OUTPUT GAPS IN REAL TIME: THE GOOD, THE BAD AND THE UGLY

David Gruen, Tim Robinson and Andrew Stone

Research Discussion Paper
2002-??

August 2002

Economic Research Department
Reserve Bank of Australia

The authors are grateful to colleagues at the Reserve Bank, and to Richard Dennis, Simon van Norden, Athanasios Orphanides, and Adrian Pagan, for helpful comments. Responsibility for any remaining errors rests with the authors. The views expressed in this paper are those of the authors and should not be attributed to the Reserve Bank.

Abstract

The output gap – the difference between actual and potential output – is widely regarded as a useful guide to future inflationary pressures, as well as an important indicator of the state of the economy in its own right. Since the output gap is unobservable, however, its estimation is prone to error, particularly in real time. Errors result both from revisions to the underlying data, as well as from end-point problems that are endemic to econometric procedures used to estimate output gaps. These problems reduce the reliability of output gaps estimated in real time, and lead to questions about their usefulness.

We examine 121 vintages of Australian GDP data to assess the seriousness of these problems. Our study, which is the first to address these issues using Australian data, is of interest for the method we use to obtain real-time output-gap estimates. Over the past twenty-eight years, our real-time output-gap estimates show no apparent bias, when compared with final output gap estimates derived with the benefit of hindsight using the latest available data. Furthermore, the root-mean-square difference between the real-time and final output-gap series is less than 2 percentage points, and the correlation between them is over 0.8. Our general conclusion is that quite good estimates of the output gap can be generated in real time, provided a sufficiently flexible and robust approach is used to obtain them.

JEL Classification Numbers: E21, E22, E23

Keywords: Output gaps, real-time data, monetary policy

Table of Contents

OUTPUT GAPS IN REAL TIME: THE GOOD, THE BAD AND THE UGLY

David Gruen, Tim Robinson and Andrew Stone

1. Introduction

Successful macroeconomic management involves a process of continual reassessment of the state of the macroeconomy. Among many things that policymakers would like to know about the current state of the economy is the extent to which the level of aggregate economic activity exceeds (or falls short of) the economy's productive capacity. This gap, between actual output and the economy's potential output, is the output gap.

For the output gap to be a concept of much value to policymakers, however, it is not enough for output-gap estimates derived with the benefit of hindsight to provide useful information about the state of excess (or insufficient) demand in the economy *in the past*. It is instead important for estimates of the current output gap formed on the basis of current information – so-called real-time estimates – to provide a reasonable guide to the current state of excess demand in the economy. In this paper, we therefore ask the question: How well can we estimate the current output gap using currently available information?

The answer to this question is important because it determines, to a considerable extent, the appropriate strategy for the conduct of monetary policy. If, as a general rule, real-time estimates provide a bad guide to the "true" current state of excess demand in the economy, then it makes no sense for policymakers to place much weight on them in their monetary policy deliberations. In that case (and presuming that there are no alternative indicators that can give a reasonable guide to the current state of excess demand in the economy), policymakers might be best advised to ignore these estimates and aim instead to stabilise the nominal growth rate of the economy, as has been argued for example by McCallum (1995, 2001).

However, if real-time output-gap estimates provide quite a good guide to the true current state of excess demand in the economy, then an alternative monetary-policy strategy seems superior. That alternative involves responding to the estimated output gap; ensuring that, other things given, monetary policy is expansionary when the estimated gap is negative and contractionary when it is positive. The Taylor rule is the most famous monetary-policy rule incorporating this logic, and many of the papers in the volume edited by Taylor (1999) suggest that monetary-policy rules of this type dominate most simple alternatives, in terms of their capacity to stabilise output and inflation in the economy. But these rules are clearly of use only if estimates of the output gap available in real time are of reasonable quality.

Estimating output gaps is not a straightforward exercise, either in real-time or with the benefit of hindsight, simply because the level of potential output, on which they are based, is unobservable. Problems associated with estimating potential output arise from several sources, among which are: uncertainty about the true structure of the economy and hence about the relationship between potential output and observed economic data on actual output, inflation, etc.; revisions to the data, particularly to actual output; and end-point problems common to most procedures used to estimate potential output.

In a companion paper, Stone and Wardrop (2002) provided an historical examination of the extent of real-time problems with the measurement of actual output. This examination suggested that the scale and persistence of actual output mismeasurement could make accurate estimation of the output gap difficult in real-time, but did not specifically address the task of estimating potential output and, hence, the output gap.

In this paper we take up the challenge of obtaining real-time potential output estimates for 121 vintages of actual Australian GDP data, to assess explicitly the extent to which real-time problems hamper the use of output gap estimates by policymakers. The two novel features of this study are the method used to obtain real-time potential output estimates, and the fact that it is the first attempt to assess the scale of the real-time problem in estimating output gaps for Australian data. In this latter regard, it supplements earlier studies for the US undertaken by Orphanides and others, and for the UK by Nelson and Nikolov.

Australia experienced a significant productivity slowdown in the early 1970s, at a similar time to other industrial countries, and a significant productivity acceleration in the 1990s, which pre-dates the US acceleration. In light of the well-known difficulties associated with estimating output gaps in real time in the presence of changes in the trend rate of growth of potential output, these productivity developments make the Australian case of interest to researchers beyond the Antipodes.

The approach we adopt involves estimating a Phillips curve for each vintage of output data, and deriving a smooth path for potential output that generates a best fit for these Phillips curves. Variants of such a Phillips curve-based approach, but using the Kalman filter to derive results, have been recently investigated for US data by Orphanides and van Norden (2001), who report that such an approach *does not* reduce real-time problems, relative even to an unsophisticated method such as using an ordinary Hodrick-Prescott filter.

There are, however, some important differences between our approach and that of Orphanides and van Norden. Orphanides and van Norden estimate Phillips curves which assume a simple relationship between inflation and the output gap. By contrast, we use specifications which allow richer dynamics and a role for influences on inflation other than the output gap. The Phillips curves we estimate include a role for the output gap and also (possible) roles for changes in the gap, estimated inflation expectations from the bond market, oil price inflation and import price inflation. We also allow for possible changes in the Phillips curve specification for different vintages of data as identifiable shocks, such as the oil shock of the early 1970s, hit the economy. In choosing our Phillips curve specifications, however, we are careful to rely solely on information available in real time. By using a richer specification for the Phillips curve, and one which can potentially change with changes in the data vintage, we hope to be able to improve the accuracy of real-time estimates of the output gap.

Overall, our examination of the Australian data, while confirming some of the broad conclusions of Orphanides and others about the difficulties of using output gap estimates to guide monetary policy in real-time, is more encouraging than these previous studies. Our results confirm that, although significant revisions to actual output estimates occur from time to time, these revisions are not the principal source of real-time problems in the estimation of the output gap. Rather, these problems arise primarily from the end-point problem of “not knowing the future”. In general, however, our results are quite promising, and suggest that useful information can be extracted from output-gap estimates in guiding judgements about policy in real time.

2. The Construction of Real-Time Potential Output Estimates

There are two alternative approaches to generating estimates of potential output, and the output gap, in real time. The first approach involves examining the historical record to see whether any explicit potential output estimates were recorded at the time, or barring that, whether such estimates can be derived from the public pronouncements of policymakers at the time. As described in more detail below, the former of these approaches was used by Orphanides (2000) for the US, while the latter was used by Nelson and Nikolov (2001) for the UK.

The alternative approach to deriving real-time potential output estimates is to use an econometric method, and this approach is used by Orphanides and van Norden (2001) for the US, and in this paper for Australia.

There is no guarantee, of course, that these two approaches will generate similar real-time estimates of potential output. There is an obvious respect in which they might differ. It seems likely that policymakers, when confronted with a change in the macroeconomy they had never seen before, may well have taken some time to fully comprehend its implications, in particular for the *framework* within which they were forming their output-gap estimates. One example of such a change is the slowdown in the rate of potential output growth in the 1970s. The possibility of such a development, and the need to allow for it in estimating potential output, might not have been apparent to policymakers at the time, whereas an analyst in 2002, cognisant of this possibility, can design an econometric approach that is robust to it.[1]

2.1 Orphanides's Approach for the US

Orphanides generates historical estimates of the US real-time output gap from two sources. For the 1960s and 1970s, his estimates are based on those generated by the Council of Economic Advisors (CEA), while for the 1980s and 1990s, they are based on estimates available directly from Federal Reserve documents. Orphanides then uses this composite real-time series in his analysis of US monetary policy over history, and in particular, of how Taylor rules would have performed had they been implemented using data available in real time.

1 Not surprisingly, the econometric approach we employ is robust to such changes.

This approach, in turn, has been criticised by Taylor (2000) on the grounds that the CEA estimates were not accepted by serious economic analysts at the time – especially in those periods when they implied very large output gaps that did not sit comfortably with other indicators of the state of the economy. Such periods coincide roughly with those over which Orphanides's analysis concludes that monetary policy based on a Taylor rule would have performed poorly in real time. Nevertheless, Orphanides' estimates represent a concrete and reasonable starting point for an historically based real-time US potential output series.

2.2 Nelson and Nikolov's Approach for the UK

The task of assembling real-time potential output estimates from historical sources turns out to be more complicated for the UK, but still tractable. While no analogue of the CEA series is available, both the output gap and the growth rate of potential output were concepts about which policymakers at Her Majesty's Treasury and the Bank of England were prepared to hazard occasional public guesses at least as far back as the mid-1960s. By meticulously sifting back through nearly 40 years of budget papers and speeches by the Chancellor of the Exchequer and the Governor of the Bank of England, Nelson and Nikolov (2001) were thus able to reconstruct an approximate real-time series for potential output. This series uses intermittent estimates available from these sources for the output gap, and interpolates between them based on occasional estimates for the growth rate of potential output also found in these documents.[2] Although also obviously imperfect and open to dispute, the series so generated once again provides at least a plausible initial guess at a real-time potential output series for the UK.

2.3 The Unavailability of Historical Information for Australia

Neither of these historical approaches to obtaining a real-time potential output series can be implemented for Australia. No systematic estimates of Australian potential output, akin to those prepared for the US by the CEA, are available. Equally, perusal of Reserve Bank annual reports or Commonwealth budget

2 Note that this process leads to a real-time potential output series subject to occasional substantial breaks. For example, publication after several years of a new and different estimate of the growth rate of potential output over recent history, say in response to apparent persistent under- or over-performance by the economy, leads to a sudden break in the real-time estimate of potential output at that time.

papers from the 1970s onwards shows that, while reference is often made to the economy's "supply capacity", and to whether or not it is operating "at full capacity", or at "full stretch", concrete estimates are not provided of either the output gap or the economy's potential growth rate. The same is true of other possible sources of such historical information. Hence, it is also not possible to replicate Nelson and Nikolov's approach for Australia.

2.4 Generating Real-Time Potential Output Estimates for Australia

To construct real-time potential output and output gap estimates for Australia, we therefore use an econometric approach. This has some advantages over an historical approach. There will always be room for debate about whether policymakers' historical estimates of the output gap could have been improved upon at the time. By contrast, an econometric approach, designed to be as robust as possible to a range of specification problems, may enable us to come to a more informed view about the *inherent* seriousness of the real-time problems associated with estimating output gaps. Such an approach should therefore allow us to better assess whether we are likely to be plagued by these problems in the future.

