf{C}(1) = 0$. Further, given that $\boldsymbol{\beta}' = [1, -1]$, this ensures that $\mathbf{C}(1)$ takes the form

$$\mathbf{C}(1) = \begin{bmatrix} k_3 & k_4 \\ k_3 & k_4 \end{bmatrix} \quad (6.9)$$

for scalars k_3 and k_4 .

References

- Amato, J. D., and N. R. Swanson (2001) "The Real-Time Predictive Content of Money for Output", *Journal of Monetary Economics*, 48, 3-24.
- Beveridge S. and C.R. Nelson (1981), "A New Approach to Decomposition of Economic Time Series into Permanent and Transitory Components with Particular Attention to Measurement of the Business Cycle", *Journal of Monetary Economics*, 7, 2, 151-174.
- Brunner, A.D. (2000), "On the Derivation of Monetary Policy Shocks: Should we Throw the VAR Out With the Bath Water", *Journal of Money Credit and Banking*, 32, 254-279.
- Baxter, M. and King, R.G. (1999), "Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series", *Review of Economics and Statistics*, 81(4) 55-593.
- Burman, J.P. (1980), "Seasonal Adjustment by Signal Extraction", *Journal of Royal Statistical Society A*, 143, 321-337.
- Canova, F. (1998), "Detrending and Business Cycle Facts", *Journal of Monetary Economics*, 41, 3, 475-512.
- Christiano, L.J., M. Eichenbaum and C.L. Evans (1999), "Monetary Policy Shocks: What Have We Learned and to What End?", Ch. 2 in J.B Taylor and M. Woodford (eds.), *Handbook of Macroeconomics*, Vol.1A. North-Holland, Elsevier: Amsterdam.
- Cukierman, A. and S. Gerlach (2003), "The Inflation Bias Revisited: Theory and Some International Evidence", *Manchester School*, 71, 3, 541-565.
- Diebold, F. X., and G. D. Rudebusch (1991) "Forecasting Output with the Composite Leading Index: A Real-Time Analysis", *Journal of the American Statistical Association*, 86, 603-610.

- Garratt, A., K. Lee, M.H. Pesaran, and Y. Shin, Y. (2003) "Forecast Uncertainties in Macroeconometric Modelling: An Application to the UK Economy", *Journal of American Statistical Association*, 98, 464, 829-838.
- Garratt, A., K.C. Lee, M.H. Pesaran and Y. Shin, (2006), *Global and National Macroeconometric Modelling*, forthcoming, OUP, Oxford.
- Hodrick, R.J. and E.C. Prescott (1997), "Postwar U.S. Business Cycles: An Empirical Investigation", *Journal of Money, Credit and Banking*, 29, 116-130.
- Howery, E.P. (1978) "The Use of Preliminary Data in Econometric Forecasting", *Review of Economic Statistics*, 6, 193-200.
- Koenig, E.F, S. Dolmas and J. Piger (2003), "The Use and Abuse of Real-Time Data in Economic Forecasting", *Review of Economics and Statistics*, 85(3), 618-628.
- King, R.G and S. Rebelo (1993), "Low Frequency Filtering and Real Business Cycles", *Journal of Economic Dynamics and Control*, 17:1-2, 207-231.
- Mise, E., T-H. Kim and P. Newbold (2005a), "On the Sub-Optimality of the Hodrick-Prescott Filter", *Journal of Macroeconomics*, 27(1), 53-67.
- Mise, E., T-H. Kim and P. Newbold (2005b), "Correction of the Distortionary end-effect of the Hodrick-Prescott Filter: Application", *mimeo*, downloadable from <http://www.le.ac.uk/economics/staff/em92.html>.
- Orphanides, A., R.D. Porter, D. Refschneider, R. Tetlow and F. Finan (2000), "Errors in the Measurement of the Output gap and the Design of Monetary Policy" *Journal of Economics and Business*, 52, 117-141.
- Orphanides, A. (2001) "Monetary Policy Rules Based on Real-Time Data", *American Economic Review*, 91, 964-985.
- Orphanides, A. and S. van Norden (2002) "The Unreliability of Output-Gap Estimates in Real Time", *Review of Economics and Statistics*, 84(4), 569-583.

Orphanides, A. and S. van Norden (2003) “The Reliability of Inflation Forecasts Based on Output Gap Estimates in Real Time”, *CIRANO Working Paper* 2003-01.

