

# The Australian Real-Time Database: An Overview and an Illustration of its Use in Business Cycle Analysis\*

by

Kevin Lee<sup>†</sup>, Nilss Olekalns<sup>††</sup>, Kalvinder Shields<sup>††</sup> and Zheng Wang<sup>††</sup>

## Abstract

This paper describes a newly constructed macroeconomic database for Australia including measures of GDP, its components, prices, and key monetary and labour market statistics over the last fifty years as published and revised in real time. Data vintages are collated from various sources and accommodate multiple definitional changes, providing a comprehensive description of the macroeconomic environment actually experienced by Australian policy- and decision-makers. The database exposes the difficulties in drawing inferences and decision-making based on macroeconomic data that is subsequently revised. Methods are described that can exploit the real-time dataset and they are illustrated through an analysis of the Australian output gap.

**Keywords:** Real-time, Australian Database, Revisions, Business Cycles, Output Gap, Density Forecasts

**JEL Classification:** C32, D84, E32

---

\*<sup>†</sup>University of Nottingham, UK; <sup>††</sup>University of Melbourne, Australia. Version dated August 2011. Financial support from the ARC (Discovery Grant DP0988112) is gratefully acknowledged. Corresponding author: Kalvinder Shields, Department of Economics, University of Melbourne, Victoria 3010, Australia (k.shields@unimelb.edu.au).

## 1 Introduction

There is now a great deal of interest from macroeconomic and finance researchers in using real-time data sets; that is, data sets that include not only the most recently published set of historical data on macroeconomic and financial variables (the “final vintage” dataset), but all the previous vintages of data that were available in the past. The interest in the different vintages arises because of data revisions and rebasing. Since revisions can be large and rebasing of macroeconomic data is common, the measurements of past observations published today may differ substantially from those that were actually used by the decision-makers in ‘real time’ (i.e. at the time that decisions are made). As a consequence, ex post empirical policy and forecast analyses based on currently available data can be misleading. Equally, decision-makers of today recognise that the most recently released data will also be subject to revision in the future. Interest in real-time datasets is fuelled, then, by an interest in properly understanding why decision-makers made the decisions they did at the time and by the desire to uncover and exploit any systematic patterns in revisions to make sure that current decisions are robust to data revision.

Empirical studies using real-time data have been especially helpful in explaining historical interpretations of business cycle dynamics and monetary policy decisions. For example, Amato and Swanson (2001) find that the predictive content of money for output in the US disappears when real time data are considered; Garratt, Koop, Mise and Vahey (2007) emphasise the need to use real time data in the analysis of the UK’s monetary targeting regime of the eighties; Diebold and Rudebusch (1991) find a substantial deterioration of forecasting performance of US output using a composite leading index in real time in place of post-revision data; and so on. The estimation and analysis of monetary policy rules have also been the subject of real-time data applications, showing that past policy decisions would have been substantially different if the final-vintage data had been available at the time (see, among others, Orphanides, 2001 and Croushore and Evans, 2006). A large part of this change relates to the central role played by the output gap in such rules and the fact that potential output estimates are particularly sensitive to revisions; see Orphanides et al (2000; US), Orphanides and van Norden [OvN] (2002;

US), Rünstler (2002; Euro area), Nelson and Nikolov (2003; UK) and Garratt, Lee, Mise and Shields [GLMS] (2006; US) and (2007; UK).

Real-time data are also important in the generation and use of forecasts. Specifically, real-time data sets provide important insights in forecasting exercises on the role of the timing of data releases, focusing attention on publication lags, the role of data revisions and end-of-sample issues. These features of the data are frequently ignored in standard analyses of final-vintage datasets but can be crucial in using forecasts in decision-making where access to the latest information is essential. Indeed, there is a burgeoning literature focusing on the “nowcasting” of macroeconomic magnitudes and events, the purpose of which is to provide today’s decision-makers with the most up-to-date statement on the current state of the economy before the data on today’s position is actually released (let alone revised); see Giannone *et al.* (2004, 2006). In a similar vein, the improved availability of real-time data since 2002 provided impetus to UK research to quantify the uncertainties surrounding data revisions (see, for example, Garratt and Vahey, 2006, Patterson, 2003, and Mitchell, 2004, Garratt, Lee and Vahey, 2008). Following this, the Bank of England now not only provides an indication of the uncertainty surrounding forecasts of future inflation and output growth through their published fan-charts, but also illustrate there the uncertainty surrounding the measures of the series over the recent *past* to highlight the potential effects of data revisions.

Research on the impact of revisions has also been undertaken on Australian data. For example, Lim (1985), Brookes *et al.* (1998), and Bajada (2002) considered Australian GDP data and found systematic patterns in the revisions which would impact on decision-making if, as suggested by Bajada, the markets act on the preliminary estimates as though they were final and complete. Gruen *et al.* (2005) considered the effect of revisions of GDP data on estimates of the output gap and de Brouwer and Gilbert (2005) explicitly considered the importance of the use of real time data in estimating monetary policy reaction functions and explaining Australian monetary policy decisions. And, as elsewhere, the Australian media pay considerable attention to revisions published by the Australian Bureau of Statistics [ABS] since they are believed to have an important influence on the timing and conduct of monetary policy by the Reserve Bank of Australia [RBA].

This paper describes the construction of a real-time macroeconomic database for Australia which includes measures of GDP, its components, prices, and key monetary and labour market statistics over the last fifty years as published and revised in real time. The vintages of data are collated from various sources and accommodate multiple definitional changes, providing a comprehensive description of the macroeconomic environment as actually experienced by Australian policy- and decision-makers at the time decisions are made. The database is available through the University of Melbourne ([http://www.economics.unimelb.edu.au/Real-Time\\_Macroeconomic\\_Database\\_for\\_Australia/Home.html](http://www.economics.unimelb.edu.au/Real-Time_Macroeconomic_Database_for_Australia/Home.html)) and the ABS websites along with a manual describing the sources and definitions of the series in more detail than is possible in this paper; see Lee *et al.* (2011). The description of the database provided here highlights the complexities in drawing inferences and in decision-making on the basis of macroeconomic data that are subject to revision. For this reason, we also provide an illustrative example of how the data can be manipulated to deal with and to exploit the richness of the real-time datasets. The example focuses on the calculation and representation of the Australian output gap, making use of methods that jointly model the growth of output and revisions in its measurement and that can accurately convey the uncertainties associated with current and future output levels when data are published subject to revision.