To implement this approach, we specify a method of generating vintages of potential output which requires only data for economic indicators, such as actual output, for which we have real-time information over history. Then, by applying this procedure successively to the real-time data sets in each quarter, we create corresponding implied real-time potential output estimates for Australia.

An outline of the procedure follows, with technical details relegated to Appendix A. Consider an expectations-augmented Phillips curve of the generic form

$$\pi_t = \pi_t^e + \gamma(y_t - y_t^*) + \theta Z_t + \varepsilon_t \tag{1}$$

where π_t denotes quarterly (core consumer price) inflation, π_t^e denotes inflation expectations, y_t and y_t^* denote actual and potential output (in logs), Z_t represents a vector of other variables (which may include changes in the output gap), and ε_t denotes an error term.

The Z_t variables are constructed so that they are zero in the long-run steady state. As a consequence, the Phillips curve defined by equation (1) is vertical in the long run, with output at potential when inflation is equal to expected inflation.

For each vintage of data, we seek the *smooth* path for potential output that gives the best fit to this Phillips curve equation. Formulated mathematically, we find the values for the parameter γ, the parameter vector θ, and the potential output series $\{y_t^*\}$ which minimise the loss function

$$\mathcal{L} = \sum_{t=1}^{n} \varepsilon_t^2 + \lambda_{PC} \sum_{t=2}^{n-1} \left((y_{t+1}^* - y_t^*) - (y_t^* - y_{t-1}^*) \right)^2 \tag{2}$$

where, as in the usual Hodrick-Prescott filter, λ_{PC} is a smoothing parameter to be chosen.[3]

2.5 Phillips Curve Specifications

To choose appropriate specifications for our Phillips curves, we adopt a general-to-specific approach. The general specification includes possible roles for lags of inflation, lagged inflation expectations from the bond market, the output gap, current and lagged changes in the gap ("speed-limit" terms), and current and lagged oil price and import price inflation.

In terms of the generic form of the Phillips curve, equation (1) above, the general specification assumes that inflation expectations are a linear combination of lagged inflation and inflation expectations from the bond market,

$$\pi_t^e = \sum_{i=1}^{k_\alpha} \alpha_i \pi_{t-i} + \sum_{i=1}^{k_\beta} \beta_i \text{bond}_{t-i} \tag{3}$$

and that the vector of other variables is of the form

$$\theta Z_t = \sum_{i=0}^{4} \delta_i \Delta(y_{t-i} - y_{t-i}^*) + \sum_{i=0}^{k_\eta} \eta_i \text{oil}_{t-i} + \sum_{i=0}^{k_\xi} \xi_i \text{import}_{t-i} \tag{4}$$

where $\text{bond}_t = (\pi_t^b - \Delta_4 p_{t-1})/4$ is the excess of bond market inflation expectations over lagged year-ended inflation, expressed in per-quarter terms; $\text{oil}_t = \pi_t^{oil} - \pi_{t-1}$

3 We choose $\lambda_{PC} = 80$, which leads to derived potential output series that are much smoother than those derived from an H-P filter of the output data with the standard smoothing parameter of $\lambda_{HP} = 1600$. In the notation of Laxton and Tetlow (1992), our minimisation procedure corresponds to a multivariate Hodrick-Prescott filter of type $(0, 1, 0, \lambda_{PC})$. We examine the sensitivity of our results to a change in the value of λ_{PC} later in the paper.

is the excess of quarterly oil price inflation over lagged inflation; and $\text{import}_t = \pi_t^{import} - \pi_{t-1}$ is the excess of import price inflation over lagged inflation. For all the price series (core consumer prices, oil prices and import prices), we use the first difference of the log price level to approximate the quarterly inflation rate. We impose the constraint $\sum_{i=1}^{k_\alpha} \alpha_i = 1$, to ensure that when expected inflation, π_t^e, is expressed as a linear combination of lags of inflation and bond market inflation expectations, the coefficient weights sum to unity. A full description of the data is provided in Appendix D.[4]

Australian quarterly GDP data are now available from 1959:3 to the present. However, 1971:4 is the first quarter for which we have original-vintage GDP data back to 1959:3, and so our real GDP vintages run from 1971:4 to 2001:4; 121 vintages in all.

For a given vintage of GDP data, we start with the general specification defined by equations (1), (3) and (4) above. The values of k_α, k_β, k_η, and k_ξ, which define the lag lengths of the variables, are kept small (most commonly, $k_\alpha = 4$, and $k_\beta = k_\eta = k_\xi = 2$), although some searching is also undertaken of variables at longer lags (lag lengths of up to eight). Variables with coefficient *t*-statistics less than about 1.5 are sequentially eliminated, which leads eventually to a parsimonious specific specification. As it turns out, the coefficients on most variables in most specific specifications have *t*-statistics greater than 2, and the coefficient on the output gap usually has a *t*-statistic in excess of 5.[5]

4 Note that imposing the constraint $\sum \alpha_i = 1$ does not imply wholly backward-looking inflation expectations. Indeed, the general specification allows for expectations to be wholly forward looking, as could be the case if $\sum \beta_i = 1$. Therefore, our specification should not be subject to Sargent's critique of the accelerationist Phillips curve (Sargent (1971)). Note also that there are some minor aspects of the data that are not available in real time. They are discussed in Appendix D. The most substantive of them involves the construction of bond market inflation expectations before 1993, which requires the value of a parameter, and we use a value from Tanzi and Fanizza (1995). We establish in Appendix D, however, that our results are fairly insensitive to the value of this parameter.

5 Note that these *t*-statistics cannot be translated straightforwardly into levels of statistical significance because the potential output series used in the equations is generated as part of the estimation procedure, and hence the usual standard errors do not apply. It is to reduce the severity of this problem that we choose a value of the smoothness parameter, λ_{PC}, in equation (2) that leads to such smooth derived potential output series. Monte Carlo simulations on pseudo-data generated by a bootstrapping procedure suggest that the coefficient estimates from our Phillips curves are not subject to significant biases (see Appendix B).

Rather than conduct a new specification search for each new data vintage, we revisit the specific specification of the Phillips curve whenever significant deterioration is observed in the performance either of the overall equation or of its components; or, in any event, roughly every ten to twelve years. Particular emphasis is placed on the stability of the coefficient on the output gap. In analysing coefficient stability and equation performance, we use both regressions where the start date is held fixed (at 1961:2) and fifteen-year rolling regressions.

Conducting specification searches only intermittently seems to make little difference to our estimates of potential output and the output gap, except on rare occasions when inflation is subject to a major shock of a type which has not been seen before. The chief instance of this is the first oil price shock in the early 1970s, and when such events occur, we re-specify the Phillips curve more frequently.

In all, for the 121 data vintages from 1971:4 to 2001:4, five broad Phillips curve specifications are used, as outlined in Table 1. Within the five periods delineated in Table 1, minor additional changes are also sometimes made to the Phillips curve specification for particular vintages. Full details of these re-specifications, together with a brief discussion of the reasoning behind the changes, can be found in Appendix C.

3. Results

In this section we report on the performance of our Phillips curve method of estimating output gaps, both using the latest available data and in real time. We also compare this performance with that of alternative approaches to estimating output gaps, both for Australia and the US (as reported in Orphanides and van Norden (2001)).

Figure 1 shows some of the data used in the estimation of our Phillips curves. Included in the figure are the core consumer price inflation series we use, inflation expectations from the bond market, and oil and import price inflation. Also shown is the year-ended growth rate of GDP from the final vintage of data, the 2001:4 vintage, along with a ten-year moving average which shows how average GDP growth has varied over the past four decades.

Table 1: Specification of Equations	
Date of Vintage	Broad Equation Specification
1971:4 to 1973:3	$\pi_t = 0.5(\pi_{t-2}+\pi_{t-3})+\beta_2\text{bond}_{t-2}+\beta_3\text{bond}_{t-3}+\beta_4\text{bond}_{t-4}+\gamma(y_t-y_t^*)+\delta_4\Delta(y_{t-4}-y_{t-4}^*)+\xi_3\text{import}_{t-3}+\xi_4(\text{import}_{t-4}+\text{import}_{t-5}+\text{import}_{t-6})$
1973:4 to 1974:2	$\pi_t = 0.5(\pi_{t-2}+\pi_{t-3})+\beta_2\text{bond}_{t-2}+\beta_3\text{bond}_{t-3}+\gamma(y_t-y_t^*)+\eta_3\text{oil}_{t-3}$
1974:3 to 1986:2	$\pi_t = 0.25(\pi_{t-2}+\pi_{t-3}+\pi_{t-4}+\pi_{t-5})+\beta_1\text{bond}_{t-1}+\gamma(y_t-y_t^*)+\eta_2\text{oil}_{t-2}$
1986:3 to 1998:2	$\pi_t = 0.25(\pi_{t-2}+\pi_{t-3}+\pi_{t-4}+\pi_{t-5})+\zeta_2(\pi_{t-2}-\pi_{t-6})+\zeta_3(\pi_{t-3}-\pi_{t-7})+\beta_1\text{bond}_{t-1}+\beta_2\text{bond}_{t-2}+\gamma(y_t-y_t^*)+\eta_2\text{oil}_{t-2}+\eta_3\text{oil}_{t-3}+\eta_7\text{oil}_{t-7}$
1998:3 to 2001:4	$\pi_t = 0.25(\pi_{t-1}+\pi_{t-2}+\pi_{t-3}+\pi_{t-4})+\zeta_2(\pi_{t-2}-\pi_{t-6})+\beta_1\text{bond}_{t-1}+\gamma(y_t-y_t^*)+\eta_2\text{oil}_{t-2}+\eta_7\text{oil}_{t-7}+\xi_0\text{import}_t+\xi_1\text{import}_{t-1}$

Note: Start of sample for all regressions is 1961:2.

3.1 Preferred Phillips Curve Method of Estimating the Output Gap

Table 2 reports estimation results for the final Phillips-curve specification estimated on the final vintage of GDP data. The estimated equation explains a sizable fraction of the variation in quarterly inflation over the past four decades. The *t*-statistic on the output gap is about 6. It is also noteworthy that inflation expectations from the bond market play an important role in the equation – they contribute about forty per cent to inflation expectations (since the coefficient on bond_{t-1} is about 0.4), with the remaining sixty per cent being contributed by lags of inflation – and that there is no apparent serial correlation in the equation residuals.