Orphanides, A., (1998), Monetary Policy Evaluation with Noisy Information, *Board of Governors of the Federal Reserve System, Working Paper*, October 1998.

Pappell, D.H. and R. Prodan (2004), “The Uncertain Unit Root in US Real GDP: Evidence with Restricted and Unrestricted Change”, *Journal of Money Credit and Banking*, 36, 423-427.

Pesaran, M.H. and S. Skouras (2002), Decision-based Methods for Forecast Evaluation, in Clements, M.P. and D.F. Hendry (eds.) *A Companion to Economic Forecasting*, Oxford Blackwell.

Rotemberg, J.J. and M. Woodford (1997), “An Optimisation-Based Econometric Framework for the Evaluation of Monetary Policy”, *NBER Macroeconomics Annual*, 297-345.

Rotemberg, J. and M. Woodford (1999), “Interest Rate Rules in an Estimated Sticky Price Model”, pages , 57-119 in Taylor J.B. (ed.) *Monetary Policy Rules*, University of Chicago Press: Chicago.

St-Amant, P. and S. van Norden (1998) ”Measurement of the Output Gap: A Discussion of Recent Research at the Bank of Canada”, *Bank of Canada Technical Report* No. 79.

Svensson, L.E.O.,(2001), “Inflation Targeting: Should It Be Modelled as an Instrument Rule or a Target Rule?”, *European Economic Review*, 46, 771-780.

Svensson, L.E.O. (2002), “What is Wrong with Taylor Rules ? Using Judgement in Monetary Policy through Targeting Rules”, *Journal of Economic Literature*, 41, 426-477.

Swanson, N.R., E. Ghysels and M. Callan, (1999), “A Multivariate Time Series Analysis of the Data Revision Process for Industrial Production and the Composite

Leading Indicator” in *Cointegration, Causality and Forecasting: Festschrift in Honor of Clive W.J. Granger*.

Walsh, C.E. (2003), “Speed Limit Policies: The Output Gap and Optimal Monetary Policy”, *American Economic Review*, 93, 265-278.

Watson, M W. (1986), “Univariate Detrending Methods with Stochastic Trends”, *Journal of Monetary Economics*, 18, 49-75.

Table 1: Univariate Output Gap Measures: 1965q3 – 2004q3

	x_t^{ro}	x_t^{qo}	x_t^{fo}	x_t^{ruh}	x_t^{quh}	x_t^{fuh}
Mean	-0.002	0.007	0.001	-0.005	-0.003	0.001
SD	0.018	0.020	0.016	0.013	0.012	0.016
Min	-0.066	-0.042	-0.047	-0.055	-0.031	-0.047
Max	0.038	0.052	0.038	0.016	0.024	0.038
x_t^{ro}	1	0.937	0.526	0.833	0.822	0.520
x_t^{qo}	<i>0.847</i>	1	0.553	0.813	0.912	0.623
x_t^{fo}	<i>0.631</i>	<i>0.694</i>	1	0.771	0.780	0.998
x_t^{ruh}	<i>0.771</i>	<i>0.771</i>	<i>0.822</i>	1	0.928	0.769
x_t^{quh}	<i>0.796</i>	<i>0.796</i>	<i>0.796</i>	<i>0.885</i>	1	0.779
x_t^{fuh}	<i>0.631</i>	<i>0.694</i>	<i>0.987</i>	<i>0.834</i>	<i>0.809</i>	1

Notes: Output gaps are denoted by x_t . The ‘r’, ‘q’ and ‘f’ superscripts refer to real-time, quasi-real time and final measures respectively, as described in the text; the ‘o’ and ‘uh’ superscripts refer, respectively, to trend measures based on methods described in OvN and MKN, with reference to the HP filter, using an eighth-order univariate autoregression for forecasts, again described in the text. Summary statistics in the upper panel refer to the mean, standard deviation, minimum and maximum values respectively. Figures in the lower panel refer to correlation coefficients and, in italics, proportion of the sample for which there is agreement that the output gap is positive or negative.