## **2 An Overview of the Australian Real-Time Database**

### **2.1 Structure and Content**

The Australian Real-Time Database includes a total of 36 variables including real and nominal measures of key macroeconomic variables, typically in their original form and following a seasonal adjustment (SA), as reported in real time. The database focuses on macroeconomic variables as these are susceptible to revision (unlike financial variables such as interest rates, exchange rates, etc.), although real time measures of various monetary aggregates are also provided. Most of the series were originally published by the Australian Bureau of Statistics (ABS) in the form of hardcopy, microfiche, CD ROMs and, more recently, Excel workbooks accessible via the ABS's website. The monetary variables were

published by the Reserve Bank of Australia (RBA) in its Monthly Bulletin until 2010, and are now made available online through the RBA website.

The series in the database can be grouped into three blocks according to their focus and the frequency of their measurement:

1. GDP and its components, measured at a quarterly frequency and published quarterly

- (a)  $Y_t$  Gross Domestic Product - original, current price SA, real SA
- (b)  $C_t^{PR}$  Household Final Consumption Expenditure - original, current price SA, real SA
- (c)  $C_t^G$  General Government Final Consumption Expenditure - original, current price SA, real SA
- (d)  $I_t^{FCF}$  Gross Fixed Capital Formation - original, current price SA, real SA
- (e)  $I_t^{INV}$  Changes in Inventories - original, current price SA, real SA
- (f)  $X_t$  Exports of Goods and Services - original, current price SA, real SA
- (g)  $X_t^*$  Imports of Goods and Services - original, current price SA, real SA;

2. Other key macroeconomic aggregates, measured at quarterly frequency and published quarterly

- (a)  $Y_t^{MAN}$  Manufacturing Production Index - real SA
- (b)  $Y_t^{IP}$  Industrial Production Index - real SA
- (c)  $B_t$  Balance of Payments - original, SA
- (d)  $P_t^{GDP}$  GDP Implicit Price Deflator - SA
- (e)  $P_t^{CPI}$  Consumer Price Index - original

3. Labour market and monetary aggregates, measured at monthly frequency and published monthly

- (a)  $E_t$  Employed Persons, Total Persons - original, SA
- (b)  $U_t$  Unemployment Rate, Persons - original, SA
- (c)  $M_{0t}^C$  Currency - original
- (d)  $M_{0t}^D$  Current Deposits - original
- (e)  $M_{1t}$  M1 - SA
- (f)  $M_{3t}$  M3 - SA
- (g)  $M_{4t}$  Broad Money - SA
- (h)  $M_t$  Money Base - original

As noted in the list, some of these series are available in more than one form. The GDP and component series are available in their original form (i.e. measured in nominal terms and without SA), in seasonally-adjusted current price form (i.e. nominal measures with SA) and in seasonally-adjusted, constant price or chain volume form (i.e. measured in real terms with SA). The employment and unemployment figures are also available with and without seasonal adjustment. The other series are only available in the form explicitly described in the list above. Data availability also differs from series to series. For example, for the SA real measures of GDP and its components, the first available vintage of data is typically from the early seventies and covers observations back to the late fifties. Most of the other series typically have shorter sample periods however. More complete details on the coverage of the data is provided in Table 1 and in the Database Manual.

The variables in the database are presented in a common format. There is an Excel workbook for each of the 36 variables containing a summary of the details of the data (source, definition, etc.) on the first sheet and the raw data in a second sheet. If we denote a variable  $z$  at time  $t - s$  by  $z_{t-s}$  and the measure of this magnitude as published in time  $t$  by  ${}_t z_{t-s}$ , then the time- $t$  vintage of data typically includes the observations  ${}_t z_1, {}_t z_2, \dots, {}_t z_{t-1}$ . This runs from the start of the sample to  $t - 1$  because there is usually a one period delay in the release of data. The raw data presented in the second sheet is in the form of “data triangles” where each column of data relates to a data vintage so that the successive columns grow longer each period to give a triangular shape to the dataset. The rows show the published measure for the same observation at different vintages so the revisions to a particular observation can be tracked by looking horizontally across the spreadsheet. Where a series is available in more than one form (with SA and without SA, for example), the series are provided in separate workbooks.<sup>1</sup>

---

<sup>1</sup>The database also includes three tables providing information on gross value added and changes in inventories (current price and chain volume measures) disaggregated by industry. However, the first available vintage of this data is dated 2001q4 so that it is not yet easily used in time series analysis.

## 2.2 Changes in Series over Time

Real-time data are useful because the measures of a series change over time. Analyses of past decisions should be conducted on the data that were available at the time and current decisions should take into account the fact that the data currently available are subject to change. However, the measurement of a variable might change over time for various reasons and it can be useful to draw a distinction between two broad categories of change: namely, “definitional changes” and “revisions”.

Definitional changes tend to be larger once-and-for-all shifts in a series and reflect a change in the way a concept is conceived. For example, an analysis of the role played by a particular group of individuals in the labour market may prompt a change in the definition of unemployment so that the group is included among the measured unemployed where previously it had been excluded or vice-versa. The change in definition would show as a discrete shift in the next vintage of the unemployment data as the historical data are amended to take into account the retrospective reallocation of the group of individuals. Another example is the “rebasings” of the constant price measures of a series. In a constant price measure of output, for example, the prices of different goods and services are fixed at a base year and these prices provide the weights used to combine the outputs of the different goods and services in different years. However, over time, the relative price of goods change (consider the price of computers relative to other goods and services, for example), the original base year prices become inappropriate and a new set of base year weights are used. This change effectively re-defines the output concept, focusing on a different basket of goods, and introduces a discrete shift in the series as the next vintage of data applies the new weights to the historical data.

Definitional changes typically occur only periodically and the timing and nature of the changes are well-documented. This means that their effect can usually be taken into account in a straightforward way if an investigator wishes to produce a historical series that abstracts from the definitional changes. For example, depending on the nature of the re-definition, its effects might be captured by a simple scaling of the pre-change data by some additive or multiplicative factor. The effects become more difficult to model if the re-definitions take place very regularly though. This is the case where chain volume

indices are used, for example. Here, the series are re-based every year by expressing the value of the aggregate in each pair of consecutive years in the prices of the earlier year. The year-to-year indexes are then compounded to form a long, continuous time series which is referenced to the current price value of the latest base year. The procedure is appropriate if relative price movements have become rapid so that constant price measures based on a fixed price profile become inappropriate very quickly. However, the chain volume index variable is effectively re-based and re-defined every year so that it becomes more difficult to separately identify the effects of the rebasing from other influences when considering different vintages of the same series.