Figure 2 shows estimates of the growth rate of potential output and the output gap over the past four decades based on the estimated results using the final data vintage. We refer to these output gap estimates as the “final” output gap estimates. Figure 3 shows the (log) levels of actual and potential output, again based on the

Figure 1: Data

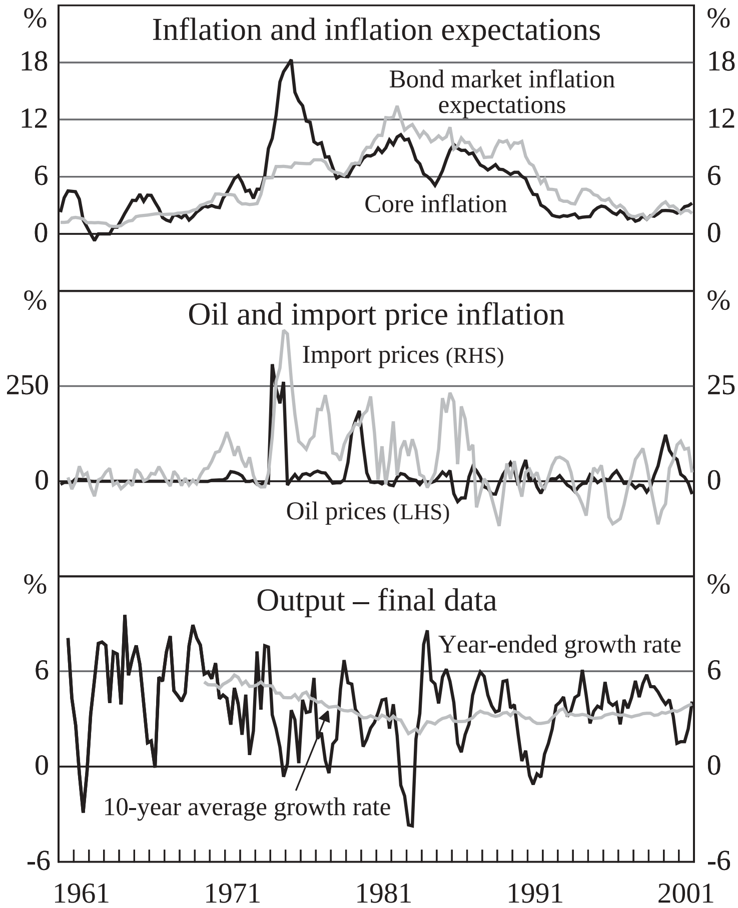

Note: The inflation rates (core consumer price, oil, and import price) and output growth are year-ended percentage changes. The estimation results, however, use log-differences to approximate percentage changes.

estimated results using the final data vintage. The smoothness of the potential output series is also clear from this figure.

The results suggest that the potential capacity of the economy grew at an annual rate of nearly five per cent through much of the 1960s before slowing over the next couple of decades to an annual rate more like three per cent in the early 1990s. In the latest few years, however, this growth rate appears to have accelerated to a little over 3.5 per cent per annum.

The estimated output gap implies that the economy was operating above its potential capacity in the last few years of the 1960s, and for much of the 1970s, and below its potential through much of the 1980s and 1990s. Not surprisingly given

Table 2: Estimation Results for the Final-Vintage Phillips Curve

$$\pi_t = 0.25(\pi_{t-1}+\pi_{t-2}+\pi_{t-3}+\pi_{t-4})+\zeta_2(\pi_{t-2}-\pi_{t-6})+\beta_1\text{bond}_{t-1}+\gamma(y_t-y_t^*)+\eta_2\text{oil}_{t-2}+\eta_7\text{oil}_{t-7}+\xi_0\text{import}_t+\xi_1\text{import}_{t-1}$$

Coefficient	Value	t-statistic
ζ_2	0.112	2.408
β_1	0.418	6.939
γ	0.064	6.081
η_2	0.007	3.102
η_7	0.007	3.324
ξ_0	0.023	2.202
ξ_1	0.020	1.845
Summary Statistics	**Value**	
R^2	0.823	
Adjusted R^2	0.817	
Standard error of the regression	0.004	
Breusch-Godfrey LM test for autocorrelation (p-value):		
First order	0.976	
First to fourth order	0.419	

Note: The sample is 1961:2 - 2001:4 ($n = 163$).

how they were generated, these patterns of capacity utilisation appear broadly consistent with the behaviour of inflation over the four decades.[6]

To examine the scale of the real-time problem, we construct a new series of output-gap estimates, the real-time estimates. To do so, we extract the *last* estimate from the output-gap series derived from each data vintage. The real-time output-gap estimates are then these last estimates, strung together into a new composite series. This real-time output gap is shown along with the final output gap estimates in Figure 4.

6 Two factors help to explain why the output gap is negative for almost the whole of the 1980s and 1990s. First, the process of disinflation requires a negative output gap. And secondly, inflation expectations from the bond market remained above actual inflation for most of the 1980s and 1990s, as shown in Figure 1. While that remains the case, the output gap must be negative to keep inflation steady (since potential output is defined as the level of output consistent with actual inflation being equal to inflation expectations, and inflation expectations are a combination of past inflation and inflation expectations from the bond market).

Figure 2: Potential Output and Gap Estimates From Final Phillips Curve

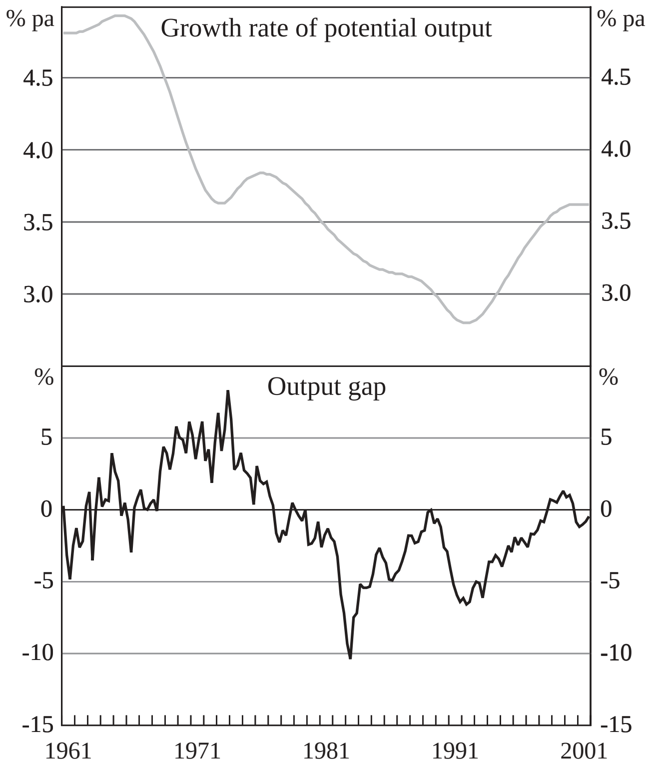

The results of this exercise are quite encouraging. There is quite a close correspondence between the real-time and final estimates of the output gap for most of the past three decades, although two periods of divergence stand out – the first few years of the 1970s and the few years surrounding the early 1990s recession. Before discussing these two periods, however, we present the results of a further exercise which proves illuminating.

Recall that the real-time estimates of the output gap differ from the final estimates both because they use different vintages of GDP data (with the former using a different data vintage for each output gap estimate, while the latter uses the final vintage throughout), and also because, in contrast to the final estimates, no information about the future is used in the construction of the real-time estimates. We can assess the relative importance of these two influences by constructing an alternative set of real-time estimates, the "quasi-real" estimates. The quasi-real estimate of the gap in period *t* is constructed using the final-vintage data up to period *t*, rather than the period-*t* data vintage. Any differences between the real-time and quasi-real output gap estimates are therefore due entirely to data

Figure 3: Actual and Potential Output From Final Phillips Curve

Quarterly output at 1999/2000 prices, Log scale

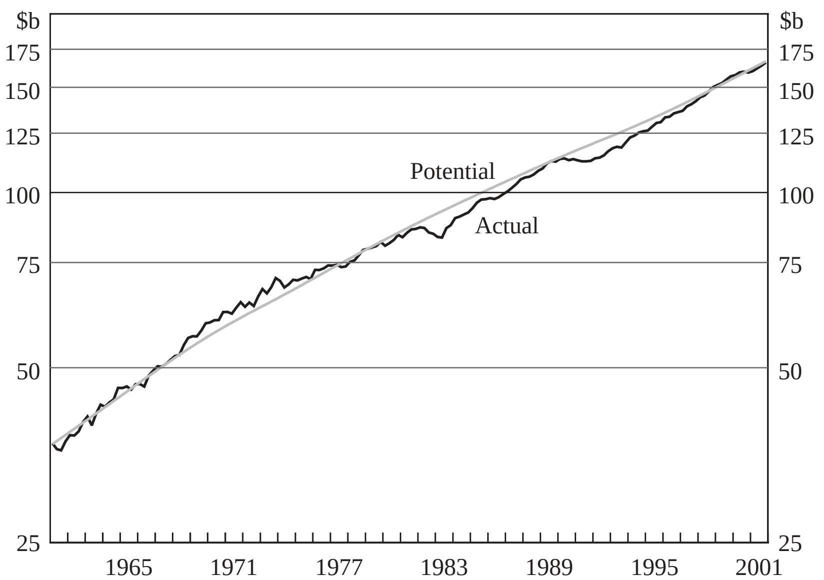

Figure 4: Real-time and Final Output Gaps

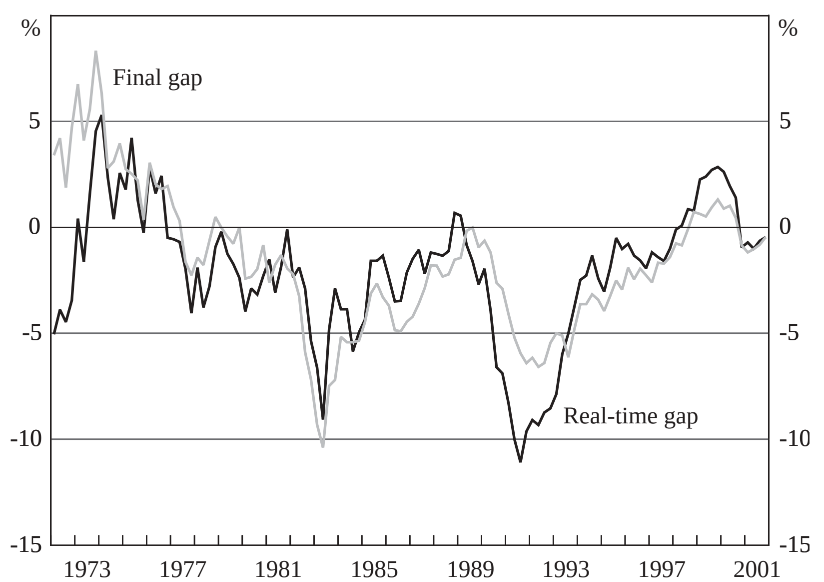

revisions, since estimates in the two series at any point in time are derived using data over identical time periods.[7]

As the results shown in Figure 5 demonstrate, the real-time and quasi-real output gap estimates are quite similar for most of the past three decades, although there are some noticeable differences in the mid 1970s.[8] It seems clear, therefore, that the predominant reason why real-time output gap estimates differ from final estimates is that (with the obvious exception of the last quarter in the sample), the final estimates use information about the future that is unavailable in real time.