Table 2: Model of Output Growth and Revisions: 1967q1 - 2004q3

Independent Variable	Dependent Variable			
	$y_{t-1} - y_{t-2}$	$y_{t-2} - y_{t-3}$	$y_{t-3} - y_{t-4}$	$y_{t-4} - y_{t-5}$
intercept	0.0034 (0.0009)	0.0005 (0.0004)	0.0002 (0.0004)	0.0002 (0.0004)
$y_{t-2} - y_{t-3}$	0.5749 (0.0934)	0.0840 (0.0453)	0.0258 (0.0378)	0.0134 (0.0381)
$y_{t-3} - y_{t-4}$	-0.0171 (0.0950)	0.0134 (0.0460)	0.0404 (0.0384)	0.0462 (0.0388)
$y_{t-4} - y_{t-5}$	-0.9579 (0.3713)	-0.5251 (0.1800)	-0.3075 (0.1503)	-0.2750 (0.1516)
$y_{t-5} - y_{t-6}$	0.7657 (0.3808)	-0.1743 (0.1846)	-0.2177 (0.1541)	-0.1570 (0.1555)
$y_{t-6} - y_{t-7}$	0.6577 (0.4295)	0.4072 (0.2081)	0.2510 (0.1738)	0.2395 (0.1754)
$y_{t-7} - y_{t-8}$	-0.7371 (0.4329)	0.2587 (0.2098)	0.1888 (0.1752)	0.1177 (0.1768)
R^2	0.2817	0.0799	0.0430	0.0335
$\hat{\sigma}$	0.0081	0.0039	0.0033	0.0033
$\chi_{LM}^2(10)$	{0.04}	{0.62}	{0.44}	{0.39}
F_{SC}	{0.06}	{0.61}	{0.77}	{0.26}
F_{FF}	{0.42}	{0.94}	{0.75}	{0.81}
F_H	{0.80}	{0.13}	{0.25}	{0.44}
F_N	{1.00}	{1.00}	{1.00}	{1.00}


Notes: Standard errors are given in (.). R^2 is the squared multiple correlation coefficient, $\hat{\sigma}$ the standard error of the regression and χ_{LM}^2 a chi-squared test statistic (with 10 d.f.) for the exclusion of the third and fourth lags of the first three dependent variables and all four lags of the fourth from each of the regression equations as described in the text. The remaining diagnostics are p-values, in {.), for F-test statistics for serial correlation (SC), functional form (FF), normality (N) and heteroscedasticity (H).

Table 3: Multivariate Output Gap Measures: 1970q1 – 2004q1

	x_t^{rmh}	x_t^{fmh}	x_t^{rmw}	x_t^{fmw}	x_t^{rme}	x_t^{fme}
Mean	-0.001	-0.001	0.0078	0.0061	-0.0003	-0.0003
SD	0.014	0.016	0.0139	0.0184	0.0085	0.0093
Min	-0.066	-0.047	-0.0419	-0.0490	-0.0400	-0.0276
Max	0.051	0.038	0.0308	0.0467	0.0280	0.0246
x_t^{rmh}	1	0.748	0.841	0.717	0.972	0.767
x_t^{fmh}	<i>0.842</i>	1	0.872	0.937	0.713	0.961
x_t^{rmw}	<i>0.798</i>	<i>0.773</i>	1	0.892	0.769	0.831
x_t^{fmw}	<i>0.815</i>	<i>0.840</i>	<i>0.882</i>	1	0.650	0.863
x_t^{rme}	<i>0.892</i>	<i>0.806</i>	<i>0.798</i>	<i>0.798</i>	1	0.784
x_t^{fme}	<i>0.791</i>	<i>0.921</i>	<i>0.798</i>	<i>0.798</i>	<i>0.842</i>	1

Notes: Output gaps are denoted by x_t . The ‘r’ and ‘f’ superscripts refer to real-time and final measures described in the text, the ‘m’ superscripts refers to trend measures based on the proposed multivariate model and the ‘h’, ‘w’ and ‘e’ superscripts respectively refer to trends according to the HP filter, Watson UC model and the exponential smoother forecast-augmented approaches, as described in the text. See also notes to Table 1.

Figure 1: Output Gap Measures at 2002q2



$x_t^{fmh} | \Omega_{2002q2}$ (with 95% confidence intervals) and $x_t^{fmh} | \Omega_{2004q4}$

**Figure 2: Cumulative Density Functions for Forecast Horizons
2001q3, 2002q2 and 2003q2.**





Figure 3: Real Time Probability Forecasts of a Positive Output Gap
One-Step Ahead



$prob(x_{t+1}^{rmh} | \Omega_{t+1} > 0)$ ——— $prob(x_{t+1}^{\overline{rm}} | \Omega_{t+1} > 0)$ - - - $prob(x_{t+1}^{ruh} | Y_{t+1} > 0)$ - - - - -