Data “revisions” based on the arrival of new information are different in nature to definitional changes in a series. Collating information, through the collection of surveys, activity indicators and other source data, is a time-consuming process. Some of the more reliable information is available only infrequently (from annual surveys, say) and the information from higher frequency but less reliable data sources has to be reconciled with this. The production of measures therefore involves a judgement in using unadjusted source data against ‘modelled’ data that anticipates the arrival of additional and/or more reliable data. The collation of information and its analysis means that, in Australia, the first release of data on a quarter is typically released only in the last month of the subsequent quarter. Moreover, even after this delay, some of the provisional information on which the series are based will be updated so that the measures are revised. It would be reasonable to assume that any systematic discrepancies between the provisional and the more reliable information that are observed over time may prompt an investigation by the data collecting authority to understand how the collection and processing of data might be improved. It should be noted, however, that the process of measuring a variable typically aims to produce a coherent and timely measure of the concept of interest using the information that is currently available, not necessarily to eliminate revisions. Revisions of measures can continue for some time, therefore, and they can be relatively large and they can contain systematic content. It is for these reasons that the analysis of revisions in real-time data is potentially so revealing.

### 2.3 Definitional Changes and Revisions in Australian Data

Figures 1 and 2 illustrate the effects of data definitions and revisions and their order of magnitudes in the Australian real output data. Figure 1 shows the measure of the level of real GDP (\$m) for the 1974q3 observation as reported over the successive vintages 1974q4 – 2010q4. The figure shows clear discrete shifts in the series in 1978q1, 1982q2, 1988q2, 1993q1 and 1998q4 corresponding to the change in base weights in these years, plus the gradual evolution of the series from 1998q4 reflecting the effects of the shift to measuring output with chain volume indices from that time.<sup>2,3</sup> Figure 2 shows measures of annualised output growth as reported in eight vintages published at five yearly intervals from 1975q1 – 2010q1 with the crosses indicating the final observation for the relevant vintage. Each growth measure is calculated using the definition relating to that vintage so the effects of the definitional changes are eliminated and the plots provide a better sense of the effects of revisions. The plots show that, abstracting from definitional changes, there have been some very large changes in measured growth, revised by up to 5%, over a ten-twenty year period.

This latter point is made even more clearly in Figure 3 which focuses on the quarterly growth of real GDP over a particular period at the beginning of our sample; namely 1974q2-1975q4. This figure illustrates how the growth measures have changed over time as reported in the successive data vintages from 1976q4 onwards. The figures show enormous variability in the measures reported in the various vintages. For example, growth in 1974q3 was initially measured to be around -3% or -4% in the vintages published soon after 1974q3. But the measure rose over time so that it was judged to be positive and in the region of +1% through the vintages of the early eighties, and is considered to be closer to +2% in the most recent data vintages. Some part of the variability is explained as output appears to have been reallocated from one quarter to another (so growth is revised down in one period and up in the next). But this is certainly not the whole explanation and a close examination of particular periods shows some dramatic examples of how the

---

<sup>2</sup>There were two releases of the national accounts in 1988q2 each corresponding to a different base year. Figure 1 reports the second of these releases measured using weights based on prices in 1984 – 1985.

<sup>3</sup>The small discrete shift in 2009q4 is due to the implementation of new international standards.

interpretation of business cycle conditions can change simply because of data revisions. For example, defining a recession as two successive periods of negative quarterly growth, we find the recession observed to take place in 1974q2/q3 according to the data vintages published at the time had simply disappeared according to the data published after 1979. Equally, a similarly-defined recession is shown to take place in 1975q3/q4 according to the vintages published up to 1992q4 and in the most recently published data, but did not take place according to data published between 1993q2 and 2002q2.

A time series analysis of the real-time data will require a stance to be taken on the treatment of the changes in the series due to definitional changes and those due to revision. One straightforward approach is to assume that the effects of a definitional change can be captured by the joint effect of an additive and a multiplicative factor. The size of the factors can be estimated by regressing the (logarithm of the) post-change series on an intercept and the (logarithm of the) pre-change series over the period for which both series are observed. Running these regressions for each vintage and applying the scaling factors provides a dataset that abstracts from the definitional changes arising from the effects of simple additive and multiplicative factors and means that all the vintages of the variable are expressed using the same unit of measurement as the final vintage. Tables 2a and 2b provide some summary statistics for the variables obtained abstracting from the effect of definitional changes in this way. They show that the effects of revisions are relatively large in every variable (although close to zero for most series on average). So, for real GDP for example, the average revision over the eight quarters following the first release of data on  $z_t$  is just -0.52 percentage points. However, with a standard deviation of 0.0112, the observed revisions in the quarterly growth rate typically lie in the range -2.7 and +1.7 percentage points (assuming the majority of revisions lie within two standard deviation of this mean). This is obviously a large range relative to the quarterly growth rate itself and a succession of similarly signed revisions could have a very large effect on the annual growth rate. The size of the revisions beyond the first quarter gradually falls, but the revisions are still reasonably large even after six quarters (with  $z_{t+6} - z_{t+9}$  having a standard deviation of 0.0055), reflecting again the effect of continued revision

captured in Figure 3.<sup>4</sup> Similar comments could also be made regarding all of the other macro variables covered in Tables 2a – 2b showing that econometric analyses that fail to take into account revisions could be seriously misspecified and conclusions drawn could be very misleading.

### **3 A Real-Time Analysis of the Australian Output Gap**

The importance of the use of real-time data is best conveyed by looking at a specific issue and in this section we discuss the measurement of the Australian output gap (OG) as obtained in real time and compared to the final vintage measurements that are typically reported. The analysis follows that of OvN and GLMS which demonstrated that the use of real-time data has important implications for the measurement of the output gap in the US and the UK and it elaborates on the analyses of the Australian gap in Gruen et al. (2005) and de Brouwer and Gilbert (2005). In the sections below, we first describe some of the issues that arise in measuring and representing the output gap in real time. We then apply the techniques used to deal with these issues to the Australian real GDP data. The data that are used are adjusted to take into account definitional changes as described in the previous section so that the analysis focuses on the role of revisions only.