Figure 5: Effect of Data Revisions

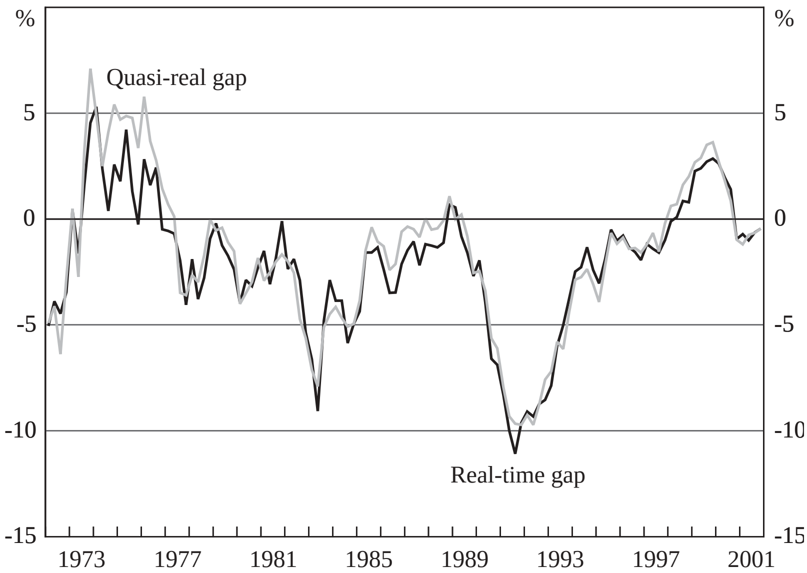

We can now return to the discussion of the two periods of divergence between the real-time and final output gaps highlighted in Figure 4 — the first few years of the 1970s and the few years surrounding the early 1990s recession — confident that the divergence is not due primarily to data revisions. Two further figures aid the discussion. Figure 6 shows the final output gap estimates together with the quasi-real estimates, to abstract from data-revision issues. Figure 7 shows

7 The term "quasi-real" is from Orphanides and van Norden (2001). For each sample period, we construct the quasi-real output gap estimates using the same Phillips curve specification as for the real-time estimates, though the coefficient estimates are different in general.

8 The mid 1970s were still very early years in the production of estimates of real GDP by the Australian Bureau of Statistics. The first real GDP estimates, which were income-based, were for the 1971:3 vintage, while expenditure-based estimates only began with the 1974:4 vintage. It is perhaps not surprising, therefore, that the greatest differences between the real-time and quasi-real output-gap estimates are found in these early, experimental years.

the unexplained residuals from the final Phillips curve, estimated over the whole sample, 1961:2 to 2001:4, using the final data vintage. The equation clearly does a poor job explaining inflation around these two periods of divergence, with the sample's three largest residuals in absolute value occurring in 1973:4, 1974:3 and 1991:1.

Figure 6: Quasi-real and Final Output Gaps

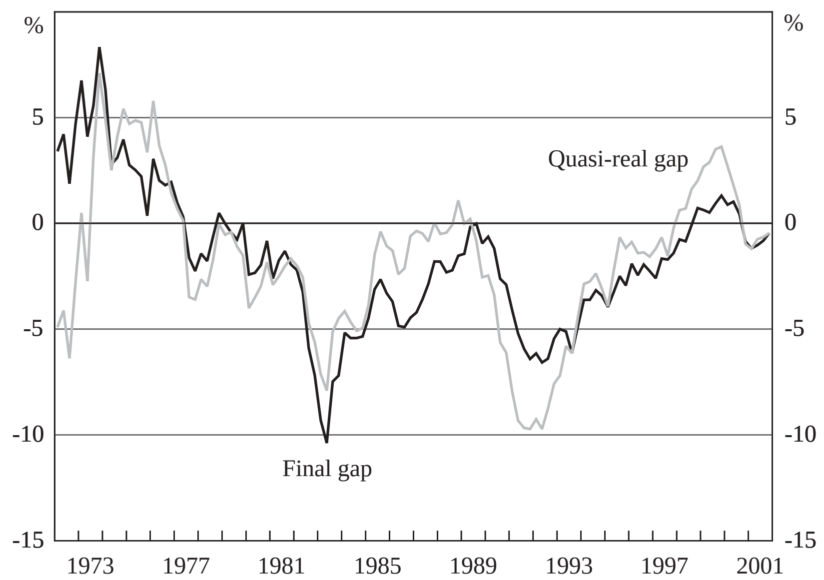

Figure 7: Final Phillips Curve Residuals

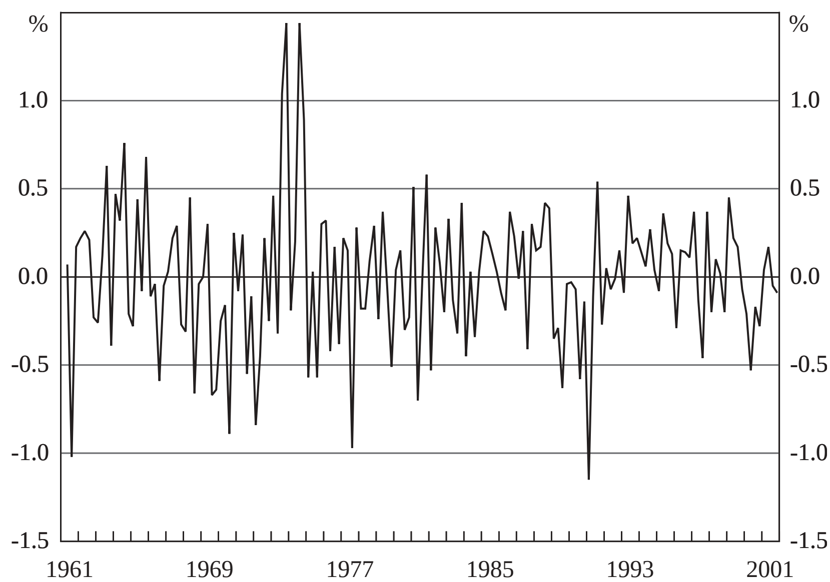

Some part of the high inflation outcomes of 1973/74 is a consequence of wage developments in the economy that should be regarded, at least to some extent, as unrelated to the state of the economy and the size of the output gap.[9] Unsurprisingly, Phillips curves estimated using either real-time or final-vintage data up to 1973 fail to anticipate this partly exogenous event. As a consequence, output gaps derived from these Phillips curves perform particularly poorly when compared with final output gaps, as Figures 4 and 6 make clear.

Given the partly exogenous nature of this event, however, we are inclined to view it as not particularly informative about general real-time problems with output-gap estimates. As a consequence, we focus much of our analysis on the twenty-eight years *after* this event, that is, from 1974:1 to 2001:4.[10] Nevertheless, for completeness, we also report results for our full set of data vintages, from 1971:4 to 2001:4.

The second period of significant divergence between the real-time (or quasi-real) and final output gap estimates occurs around the time of the early 1990s recession. Inflation fell more rapidly at that time than the average behaviour of inflation over the past four decades would lead one to expect, as is clear from the sequence of negative residuals around that time from the final estimated Phillips curve – including the largest negative residual in the sample, in 1991:1 (Figure 7). Phillips curves estimated using information up to that time interpret this unexpectedly rapid disinflation as a sign that the output gap is large and negative at that time – an interpretation that is subsequently revised as later information arrives. But this, of course, is simply a classic illustration of the real-time problem in action. The crucial issues are how large and how frequent are such real-time problems.

---

9 Australia at the time had a centralised wage-setting system in which an official body with legislated powers, the *Conciliation and Arbitration Commission*, set wages for much of the workforce. A left-of-centre government was elected in December 1972, coming to power for the first time for 23 years. Within a few months of the election, the *Conciliation and Arbitration Commission* awarded a 17.5 per cent increase in minimum wages at a time when consumer price inflation, although rising, was running at an annual rate of less than 6 per cent. Although this wage decision was undoubtedly influenced by the state of the economy at the time, it seems clear that political-economy influences also played a role.

10 Note that focusing on the period from 1974:1 onwards does not exclude the effects of the first OPEC oil shock. After a period of quiescence from mid-1971 to the end of 1973, the Australian-dollar price of oil nearly quadrupled in 1974:1.

There is also a third period of divergence between the real-time (or quasi-real) and final output gap estimates which is less pronounced, but nevertheless instructive. It occurs in the second half of the 1990s, at a time when (we would now assess that) potential output growth was accelerating from an annual rate of around 3 per cent at the beginning of the decade to above 3.5 per cent near its end (see Figure 2). We would not expect estimates of the level of potential output constructed during the transition to this faster rate of potential growth to be able fully to take it into account. As a consequence, we would expect real-time, or quasi-real, estimates of the output gap to be systematically above (more positive than) final estimates during this transition. While this pattern is clearly discernable in Figures 4 and 6, these real-time errors are, for the most part, quite moderate in size. This experience suggests that our approach to estimating output gaps in real time is reasonably robust to (at least moderate) changes in the rate of growth of potential output.

There are two further interesting ways to examine the relationship between the real-time and final output gap estimates. One is to use a scatter-plot, as shown in Figure 8. The closeness of the relationship between the two estimates over the twenty-eight years, 1974:1 to 2001:4, is reflected in the clustering of points relatively near the 45 degree line.

Figure 8: Scatter-plot of Real-time Gap versus Final Gap

Phillips curve Based, 1974:1 - 2001:4

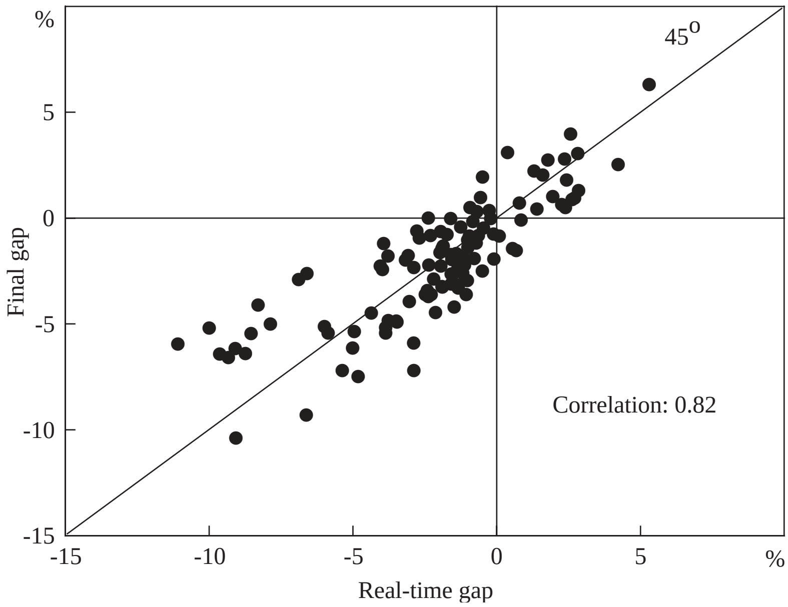

Figure 9: Differences Between Final and Real-time Gaps

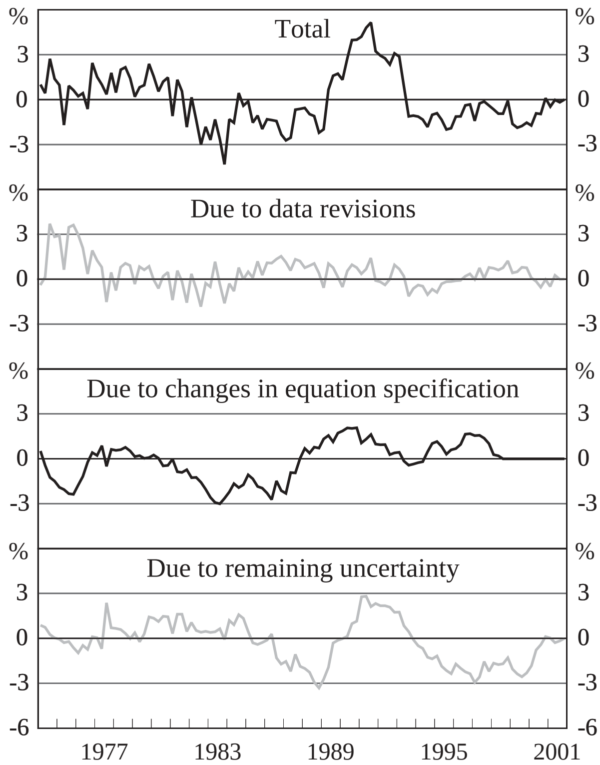

Finally, the top panel of Figure 9 shows the difference between the final and real-time output gap estimates through time, while the subsequent panels show a decomposition of this difference, constructed so that the numbers in these lower three panels sum to the number in the top panel at each time *t*. The second panel shows the difference attributable to data-revisions – that is, the difference between the real-time and quasi-real output gap estimates. The third panel shows the difference between the quasi-real estimate at time *t* and the output gap estimate also derived from final-vintage data up to time *t*, but using the final Phillips curve specification rather than the vintage-*t* specification used to derive the quasi-real estimates. The fourth panel of the figure shows the remaining difference.[11]