#### **3.1 Measuring and Representing the Output Gap in Real Time**

In the presence of data revisions, the output gap at time  $t$  is the difference between the post-revision measure of output at time  $t$  and the trend output level at  $t$  obtained by applying a chosen de-trending technique to the post-revision output series. So, denoting output by  $y_t$  and assuming that output is published with a one period delay and then

---

<sup>4</sup>The time frame over which changes occur might also be used to distinguish ‘definitional changes’ from ‘revisions’. So, for example, these statistics suggest that changes in measures that occur over the first two years following the first release might reasonably be considered as ‘revisions’ while changes that appear after longer periods might be judged the outcome of a retrospective change in the concept of interest and classed as ‘definitional’. This type of definitional change could be accommodated by a further intercept shift that aligns earlier vintages with the most recent vintage according to the first observation prior to the time frame assumed to reflect revisions.

revised  $q$  times, the post-revision output series available at time  $T$  is given by  $\{ {}_{q+2}y_1, {}_{q+3}y_2, \dots, {}_T y_{T-q-1} \}$ . For output measures between  $T - q - 1$  and  $T$ , and indeed into the future beyond  $T$ , a model is required to forecast the revisions to obtain the post-revision series. Trend output levels calculated using the data at time  $T$  applies the chosen de-trending technique to the post-revision series for output upto  $T - q - 1$  augmented with forecasts of the post-revision series at  $T - q$  and beyond. In what follows, the trend obtained in this way is denoted  ${}_T \tilde{y}_t$  for  $t = \dots, T - q, \dots, T, \dots$ , and the gap measure is denoted  ${}_T x_t$ . Of course, there is considerable controversy on the choice of technique used to define the output trend on which the gap measure is based; see the discussion of OvN, GLMS and Garratt *et al.* (2011). In what follows, we consider gaps based on trends calculated by applying the popular Hodrick-Prescott (HP) filter to the output series. By focusing on only the HP filter throughout, we side-step the discussion on how the trend should be measured allowing us to highlight the role played by the real-time data in measuring the gap.<sup>5</sup>

### 3.1.1 Forecasting post-revision outcomes

Working with post-revision series clearly requires a model with which to derive the forecasts. One possibility is to estimate a simple autoregressive model using the most recent vintage of data (a univariate approach). However, this does not take into account any information contained in the revisions of the series across past vintages. GLMS therefore suggest using a more sophisticated (multivariate) approach which models the growth in the first-release data and the first  $q$  revisions of the data. This is done straightforwardly by estimating a vector autoregressive (VAR) model of order  $p$ :

$$\mathbf{z}_t = \mathbf{a} - \mathbf{B}_1 \mathbf{z}_{t-1} - \dots - \mathbf{B}_p \mathbf{z}_{t-p} + \mathbf{e}_t \quad (3.1)$$

---

<sup>5</sup>The application of the detrending technique to a forecast-augmented series also mitigates against the ‘end-of-sample problem’ that arises when a two-sided filter is used. These filters become progressively more one-sided as the end of a finite sample is approached and this can introduce measurement problems even in the absence of data revisions. Burnham (1980) and Stock and Watson (1989), among others, suggest the use of forecast-augmented series so that the two-sided filter can be applied even at the end of the sample, and GLMS show that this is a very effective procedure when measuring the output gap in the US and UK.

for  $t = 1, 2, \dots, T$ , and where  $\mathbf{z}_t = ({}_t y_{t-1} - {}_{t-1} y_{t-2}, {}_t y_{t-2} - {}_{t-1} y_{t-2}, \dots, {}_t y_{t-q-1} - {}_{t-1} y_{t-q-1})'$ ; that is,  $\mathbf{z}_t$  is a vector containing the growth in the first-release data,  ${}_t y_{t-1} - {}_{t-1} y_{t-2}$ , and the time- $t$  revisions on output in the previous  $q$  periods,  ${}_t y_{t-i} - {}_{t-1} y_{t-i}$ ,  $i = 1, \dots, q$ .<sup>6</sup> The VAR model provides a straightforward time series representation of the variables in  $\mathbf{z}_t$  regressing each series in turn on  $p$  lags of all  $q$  variables. A VAR of this form is appropriate on the reasonably uncontentious assumptions that actual (i.e. post-revision) output is first-difference stationary and that revisions are stationary.

Having estimated this model using data available at  $T$ , say, forecasts can be obtained of the post-revision series upto and beyond  $T$ . The application of the HP filter to the forecast-augmented post-revision series can be used to provide time- $T$  measures of the past, contemporaneous and future gap  ${}_T x_{t+s}$ ,  $s = \dots, -1, 0, +1, \dots$ . These measures are straightforward to calculate, are based on past and expected future output growth and fully take into account all the data revisions that are expected to take place. On the other hand, the simple statistics obtained in this way do not convey the estimation and measurement uncertainties associated with the output gap measures. These are potentially significant here given that forecasts of the revised and unrevised series are used in various different ways in the construction of the measure.

### 3.1.2 *Density forecasts and event probability forecasts*

The uncertainties surrounding the measures can be described in a relatively straightforward way if we choose to represent the output gap through estimates of its probability density function (pdf) rather than through simple point estimates. Indeed, providing a richer probabilistic description of the output gap not only helps in conveying the uncertainties associated with the gap more clearly but also allows us to provide statements on the likely occurrence of specified events that involve the gap and that may be of interest to particular decision-makers. For example, in monetary policy decisions, it might be more realistic to assume that the monetary authority is concerned not just with the point estimates of the gap but with ‘booms’ and ‘recessions’ (i.e. whether the output gap is

---

<sup>6</sup>In fact, the model assumes simply that there is systematic content in upto  $q$  revisions. Data might continue to be revised but there is no further information to be exploited.

positive or negative over some period, irrespective of size), say, or with whether conditions are improving or deteriorating (i.e. with the gap rising or falling).<sup>7</sup> In this case, interest focuses on joint events involving the gap in successive periods, and the events are unlikely to be easily inferred from point estimates of the gap. Rather, direct statements of the likelihood of the probability of these events will be helpful and probabilistic representations of the output gap will be required.

Event probability forecasts and pdf's of this sort are straightforward to calculate using simulation methods so long as the underlying data generating process is relatively simple. This is the case with the VAR model in (3.1). Garratt *et al.* (2003) provide a detailed description of the methods, but the ideas are simple to explain. For example, assume the estimated version of the model of (3.1), denoted  $M_q$ , is the true data generating process. Then one can use this estimated model, based on the observed data  $\{\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_t\}$ , to generate  $S$  replications of the future vintages of data,  $\{\hat{\mathbf{z}}_{t+1}^{(s)}, \hat{\mathbf{z}}_{t+2}^{(s)}, \dots, \hat{\mathbf{z}}_{t+H}^{(s)}\}$  for  $s = 1, \dots, S$ , on the assumption that the model continues to hold over the forecast horizon  $t+1, \dots, t+H$ . These  $S$  simulated future vectors of variables directly describe the likelihood of observing the various values of  $\mathbf{z}_{t+1}, \mathbf{z}_{t+2}, \dots, \mathbf{z}_{t+H}$  conditional on the observations available at the end of period  $t$  and on model  $M_q$ . In particular, the simulations generate values of the forecast post-revision output levels, so that we can generate measures of trend output and the output gap in each simulation. The simulated distribution of the output gap measures obtained in this way is the estimated pdf. Equally, counting the number of times an event occurs in these simulations provides a forecast of the probability that the event will take place; for example, if a recession is defined to occur when a negative gap is observed in say four consecutive periods, then simply noting the fraction of the simulations in which this occurs provides a real time estimate of the forecast probability that there is a recession.

The simulation exercise described above can be extended to accommodate parameter uncertainty for the given model by carrying out an additional iteration of the simulation procedure in which replications of the historical data and of the model parameters are also produced (see Garratt *et al.* (2003) for more details). And the process can also

---

<sup>7</sup>See Svensson (2001, 2002), Cukierman and Gerlach (2003), Walsh (2003) and Lee and Shields (2011) for further discussion of the form of monetary authorities' objective functions.

be extended to take into account the uncertainty surrounding the choice of forecasting model by conducting the exercise with many alternative models and amalgamating the simulations from the models with appropriate weights.<sup>8</sup> In every case, the process can be conducted recursively moving through the sample of data, up to the most recently-available data, so that final vintage estimates of the pdf and probability forecasts involving the gap can be simulated too. In this way, a complete characterisation of the output gap measure can be obtained, accommodating the various elements of data and estimation uncertainty.

## 3.2 Australian Real-Time Output Gap Measures

### 3.2.1 *The unreliability of gap measure obtained in real-time*

Figure 4 shows three output gap series obtained using the Australian real time database measures of real output. The first vintage of output data is dated in 1971q4 and this runs from 1959q3 – 1971q3 while the final vintage is dated  $T = 2010q4$ , running to 2010q3. The ‘*Final OG*’ series,  ${}_T x_t^{fo}$ , is the gap obtained by applying the HP filter to the 2010q4 vintage of data and subtracting the trend from the observations on output as provided by the final vintage. This is the series that would usually be used by an econometrician studying the output gap in 2010q4 paying no attention to the earlier vintages, and ignoring the forecast-augmentation issues raise above. The ‘*OvN Real-Time OG*’ series,  ${}_t x_{t-1}^{ro}$ , was introduced in OvN and uses the successive vintages of data, applying the HP filter to each data vintage in turn and calculating the difference between the trend and actual output series of the vintage to obtain the output gap for the *final observation* covered by that vintage. This series is constructed as if we ask the time- $t$  econometrician to produce an output gap series based on her most up-to-date vintage- $t$  data and record only the final observation (dated at  $t - 1$  given the one-period delay in data releases). The ‘*Forecast-Augmented Real-Time*’ *OG*’ series,  ${}_t x_{t-1}^{ru}$ , is obtained in a similar way, but here the time- $t$  econometrician estimates a univariate autoregressive model of order eight for output growth based on vintage- $t$  data and produces forecasts of future output growth beyond

---

<sup>8</sup>This entails a ‘model averaging’ approach as explained in Draper (1995), Hoeting et al. (1999) and Garratt et al. (2003), and illustrated in the empirical work below.

the end of the vintage- $t$  sample period. The trend measure is then obtained applying the HP filter to the forecast-augmented series to take into account the end-of-sample issues discussed above. Again, it is the gap measure for the final observation covered by that vintage which is recorded.

The figure illustrates the considerable differences between the real time measures and the final vintage measure of the gap arising out of data revisions and the end-of-sample effects. There are occasions where the gap measures differ by as much as 3–5 percentage points, which is very large given that the gaps themselves have standard deviation of around 1 percentage point. Table 3 shows that the correlation between the real time and the final gap measures following the OvN procedure is just 0.493 and that the two measures agree on whether the output level is above or below the trend in only 65.0% of the sample period - not impressive when one realises that there would be 50% agreement if the series were entirely uncorrelated. If the final output gap series is considered to be the best indicator of the true output gap available, these results accord with OvN's and GLMS's conclusions for the US and UK that real time output gap measures are very unreliable. This means that it is important to use all the available methods for dealing with revisions in calculating the gap measures and to properly take into account the uncertainty surrounding the measures before they are used in policy decisions.

The table also shows the effect on the gap measures of employing the (univariate) forecast-augmentation method of calculating the trends described earlier. While the augmentation has a marked impact on the variability of the output gap series, reducing the standard deviation of the real time output gap measure from 0.14 to 0.10, it has relatively little impact on the gap as it compares to the Final OG measure. The correlation between the  ${}_t x_{t-1}^{ru}$  and  ${}_T x_{t-1}^{fo}$  rises only slightly to 0.499 with the agreement on the occurrence of booms and recessions remaining in the low 60%'s. This contrasts with the US and UK cases where the effect of the forecast-augmentation is more pronounced even in this univariate exercise which takes no account of any systematic information contained in the revisions data.