11 Note that the decomposition shown in Figure 9 is not unique. To move from the real-time to the final output-gap estimates involves three changes: to the data; to the Phillips curve specification; and to the period of estimation. There is, however, some flexibility in the order in which these changes are implemented. For example, the second panel could show the difference between the real-time estimates and estimates derived using the real-time data but with the final Phillips curve specification.

As previously discussed, data revisions generate only small changes in the output gap estimates (panel two). Note that both panels three and four are relevant to the problem of "not knowing the future". The "final" specification of the Phillips curve can only be discovered after observing the economy's evolution over the full sample (relevant to panel three). Likewise, "final" coefficient estimates only become clear with the full sample results (relevant to panel four).

3.2 Alternative Methods of Estimating the Output Gap

We now turn to a comparison between our preferred method of estimating real-time output gaps and alternative approaches. The alternatives we consider are two univariate methods (which assume that potential output is either a linear trend of actual output or a Hodrick-Prescott filter of actual output) and a variant of our Phillips curve approach in which we assume that potential output grows at a constant rate, rather than allowing for changes in the growth rate as in our preferred method. (There is also a fourth alternative approach which we will discuss shortly.)

As with our preferred method, each of these alternative methods can be applied to each of the real-time data vintages, and a real-time output gap series can then be constructed by stringing together the last output-gap estimate from each data vintage.

We can then compare each of these real-time estimates with final output gap estimates (using the full sample of final-vintage data) derived in one of two ways: either by using the same method as for the real-time estimates (with results shown in Table 3), or by using the preferred Phillips-curve method (Table 4). Table 3 also shows some of the results presented by Orphanides and van Norden (2001) for the US.

For each method, each table shows (over the relevant sample) the mean difference between the final and real-time output gap series, the correlation between them, their root-mean-square difference (RMSD), and the first-order serial correlation of the difference (AR). Thus, for example, using a linear trend through (log) Australian output to estimate (log) potential output generates a real-time output gap series which is on average 6.60 percentage points (ppt) below the final output gap series derived using the same method (from Table 3). These real-time output-gap estimates can only be described as ugly. The huge differences between the real-time and final output-gap estimates using this linear-trend method occur, of course, because the growth rate of actual and potential output has been subject to significant, long-lived, changes over the four decades from the early 1960s to 2001.[12]

Applying a standard H-P filter (with $\lambda_{HP} = 1600$) to the Australian data generates real-time output gap estimates that differ by an average of only 0.19 percentage points from the final H-P filter estimates, a huge improvement over the performance of the estimates derived from linear trends. But these H-P filter estimates are nevertheless unsatisfactory in several respects. Figure 10 shows the comparison between the real-time and final output gap estimates derived using H-P filters, and Figure 11 shows a scatter plot of the two series (corresponding to Figures 4 and 8 for our Phillips curve method).

Figure 10: Real-time and Final Output Gaps

HP Filter Based, $\lambda_{HP} = 1600$

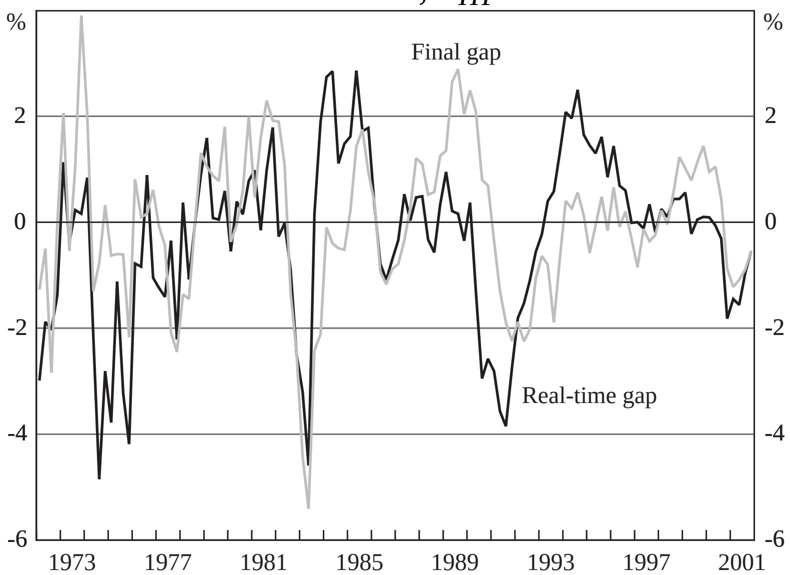

12 Not surprisingly, this method generates real-time output gap estimates which also differ very significantly from the final estimates derived using the preferred method (see Table 4). It is interesting to note that Taylor (1993), in the paper that introduced the Taylor rule, generated estimates of the output gap by removing a linear trend from actual output. The approach arguably produced quite good estimates in that case because it was applied over a short sample of only nine years from 1984:1 to 1992:3, during which the trend rate of potential output growth was probably fairly stable.

Table 3: Comparing the Final and Real-time Output Gap Estimates

Method	Mean Diff (ppt)	Corr	RMSD (ppt)	AR
Australian Results				
Univariate				
Linear trend	6.60	0.27	7.72	0.96
Hodrick-Prescott filter ($\lambda_{HP} = 1600$)	0.19	0.50	1.53	0.79
Phillips curve-based				
Potential output growth: constant[a]	8.44	0.76	9.73	0.83
Potential output growth: variable				
Preferred Phillips curve method[b]	-0.01	0.82	1.82	0.80
"Simple" Phillips curve[c]	2.86	0.65	4.24	0.74
US Results				
Univariate				
Linear trend	4.78	0.89	5.12	0.91
Hodrick-Prescott filter ($\lambda_{HP} = 1600$)	0.30	0.49	1.83	0.93
Phillips curve-based				
Potential output growth:				
Constant + "noise" (Kuttner)[d]	3.57	0.88	3.97	0.92
Potential output growth:				
Variable (Gerlach-Smets)[d]	1.64	0.75	2.17	0.80

Notes: Results are calculated over the relevant sample of data vintages, which is 1974:1 - 2001:4 for the Australian results and 1966:1-1997:4 for the US results. The US results are from Orphanides and van Norden (2001). "Mean Diff" is the mean difference between the final and real-time output gap series; "Corr" is the correlation coefficient between the two series; RMSD is the root-mean-square difference between them; and AR is the the first-order serial correlation of the difference.

[a]These results are derived assuming that $y_t^* = a + bt$, where a and b are freely estimated as part of the Phillips curve estimation for each data vintage. The equation specifications used for these Phillips curves are optimised in a similar way to our preferred approach. Appendix C provides the equation specifications used for each data vintage, and estimation results for the final data vintage.

[b] Over the longer sample of data vintages, 1971:4 to 2001:4, Mean diff is 0.5 ppt, Corr is 0.71, RMSD is 2.6 ppt, and AR is 0.89.

[c]This Phillips curve specification is defined by equation (5) in the text. These results exclude the quarters of 1974:3, 1974:4, 1975:2 and 1975:4, in which the optimisation procedure fails to converge, or generates a negative or zero coefficient on the output gap, γ.

[d]The Kuttner and Gerlach-Smets approaches use Phillips curves based solely on information about inflation and output, and use the Kalman filter to derive results.

Table 4: Comparing the Final Preferred Gap with Alternative Real-time Gaps				
Method	Mean Diff (ppt)	Corr	RMSD (ppt)	AR
Univariate				
Linear trend	4.20	0.73	4.75	0.86
Hodrick-Prescott filter ($\lambda_{HP} = 1600$)	-1.86	-0.04	3.83	0.95
Phillips curve-based				
Potential output growth: constant	10.93	0.76	11.90	0.85
Potential output growth: variable				
Preferred Phillips curve method	-0.01	0.82	1.82	0.80
"Simple" Phillips curve[a]	1.04	0.55	5.66	0.19

Notes: Results are calculated over the 1974:1 - 2001:4 sample of Australian data vintages. "Mean" diff is the mean difference between the final preferred output gap series and the real-time gap series generated by the method shown; "Corr" is the correlation coefficient between the two series; RMSD is the root-mean-square difference between them; and AR is the first-order serial correlation of the difference.
[a]See note *c* in Table 3.

Figure 11: Scatter-plot of Real-time Gap versus Final Gap

HP Filter Based, 1974:1 - 2001:4

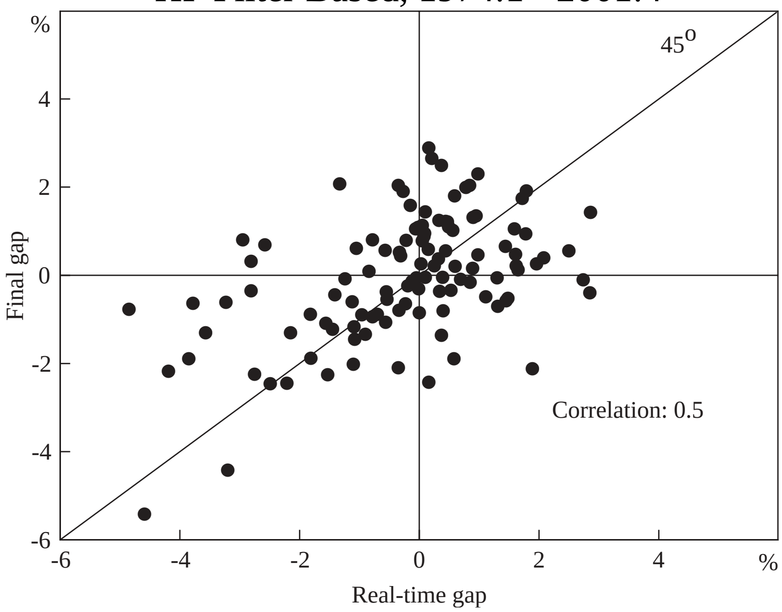

The relationship between the real-time and final H-P-filter output-gap estimates is much less clear than between the corresponding series when they are derived using the preferred Phillips-curve method, as a comparison of the relevant figures makes clear. Real-time H-P-filter output-gap estimates also appear to bear almost no resemblance to our preferred final output gap estimates – indeed, the correlation between the two series is slightly negative (see Table 4).