### 3.2.2 Exploiting the revision data

The information contained in the real time dataset is only fully exploited if any systematic patterns in revisions are identified and used in constructing the gap measure. We make use of models of the form in (3.1) to do this although in what follows we acknowledge that there is ambiguity on the precise model that should be used and that this model uncertainty contributes to the uncertainty surrounding the gap measure itself. In deciding on the appropriate multivariate model, our *a priori* view was that the revision process for Australia is protracted and complex so that a relatively sophisticated model of the data might be required (i.e. large  $p$  and large  $q$  in (3.1)). We also recognise that the process might change over time so that the model used should be re-estimated in each period. To accommodate our uncertainty on the appropriate model, we consider a set of nine distinct models for Australian output growth and revisions at each point in our sample. The nine models are defined according to the revision horizon; i.e. the quarter after which no further revision is assumed to take place or, if revisions occur, after which the revisions have no systematic content. Hence, the models each take the form given in (3.1) with  $q = 0, \dots, 8$ , and are denoted  $M_{0t}$  to  $M_{8t}$  respectively estimated on the vintages of data up to and including  $\mathbf{z}_t$ ,  $t = 1, \dots, T$ . Model  $M_{0t}$  represents the model in which no revision data are included while model  $M_{8t}$  allows for a protracted revision process of up to two years after the first release of data. The maximum lag length considered in each model is  $p = 4$  but, to deal with potential over-parameterisation, we also obtain a set of restricted models following a specification search on each of  $M_{0t}$  to  $M_{8t}$ . In this search, we impose a zero coefficient restriction on the variable with the smallest (absolute-valued) t-ratio in each model until all the remaining variables have t-ratios of 1.25 or more.

In addition to the nine distinct models  $M_{0t} - M_{8t}$ , an ‘aggregated’ counterpart  $\overline{M}_t$  can be constructed as a weighted average of the individual models. So, for example, having estimated the individual models and generated simulated futures for each model, the simulations can be pooled to provide a single pdf which takes into account the uncertainty across models. This pooling can give different weights to the simulations from the different models according to the fit of the models in the spirit of Bayesian model averaging

methods. Draper (1995) for example suggests using weights  $w_{q,t}$  given by

$$w_{q,t} = \frac{\exp(SBC_{q,t}^*)}{\sum_{j=1}^9 \exp(SBC_{j,t}^*)}, \quad q = 0, \dots, 8, \quad (3.2)$$

where  $SBC_{q,t}^* = SBC_{q,t} - \max_j(SBC_{j,t})$  and  $SBC_{q,t}$  is the Schwarz Bayesian information criterion based on the maximized value of the log-likelihood function for model  $M_{qt}$ .<sup>9</sup> The density forecasts that convey the uncertainty surrounding a gap measure could change over time even if we knew the number of revisions that take place and restricted attention to a particular revision horizon,  $r$ , since model  $M_{rt}$  is re-estimated at each period. But the use of the weighted average of the models provides a further source of time variation as the revision horizon, and weights, could also change over time.

The largest model we consider at each time,  $M_{8t}$ , includes 37 ( $= 1 + 9 \times 4$ ) explanatory variables in each of the unrestricted equations and, to ensure reasonable degrees of freedom, we only considered models based on at least 90 data points. Since the first data vintage available is dated 1971q4, this means our multivariate analysis of the output data produces real time measures of the gap from 1994q1-2010q3 (=67 observations in total). To illustrate the relative levels of support for the various alternative multivariate models, Table 4 reports the average SBC weights for the nine models based on (3.2), calculated over the 67 recursions 1994q1 – 2010q3. According to these averages, model  $M_{0t}$  has the largest support, but the remaining models  $M_{1t} - M_{8t}$  account for 61% of the weight reflecting the importance of the revisions data. Figure 5 illustrates that, in fact, the support for the alternative models changes quite substantially over time and that support for model  $M_{0t}$  declines over time with increasing weights placed on models  $M_{1t}$ ,  $M_{2t}$ ,  $M_{3t}$  and  $M_{4t}$ . Models  $M_{5t} - M_{8t}$  attract relatively little weight at any time, suggesting that the important part of the systematic revisions occurs within a year of the first-release data. The broad pattern, with the balance of weight moving from the simplest model to the more complicated models could be a reflection of the shift to a relatively more complex measurement of output as detailed in section (2.2). In any case, the time-variation in the

---

<sup>9</sup>To ensure that the likelihoods of each system are comparable, which requires that all the systems provide an explanation for all nine series, model  $M_i$  ( $i = 0, \dots, 8$ ) is supplemented by the  $(i + 1)^{th}$  equation from model  $M_{i+1}$ . The SBC statistics reflect the impact of dropping the  $(i + 1)^{th}$  revisions as regressors from the equations explaining the first-release data and the  $i^{th}$  revision therefore.

weights provides a further argument for the use of the model averages in preference to any individual model since the time-varying weights can accommodate this aspect of model uncertainty.

### 3.2.3 *The Australian gap measures*

For each recursion, we are able to compute forecasts of the post-revision output series and to estimate output gap measures using any one of the models  $M_{0t} - M_{8t}$  or their averages. Table 4 provides summary statistics on the output gap measures derived on the basis of the nine individual models plus measures based on the forecasts from the *SBC*-weighted model average. For models  $M_{0t} - M_{8t}$ , the correlations between the real time measure of the gap and the corresponding final measure are in the range  $[0.298, 0.413]$  and the agreement on booms and recessions between the real time and final vintage gap measures is between 60% and 71% across all the models.<sup>10</sup> The *SBC*-weighted average performs relatively well compared to the individual series with the correlation between real time and final-vintage measures taking a value of 0.396 and a level of agreement on the sign of gaps at 72%, matching the best of the individual models in each case. Figure 6 shows output gap measures based on the selected alternative models  ${}_t x_{0t}^{rm}$ ,  ${}_t x_{8t}^{rm}$  and  ${}_t \bar{x}_t^{rm}$  denoting the real time gap measures based on model  $M_{0t}$ , on model  $M_{8t}$  and on the *SBC*-weighted-average model, respectively. The figure also plots the final output gap measure,  ${}_T \bar{x}_t^{fm}$ , obtained using the final vintage of data, augmented at the end of the sample with forecasts based on the *SBC*-weighted-average model. The figure shows that a reasonably consensual picture of the state of the macroeconomy would have been obtained in real time using any of the alternative gap measures based on the multivariate models. However the figure also shows the variability in gap measures arising out of the model uncertainty and the advantage of the aggregate real-time measure over both of the single models (as the aggregate performs reasonably well in tracking  ${}_T \bar{x}_t^{fm}$  throughout while the individual series do well only periodically).

The gap measure based on the *SBC*-weighted average model represents our preferred

---

<sup>10</sup>Note that these figures relate to the period 1994q2 – 2010q4 and so are not comparable to those in Table 3 relating to 1971q3 – 2010q4.

real-time measure therefore although the modelling exercise demonstrates and quantifies the considerable uncertainty surrounding the measure. It also shows that the gap measure obtained in real time can differ substantially from that obtained in retrospect using final vintage data; see, for example, the difference of around one percentage point between  ${}_t\bar{x}_t^{rm}$  and  ${}_T\bar{x}_t^{fm}$  during 2007 and 2008. A complete representation of the real-time gap measure should convey accurately the uncertainties surrounding the measure therefore. The density forecasts and event probabilities of the following section serve this purpose.

#### 4 Event Probability Forecasts based on the Output Gap

In Section 3.1, we noted that information on the size and the precision of measures of the gap can be conveyed directly through the use of pdf's of the gap measured at different forecast horizons and, using these, through the use of forecasts of the probability of the likelihood of specified events involving the output gap. Figure 7 provides an illustration of the corresponding cumulative distribution functions (cdf's) that are obtained using the methods described in Section 3.1 showing plots relating to measures of the output gap in 2007q2, 2008q1 and 2009q1 (based on the SBC-average multivariate model) as were calculable using real-time data in 2008q1.<sup>11</sup> The cdf for 2007q2 relates to a measure of the gap experienced a year earlier, the 2008q1 cdf is concerned with the 'nowcast' of the gap measure and the 2009q1 cdf relates to a one-year ahead forecast.

This is a particularly interesting recent episode in RBA decision-making. The measures of annual growth, published monthly throughout 2007, ran at between 3.6% and 4.0% throughout the year.<sup>12</sup> The published minutes of the monetary policy meetings of the RBA in 2007 commented on high capacity utilisation, healthy labour markets and wage pressures in Australia and strong world growth. Comments on turbulent financial markets are made in the minutes of the final months of the year but the data on growth and the

---

<sup>11</sup>The results are based on 10,000 replications where the innovations were obtained from a multivariate normal distribution chosen to match the observed correlation of the estimated residuals for each of the models over the full sample period (a parametric bootstrap).

<sup>12</sup>These calculations, and the discussion below, are based on the data obtained abstracting from the effects of definitional changes although the commentary would be qualitatively the same if it was based on the raw data.

commentary available from the RBA minutes through 2007 suggest an economy that was unambiguously growing strongly and in which prospects for rising inflation were high. It seems surprising then that the cash rate was increased just once, by 0.25 percentage points, during the year to November. An increase in the rate of 0.25 percentage points in November 2007 was then followed by a further rise of 0.25 percentage points in February 2008 and then again in March, with the RBA minutes of these months reporting that inflation was rising, running ahead of what had been expected in late 2007, and expected to rise further in the near term.<sup>13</sup> The long period of inaction, followed by a quick succession of rises and the anxiety over the build-up of inflationary pressures at the beginning of 2008, provoked some considerable debate on the conduct of monetary policy at the time.

The plots in Figures 7 and 8 provide a potential explanation for the behaviour. Monetary policy decisions are typically based not on simple growth rates but on the output level relative to trend. When revisions are involved, estimates of the gap require forecasts of the post-revision measures of output over the recent past as well as into the future. So, for example, when the data are published with a one quarter delay and eight revisions are involved, the gap for 2007q2 involves comparing the 2008q1 forecast of  ${}_{2009q3}y_{2007q2}$  with the trend measure  ${}_{2008q1}\tilde{y}_{2007q2}$  where the latter is in turn based on post-revision forecasts of output upto and after 2007q2. The cdf's for the 2007q2 gap in Figure 7 shows that there would have been considerable uncertainty in 2008q1 about the size of the gap measure even one year earlier: while the most likely gap value was positive and in the region 0.3% according to the data available in real time, there was still an approximate 20% probability that the past gap measure had been negative. Estimated gap measures for 2008q1 and 2009q1 are still positive on average but the flatter cdf's reflect still more uncertainty surrounding these measures. These magnitudes are also shown in Figure 8 which depicts a 'fan chart' for the gap measure which emphasises both the considerable uncertainty over the past gap measures and the very rapidly growing uncertainty over the nowcast and into the future.

In the event, the accumulating evidence of a positive gap over a protracted period might have provided the evidence required to justify the cash rate hikes in early 2008.

---

<sup>13</sup>The RBA's inflation gauge was at 3.6% in March 2008, well above its preferred band of 2%-3%.

This evidence could be considered in probabilistic terms by evaluating the likelihood of the economy experiencing "excess demand", defined (arbitrarily) as being where a five-period moving average of the gap (centred around the period of interest) is positive. This phenomenon involves a very complicated joint event involving not only the forecasts of post-revision outputs to construct gap measures at various horizons but also the joint probabilities of successive gaps being positive. But the simulation methods described earlier can be used in a straightforward way (simply counting the number of times the event occurs in the repeated simulations) to obtain a forecast of the likelihood of the event occurring in real time. Figure 9 plots this probability, calculated in real time over the period 1994q2-2010q4, and set against the corresponding estimated gap measures. The diagram shows that, if this was the concept of interest to policy-makers, then the event probability forecast captures the complexity in a very straightforward way and the plot does indeed show a rise in the likelihood of excess demand in 2007, rising to around 70% by the end of the year.

The 2007/2008 experience provides a good illustration of the wealth of information contained in the real time data and the usefulness of density forecasts and event forecasts in capturing and evaluating the uncertainties surrounding decision-making in the presence of data that are subject to revision. However, the period also provides a further final illustration of the importance of the use of real-time data because of a number of significant definitional changes that occurred in the 2009q4 release of data. These were based on a move to new international standards in the system of national accounts and the adoption of a new industry classification system.<sup>14</sup> These changes resulted in revisions to the entire National Accounts time series. The changes occurred outside our assumed two year revision horizon so that, despite the use of the methods used to eliminate the effect of data redefinitions as discussed in Section 2, the changes are associated with very substantial revisions in the gap measure in 2007 and 2008. This was captured in Figure 6 where,

---

<sup>14</sup>Specifically, this vintage moved to the international standards System of National Accounts 2008 (SNA08) and Balance of Payments and International Investment Position Manual, sixth edition (BPM6); and the Australian and New Zealand Standard Industrial Classification, 2006 (ANZSIC06) and Standard Economic Sector Classification of Australia, 2008 (SESCA08) were also adopted at this time.

as noted earlier, the gap measures based on the final vintage data are approximately one percentage point higher than the real-time measures. Gap measures of this order of magnitude are consistent with the cdf's reported in Figure 6 since gaps in excess of one percentage point are in the tail but nevertheless have a non-zero probability. But the size of the gap that is apparent in the most recent vintage of data would make the inactivity of the RBA during 2007 even more difficult to understand if no attention was paid to the real-time data that describes the environment in which decisions were actually made or the uncertainty surrounding the gap measures obtained on the basis of data subject to subsequent revision.