Furthermore, the inflation experience over the four decades from 1960 to 2001 would be particularly hard to understand on the basis of output gaps derived from the H-P filter. The average values of the final H-P-filter output gaps in the four decades of the 1960s, 70s, 80s and 90s, are, in percentage points, –0.1, 0.1, 0.2, and –0.1, a pattern of capacity utilisation that clearly gives no hint about the longer-run inflation developments over this time.[13]

Turning to the Phillips-curve based approach applied to Australian data, very different results emerge depending on whether or not the growth rate of potential output is allowed to vary. Assuming a constant rate of potential growth generates bad results, just at it did in the univariate case (as the results in Tables 3 and 4 make clear). By contrast, our preferred Phillips curve approach, which allows for gradual changes in the rate of potential growth, generates very substantial improvements in the performance of the real-time output gap estimates.

It is also instructive to examine the performance of a "simple" Phillips curve. The specification we assume for this simple Phillips curve, for all data vintages, is

$$\pi_t = \pi_{t-1} + \gamma(y_t - y_t^*) + \varepsilon_t. \tag{5}$$

Results for this simple Phillips curve are derived in the same way as for the preferred Phillips curve.[14] Summary statistics are shown in the final row of Australian results in both Table 3 and 4, and a comparison of the real-time and final output gap estimates derived using this simple Phillips curve are shown in Figure 12.

A comparison of the results from the preferred Phillips curve with those from this simple Phillips curve suggests that, in order to generate reasonably accurate

---

13 The corresponding average final output gaps using our preferred Phillips curve method in the four decades are 0.7, 2.6, –3.6 and –2.9. These decadal averages sit much more comfortably with the observed inflation outcomes – with strongly rising inflation and inflation expectations in the 1970s, and the opposite in the 1980s and 1990s.

14 That is, for each data vintage, we find the values for γ, and the potential output series $\{y_t^*\}$ which minimise the loss function, equation (2), using the same smoothness parameter as for the preferred Phillips curve, $\lambda_{PC} = 80$.

Figure 12: Real-time and Final Output Gaps

Based on the simple Phillips curve

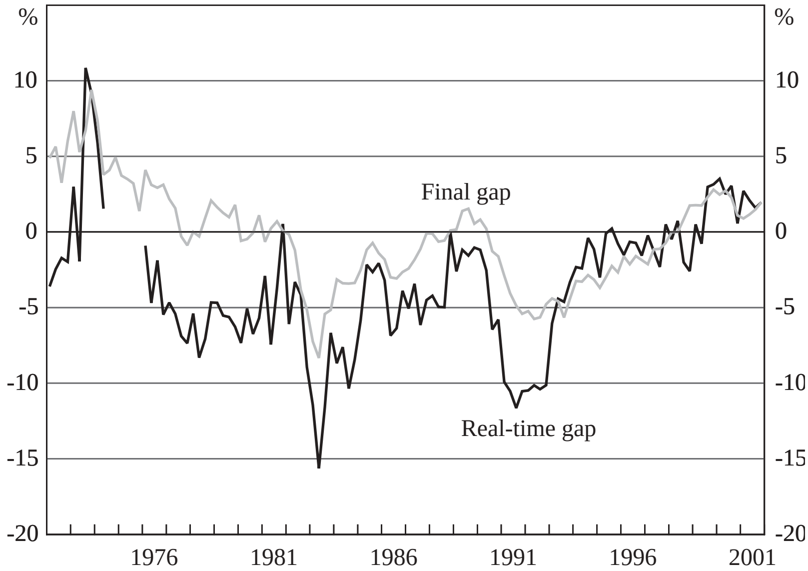

Note: The real-time output gap estimates are not shown over the period, 1974:3 to 1975:4, for reasons explained in note *c* in Table 3.

real-time output-gap estimates, it is important to use an information-rich, and fairly well-specified, Phillips curve equation. The real-time output-gap estimates from the simple Phillips curve differ very substantially from final estimates derived either using the same method or using the preferred method (as evidenced, for example, by the large root-mean-square differences reported in the two tables).

Turning to the US results, the univariate linear-trend and H-P-filter methods generate fairly poor results, just as they do when applied to the Australian data. For the Phillips-curve based approaches, the Kuttner method also seems to work poorly, just as the Phillips curve with a constant-potential-growth-rate does for the Australian data.[15]

The Gerlach-Smets approach allows for a variable rate of potential output growth. On the basis of the Australian results, one might expect this assumption to generate

[15] The Kuttner method assumes that potential output evolves as a random walk with constant drift (and hence potential growth differs from a constant by an i.i.d shock). This clearly introduces more flexibility into the assumed process for potential output than a linear trend, but apparently not enough.

a substantial improvement in performance, relative to the alternatives. There is some improvement, but nevertheless the Gerlach-Smets approach, at least when applied to US data, still seems prone to more serious real-time errors than our preferred Phillips curve approach applied to Australian data.

It may be that the simplicity of the Phillips-curve specification used by Gerlach-Smets is responsible for this relatively poor performance. The Gerlach-Smets Phillips curve is very similar to the "simple" Phillips curve defined by equation (5) above, which as we have seen, generates poor real-time output gap estimates for Australia.[16]

4. Sensitivity Analysis and Confidence Intervals

In this section, we present sensitivity analysis and confidence intervals for the results derived from our preferred Phillips curve approach.

4.1 Sensitivity to Changes in the Smoothing Parameter, λ_{PC}

In choosing a value for the smoothness parameter, λ_{PC}, in equation (2), we have been guided by a desire to allow long-lived changes in the rate of potential output growth to manifest themselves, without generating high-frequency "noise" in the resulting potential-output series. The value we choose, $\lambda_{PC} = 80$, seems to satisfy these criteria, as it generates a smooth path for potential output which nevertheless displays significant long-lived changes in its growth rate, as Figure 2 demonstrates.

It is of interest, however, to see how sensitive our results are to this choice of parameter value. Figure 13 shows a comparison of the estimated rate of growth of potential output through time from results derived from the final

[16] The Phillips curve specification used by Gerlach-Smets has quarterly inflation on the left-hand side, and a constant, lagged quarterly inflation, and the output gap as explanatory variables. This is essentially identical to our "simple" Phillips curve specification, equation (5). (There is no need for a constant in our simple Phillips curve because our optimisation procedure allows the whole potential output series, $\{y_t^*\}$, to shift up or down to give a best fit to the Phillips curve). The only substantive difference between the two specifications is the assumption of an MA(3) error process in the Gerlach-Smets Phillips curve. The Gerlach-Smets approach also assumes that the output gap follows a (stationary) AR(2) process, which would not be a good empirical description of the output gap series derived using our preferred Phillips-curve method on Australian data.

vintage of data using $\lambda_{PC} = 20$ and $\lambda_{PC} = 80$. Figure 14 shows real-time output gap estimates (upper panel) and final estimates (lower panel) assuming these alternative parameter values.[17]

Figure 13: Effect of Varying λ_{PC} on the Growth Rate of Potential Output

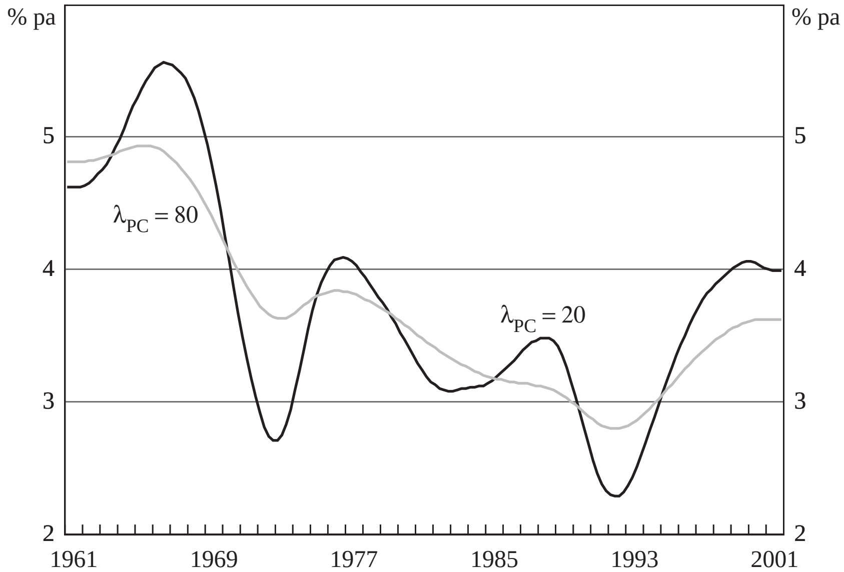

Reducing the value of the parameter λ_{PC} to 20 generates a path for potential output growth with many more wiggles, as expected. It is striking, however, how insensitive are both the real-time and final output gap estimates to the choice of this parameter (the root-mean-square difference between the two real-time series is 1.3 percentage points, and between the two final series, 0.7 ppt). Over plausible ranges for λ_{PC}, the output gap estimates are almost invariant to its value.[18]

---

17 The results assuming $\lambda_{PC} = 80$ are those that have been shown throughout the paper. For each data vintage, the results derived using $\lambda_{PC} = 20$ use the same Phillips-curve specification as for the $\lambda_{PC} = 80$ results (that is, the specifications summarised in Table 1 and Appendix C) although the estimated parameter values in these Phillips curves would be different.

18 It is also of interest to report how large the smoothness parameter, λ_{HP}, in a Hodrick-Prescott filter of actual output needs to be to generate the same smoothness in the derived potential output series as for these values for λ_{PC} in our Phillips curves. We quantify "smoothness" on the basis of the sum $\sum\left((y^*_{t+1} - y^*_t) - (y^*_t - y^*_{t-1})\right)^2$ used in the loss function, equation (2), and report results for the final data vintage. To generate the same potential-output smoothness as for our Phillips curve results with $\lambda_{PC} = 80$, requires a value for λ_{HP} of about 40,000. To generate the same smoothness as for our Phillips curve results with $\lambda_{PC} = 20$, requires a value for λ_{HP} of about 1,600.