## **5 Concluding Comments**

The analysis of Sections 3 and 4 illustrate the complexities involved in the use of real-time data in the analysis of the business cycle. The data-sets will typically require the application of techniques that can capture and reflect both the macroeconomic processes of primary interest and the measurement processes that underlie the publication of the first-release measures of the series and their subsequent revision. The multivariate modelling methods and the simulation methods used to produce gap measures and density forecasts are good examples of such techniques.

The relatively sophisticated methods required to fully exploit the real-time dataset should not obscure the power of the data though and the analysis of Sections 3 and 4 also demonstrates the enormous potential in real-time data for helping to understand policy episodes of the past and for enhancing current decision-making. This was in an analysis that focused on output alone, however, and the scope for improved macroeconomic analysis will be still greater when the real-time relationships between variables is properly explored. The Australian Real-Time Database should provide an invaluable resource in this regard.

## 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.
- Bajada, C. (2002), “The Effects of Inflation and the Business Cycle on Revisions of Macroeconomic Data”, *Australian Economic Review*, 35: 276–286.
- De Brouwer, G. and J. Gilbert. (2005), “Monetary Policy Reaction Functions in Australia”, *Economic Record*, 81, 124–134.
- Brooks, R.D, G.E. Bermana and S.R. Davidson (1998), “The Nature and Extent of Revisions to Australian Macroeconomic Data”, *Applied Economic Letters*, 5, 3, 169-174.
- Croushore, D. and C. Evans (2006), “Data Revisions and the Identification of Monetary Policy Shocks”, *Journal of Monetary Economics*, 53, 1135-1160.
- 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.
- Draper, D. (1995), “Assessment and Propagation of Model Uncertainty,” *Journal of Royal Statistical Society Series B*, 58, 45-97.
- Garratt, A., G. Koop, E. Mise and S. Vahey (2007), “Real Time Prediction with UK Monetary Aggregates in the Presence of Model Uncertainty”, Downloadable at <http://www.ems.bbk.ac.uk/faculty/garratt/index.html>.
- Garratt, A., K. Lee, E. Mise and K. Shields (2008), “Real Time Representations of the Output Gap”, *Review of Economics and Statistics*, 90, 4, 792-804.
- Garratt, A., K. Lee, M.H. Pesaran and Y. Shin (2003), “Forecast Uncertainty in Macroeconometric Modelling: An Application to the UK Economy”, *Journal of the American Statistical Association*, 98, 464, 829-838.

- Garratt, A., K. Lee, and K. Shields (2011), “Measuring the Natural Output Gap using Actual and Expected Output Data”, *University of Nottingham Discussion Paper*, March 2011.
- Garratt, A., K. Lee, and S. Vahey (2008), “Real Time Probability Forecasts of UK Macro Events”, *National Institute Economic Review*, 203, 78-90.
- Garratt, A. and S. Vahey (2006), “UK Real-Time Macro Data Characteristics” *Economic Journal*, 2006,116, F119-F135.
- Giannone, D., L. Reichlin and L. Sala, (2004), “Monetary Policy in Real Time” in M. Gertler and K. Rogoff (eds.) *NBER Macroeconomics Annual*, pp. 161-200.
- Giannone, D., L. Reichlin and L. Small, (2006), “Nowcasting GDP and Inflation: The Real-Time Informational Content of Macroeconomic Data Releases” *ECB Working Paper 633*.
- Gruen, D., T. Robinson. and A. Stone (2005), “Output Gaps In Real Time: How Reliable Are They?”. *Economic Record*, 81: 6–18.
- Hoeting, J. A., Madigan, D., Raftery, A. E. and Volinsky, C. T. (1999). “Bayesian Model Averaging: A Tutorial”, *Statistical Science*, 14, 382-417.
- Lee, K.C., N. Olekalns, K. Shields and Z. Wang (2011), *The Australian Real-Time Database: A Users Manual*, available at [www.unimelb.edu.au/.../.../](http://www.unimelb.edu.au/.../.../), April 2011.
- Lee, K.C. and K. Shields (2011), “Decision-Making in Hard Times: What is a Recession, Why Do We Care and When Do We Know We Are in One?”, *North American Journal of Economics and Finance*, 22, 1, 43-60.
- Lim, G.C. (1985), “GDP Growth Rates Calculated from Quarterly National Accounts, Discrepancies and Revisions”, *Australian Economic Review*, 4, 21-7.
- Mitchell, J. (2004) *Revisions to Economic Statistics*, National Institute of Economic and Social Research, April.

Orphanides, A. (2001), “Monetary Policy Rules Based on Real-Time Data”, *American Economic Review*, 91, 4, 964-985.

Orphanides, A. (2003), “Historical Monetary Policy Analysis and the Taylor Rule”, *Journal of Monetary Economics*,

Orphanides, A. and S. van Norden (2002), “The Unreliability of Output Gap Estimates in Real Time”, *Review of Economics and Statistics*, 84, 4, 569-83.

Orphanides, A., R.D. Porter, D. Refsneider, 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.

Nelson, E.T. and K. Nikolov (2003) “UK Inflation in the 1970s and 1980s: the Role of Output Gap Mismeasurements”, *Journal of Economics and Business*, 55, 353-370.

Patterson, K.D. (2003) “Exploiting Information in Vintages of Time-series Data”, *International Journal of Forecasting*, 19, 177-197.

Rünstler, G. (2002) “The Information Content of Real-time Output Gap Estimates: An Application to the Euro Area”, ECB Working Paper no 182, September.