Figure 14: Effect of Varying λ_{PC} on the Output Gap

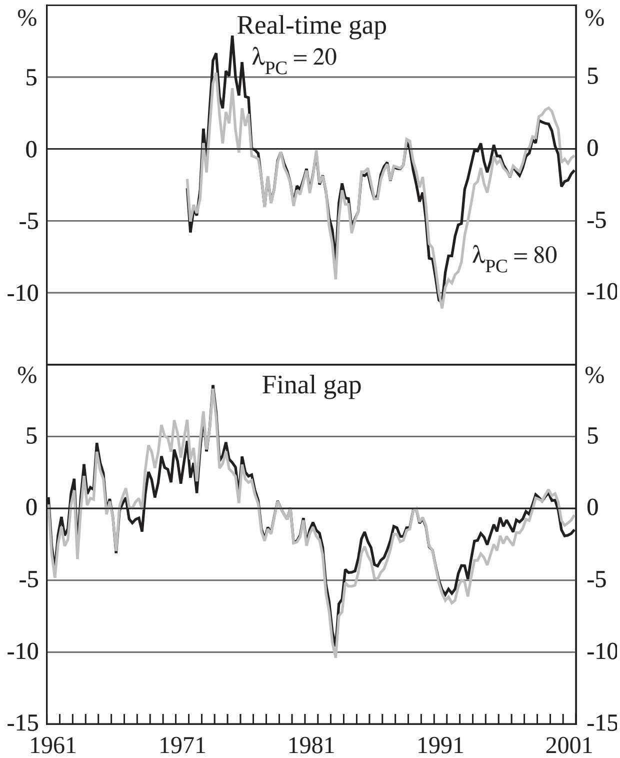

4.2 Confidence Intervals

It is also of interest to ask how large are the confidence intervals around our estimates of the output gap and the rate of growth of potential output. We use a Monte Carlo technique, summarised in Appendix B, to address these questions.

Figure 15 shows two sets of estimated standard errors for the final output gap (upper panel) and for the rate of growth of potential output (lower panel). For each line in the figure, the estimated standard errors rise considerably at each end of the sample, which is simply a manifestation of the end-point problems endemic to estimation procedures of this type.

The results derived using the full set of residuals imply that the –0.5 per cent estimate of the output gap in 2001:4 has an associated standard error of 2.1 per cent, compared with a standard error of 1.1 per cent for output gap estimates near the middle of the sample. The standard error of the annualised growth rate of

Figure 15: Standard Errors

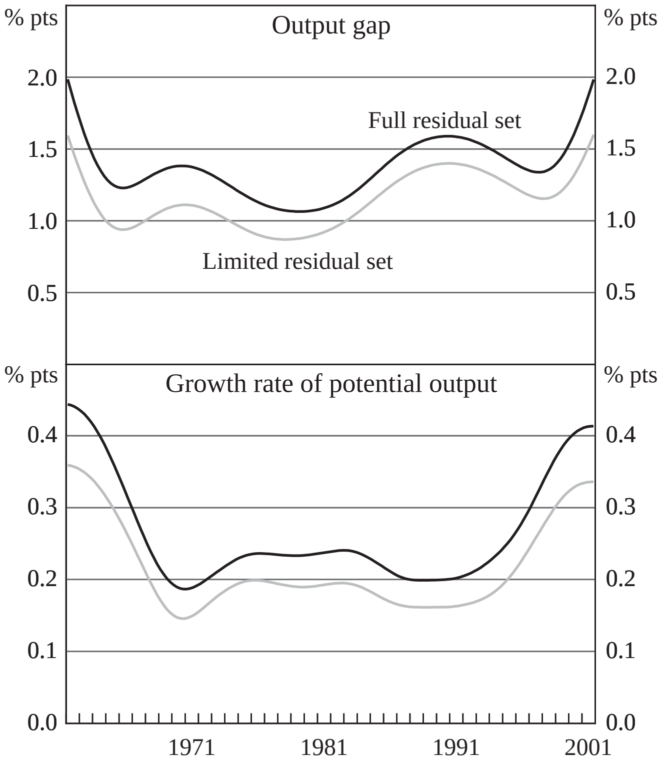

Note: The higher set of standard error estimates in each panel are derived from Monte Carlo simulations in which the full set of Phillips-curve residuals are used in generating pseudo data sets for inflation (see Appendix B). The lower estimates are derived from simulations using only the post-1976 residuals, to reduce the effects of the partly exogenous wage shock in 1973 on the results. The results are derived assuming our standard value $\lambda_{PC} = 80$ for the smoothness parameter. Alternatively, assuming $\lambda_{PC} = 20$ gives a very similar profile of results for the standard error of the gap, but a somewhat higher profile of estimates for the standard error of the growth rate of potential output.

potential output in 2001:4 is 0.4 per cent (the point estimate for potential growth at this time is 3.6 per cent), compared with 0.2 per cent in the middle of the sample.[19]

Figure 16 shows the 95 per cent confidence interval for the final output gap estimates, based on Monte Carlo simulations using the full set of Phillips curve residuals. At this level of confidence, the results imply that the output gap is non-zero in 95 of the 163 quarters of the sample; that is, 58 per cent of the time.[20] These results are very different from those of Orphanides and van Norden (2001), who report that, using the Gerlach-Smets approach, their final output gap estimates

19 When results are derived using only the post-1976 residuals, the estimated standard error of the output gap falls to 1.7 per cent in 2001:4, and 0.9 per cent in the middle of the sample, while for the annualised growth rate of potential output, the standard error falls to 0.3 per cent in 2001:4, and 0.2 per cent in the middle of the sample.

Figure 16: Final Output Gap with 95 Per Cent Confidence Bands

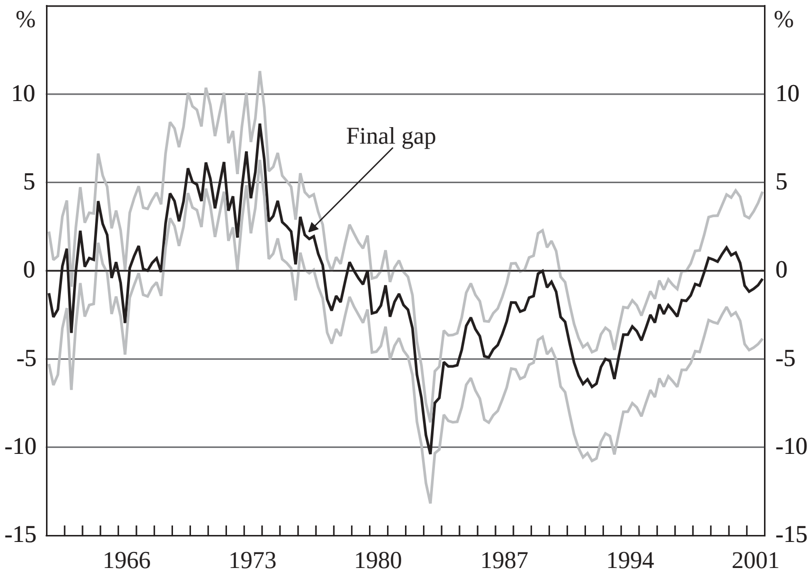

for the US are virtually never significantly different from zero over the period 1966 to 1997.

20 Using the reduced set of residuals, the output gap is non-zero 66 per cent of the time, with 95 per cent confidence. It is of interest to note that the 95 per cent confidence interval in Figure 16 is not symmetrically distributed around the output gap estimates derived from the actual data.

5. Conclusions

We have examined how serious are the problems associated with estimating output gaps in real time. On the basis of results derived from 121 vintages of Australian GDP data, from 1971:4 to 2001:4, we have addressed the questions: How well can we estimate the output gap based solely on information available at the time? And how different are these real-time estimates from estimates generated with the benefit of hindsight?

Our broad conclusion is that quite good estimates of the output gap can be generated in real time. Over the past twenty-eight years, the root-mean-square difference between real-time output-gap estimates derived using our preferred approach and final output gaps estimated with the benefit of hindsight using the latest available data, is less than 2 percentage points. Furthermore, the correlation between these two output-gap series over this time is about 0.8, and the average difference between them is virtually zero, indicating no apparent tendency for the real-time estimates to be biased.

There are three main features of our preferred approach to estimating output gaps in real time. Firstly, we allow for long-lived changes in the rate of growth of potential output. Secondly, we use a Phillips-curve relationship to inform our estimates of potential output, rather than relying on a univariate approach based solely on the behaviour of actual output. And thirdly, we take considerable care in modelling the Phillips curve, allowing for a range of possible influences on inflation in addition to the output gap, and recognising the possibility that the inflation process (or our understanding of it) may change through time as identifiable shocks, of a type not previously seen, hit the economy.

The evidence we have presented suggests that all three of these features are useful in enhancing the prospects of generating reasonably accurate real-time output-gap estimates. The assumption of a constant rate of potential output growth may be an innocuous one over sufficiently short horizons, but can lead to serious error over longer periods. Similarly, estimating potential output using a univariate de-trending method, such as a Hodrick-Prescott filter, may generate quite good output-gap estimates at times, but there should be no general presumption that it will do so in real time. Over the past twenty-eight years, the correlation between real-time output-gap estimates derived using a Hodrick-Prescott filter and our best final estimates of the output gap is virtually zero (in fact, it is slightly negative!).

Finally, if a Phillips-curve approach is to be used to generate output-gap estimates in real time, it appears to be important that the Phillips curve be information-rich and as well specified as possible. Relying on a simple Phillips curve, which uses only limited information about the relationship between inflation and the output gap, can lead to poor real-time estimates of the output gap.

Notwithstanding our general optimism about the possibility of generating quite good output gap estimates in real time, it is appropriate to end on a note of caution. Despite the apparent robustness of our approach to the range of changes in the Australian macroeconomic environment over the past twenty-eight years, there remains an irreducible degree of uncertainty associated with output gaps generated in real time. The problem of "not knowing the future" is still an important one and there will always be times when the best available estimates of the output gap made in real time will turn out, with the benefit of hindsight, to have been badly flawed.

Appendix A: Implementation of the Phillips Curve Approach

Suppose that, over a sample period $t = 1,\dots,n$, we have a model of inflation of the form

$$\Delta\pi_t = \Gamma_t + \gamma(y_t - y_t^*) + \sum_{i=0}^{4} \delta_i \Delta(y_{t-i} - y_{t-i}^*) + \varepsilon_t \tag{A1}$$

where Γ_t is some linear combination of past changes in inflation together with bond market inflation expectations, oil price inflation and import price inflation:

$$\Gamma_t = \sum_j \kappa_j \Delta\pi_{t-j} + \beta_j \text{bond}_{t-j} + \eta_j \text{oil}_{t-j} + \xi_j \text{import}_{t-j} \ . \tag{A2}$$

Note that this specification is simply a re-writing of our generic Phillips curve, equation (1).

With such a model, we wish to minimise the loss function

$$\mathcal{L} = \sum_{t=1}^{n} \varepsilon_t^2 + \lambda_{PC} \sum_{t=-3}^{n-1} ((y_{t+1}^* - y_t^*) - (y_t^* - y_{t-1}^*))^2 \ . \tag{A3}$$

Note that the latter sum in equation (A3) is taken to run from $t = -3$ because equation (A1) involves 4 lags of "change in the output gap" terms. This equation therefore requires values for potential output over the 5 periods (t = -4, -3, -2, -1 and 0) prior to the start of the sample over which it is estimated. Given this, the usual Hodrick-Prescott "smoothing penalty" built into $\mathcal{L}$ is here computed over the period $t = -3, ..., (n-1)$, rather than simply the period $t = 2, ..., (n-1)$.

Overall then we seek values for the $(n+5)\times 1$ vector $Y^* \equiv (y_{-4}^*, y_{-3}^*, \dots, y_n^*)^T$ and for the parameters $\{\kappa_j\}$, $\{\beta_j\}$, $\{\eta_j\}$, $\{\xi_j\}$, $\{\delta_i\}$ and γ which minimise $\mathcal{L}$. Note that $\mathcal{L}$ may itself be written in the form

$$\mathcal{L} = \varepsilon^T\varepsilon + \lambda_{PC} S^T S \tag{A4}$$

where ε denotes the $n\times 1$ vector $\varepsilon \equiv (\varepsilon_1, \varepsilon_2, \dots, \varepsilon_n)^T$ and S denotes the $(n+3)\times 1$ vector $S \equiv ((y_{-2}^* - 2y_{-3}^* + y_{-4}^*), \dots, (y_n^* - 2y_{n-1}^* + y_{n-2}^*))^T$.

The 4 step iterative procedure we employ for computing these values is as follows.

Step 1: Guess at initial values for the parameter γ and for the parameters $\{\kappa_j\}$,

$\{\beta_j\}$, $\{\eta_j\}$, $\{\xi_j\}$ and $\{\delta_i\}$. To do this we simply use the usual Hodrick-Prescott filter to generate a preliminary potential output series, and then, using this series in our model, estimate the corresponding model parameters via ordinary OLS to get initial guesses for these parameters.

Step 2: Using these initial parameter values, solve for the values $\{y_t^*\}_{t=-4}^{n}$ via the appropriate analogue (see below) of the usual "Hodrick-Prescott filter"-type procedure of minimising the loss function $\mathcal{L}$.

Step 3: With these $\{y_t^*\}_{t=-4}^{n}$ re-estimate the inflation equation to get new values for the parameter γ and for the parameters $\{\kappa_j\}$, $\{\beta_j\}$, $\{\eta_j\}$, $\{\xi_j\}$ and $\{\delta_i\}$.

Step 4: Repeat step 2 with these new parameter values, then repeat step 3, and keep doing this until "convergence" is achieved in some suitable sense (that is, until the values of the $\{y_t^*\}_{t=-4}^{n}$ and the parameters in the inflation equation stop changing, to within some pre-specified tolerance threshold).

Technical Details of Step 2

It is useful to begin by introducing some notation. For each $j = 0, 1, 2, \ldots$, let H_j denote the $(n+5) \times (n+5)$ matrix given by

$$(H_j)_{k,l} = \begin{cases} 1 & , \quad k = j + l \\ 0 & , \quad k \neq j + l \end{cases}$$

and let G_j denote the $(n+5-j) \times (n+5)$ matrix given by

$$(G_j)_{k,l} = \begin{cases} 1 & , \quad l = j + k \\ 0 & , \quad l \neq j + k \, . \end{cases}$$

From these core matrices we may then construct, first of all, the $(n+5) \times (n+5)$ "lagged first differencing" matrices D_0, D_1, D_2, D_3 and D_4 given by $D_i \equiv H_i - H_{i+1}$, $i = 0, \ldots, 4$. Thence in turn we may form the trimmed $n \times (n+5)$ versions of these matrices, $\tilde{D}_i$, defined by

$$\tilde{D}_i \equiv G_5 D_i \quad , \; i = 0, \ldots, 4 \, . \tag{A5}$$

The importance of these $\tilde{D}_i$ matrices derives from the fact that, over the sample period $t = 1, \ldots, n$, we may now write equation (A1) in vector form as

$$\Delta\pi = \Gamma + \gamma G_5 (Y - Y^*) + \sum_{i=0}^{4} \delta_i \tilde{D}_i (Y - Y^*) + \varepsilon \tag{A6}$$

where $\Delta\pi$ denotes the $n \times 1$ vector $(\Delta\pi_1, \Delta\pi_2, \ldots, \Delta\pi_n)^T$, Γ the $n \times 1$ vector $(\Gamma_1, \Gamma_2, \ldots, \Gamma_n)^T$, and Y the $(n+5) \times 1$ vector $(y_{-4}, y_{-3}, \ldots, y_n)^T$. Re-arranging equation (A6), and using also definition (A5), then yields that $\varepsilon = \Delta\pi - \Gamma - A(Y - Y^*)$, where A denotes the $n \times (n+5)$ matrix

$$A = G_5 \left(\gamma I_{n+5} + \sum_{i=0}^{4} \delta_i D_i \right) . \qquad \text{(A7)}$$

Now define a new $n \times 1$ vector, Ψ, by $\Psi \equiv \Delta\pi - \Gamma - AY$. Then ε may now be written in the form $\varepsilon = \Psi + AY^*$, where Ψ is independent of potential output. We then obtain the following formula for the first of the two terms on the right hand side of formula (A4) for $\mathcal{L}$:

$$\varepsilon^T \varepsilon = (\Psi + AY^*)^T (\Psi + AY^*) = \Psi^T \Psi + 2\Psi^T AY^* + Y^{*T} A^T AY^* . \qquad \text{(A8)}$$

Turning to the second of the two terms on the right hand side of formula (A4) for $\mathcal{L}$, observe that we may write $S = G_2 (D_0 - D_1) Y^*$, whence also we have that

$$S^T S = Y^{*T} (D_0 - D_1)^T G_2^T G_2 (D_0 - D_1) Y^* . \qquad \text{(A9)}$$

Combining equations (A4), (A8) and (A9) we therefore derive that

$$\mathcal{L} = \Psi^T \Psi + 2\Psi^T AY^* + Y^{*T} A^T AY^* + \lambda_{PC} Y^{*T} (D_0 - D_1)^T G_2^T G_2 (D_0 - D_1) Y^* .$$

The first order conditions for minimising $\mathcal{L}$ then yield that $FY^* = -A^T \Psi$, where F denotes the $(n+5) \times (n+5)$ matrix

$$F \equiv \left(A^T A + \lambda_{PC} (D_0 - D_1)^T G_2^T G_2 (D_0 - D_1) \right) . \qquad \text{(A10)}$$

Therefore, the unique solution for the potential output vector Y^* which minimises $\mathcal{L}$, for the given values of the parameters γ, $\{\kappa_j\}$, $\{\beta_j\}$, $\{\eta_j\}$, $\{\xi_j\}$ and $\{\delta_i\}$, is simply

$$Y^* = -F^{-1} A^T \Psi \qquad \text{(A11)}$$

where F, A and Ψ are as defined above.

Appendix B: Statistical Properties of the Phillips Curve Approach

In this appendix we examine the statistical properties of our iterative procedure for simultaneously estimating our output gap series and the coefficients of our chosen Phillips curve equation. As noted in footnote (5), the usual OLS standard errors and test statistics do not give a true measure of the uncertainty surrounding the coefficient estimates for equation (A1). It is possible *a priori* that our iterative procedure may be a biased estimator of these parameters, given that the potential output series simultaneously constructed by the procedure is itself subject to error.

To test the statistical properties of our iterative estimation procedure we use a bootstrapping approach to construct multiple sets of "pseudo data" for inflation. We can then generate distributions for the estimated values of the parameters in equation (A1), and for the potential output and output gap series, by applying our iterative procedure to each pseudo data set.

Details of the Bootstrapping Procedure

The approach we use is as follows. We describe this for the case of our 2001:4-optimised Phillips curve equation (see Table 2).

Step 1: First, we take the potential output series $\{\hat{y}_t^*\}$, parameter estimates $\{\hat{\gamma}, \hat{\delta}_i, \hat{\kappa}_j, \hat{\beta}_j, \hat{\eta}_j, \hat{\xi}_j\}$, and set of inflation residuals $\{\hat{\varepsilon}_t\}$, obtained from iterative estimation of equation (A1) using the actual inflation data.

Step 2: We then generate 1000 alternative sets of pseudo data for inflation residuals, $\left\{\{\hat{\varepsilon}_{t,i}\}\right\}_{i=1}^{1000}$, each set obtained by making 163 random draws (with replacement) from the set of initial inflation residuals $\{\hat{\varepsilon}_t\}$.

Step 3: For each $i = 1, \ldots, 1000$ we then generate a corresponding set of pseudo data for inflation by using equation (A1) – with the $\{\hat{y}_t^*\}$ for $\{y_t^*\}$, with the $\{\hat{\gamma}, \hat{\delta}_i, \hat{\kappa}_j, \hat{\beta}_j, \hat{\eta}_j, \hat{\xi}_j\}$ for $\{\gamma, \delta_i, \kappa_j, \beta_j, \eta_j, \xi_j\}$, but with the $\{\hat{\varepsilon}_{t,i}\}$ rather than the $\{\hat{\varepsilon}_t\}$ for $\{\varepsilon_t\}$ – to successively generate new pseudo data from 1961:2 onwards. In applying equation (A1) we use the actual data for other variables, as well as for the pre-1961:2 data for inflation.

Step 4: For each of these 1000 sets of pseudo data for inflation we then re-run our iterative procedure for estimating equation (A1), thereby obtaining for each a corresponding alternative potential output series $\{\hat{y}_t^{*(i)}\}$; alternative set of parameter estimates $\{\hat{\gamma}^{(i)}, \hat{\delta}_i^{(i)}, \hat{\kappa}_j^{(i)}, \hat{\beta}_j^{(i)}, \hat{\eta}_j^{(i)}, \hat{\xi}_j^{(i)}\}$; and alternative set of estimated inflation residuals $\{\hat{\varepsilon}_t^{(i)}\}$. We can then study the resulting distributions obtained from these 1000 simulations for each parameter value and for the $\{y_t^*\}$ series.

For reasons discussed in Section 4, it is also of interest to examine Monte Carlo simulations where only the more limited set of (say) post-1976 residuals are used, rather than the full set of residuals.

Results of the Monte Carlo Simulations

Figures 15 and 16 from Section 4 and Figure B1 below summarise our findings.

Figure B1: Histogram of Coefficient Estimates on the Output Gap
Based on 1000 Monte Carlo Simulations (Using the Full Set of Residuals)

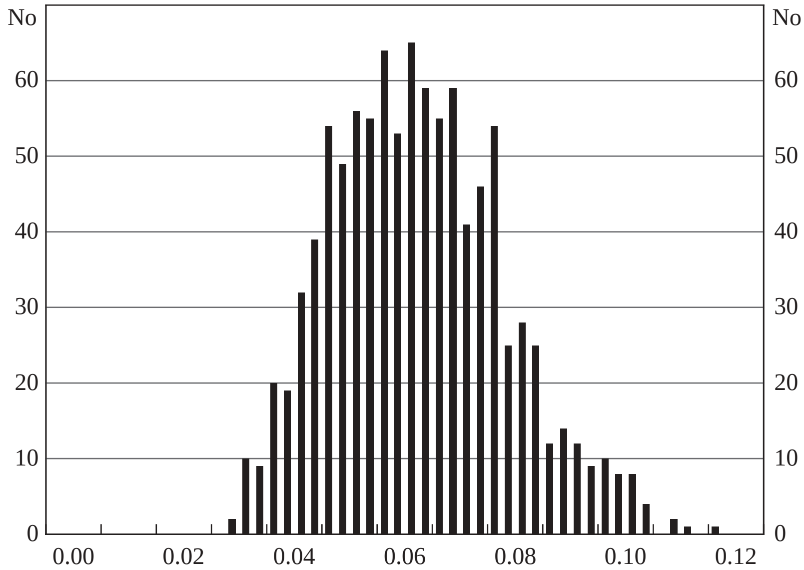

Note: While the simulated γ-values do not appear to be normally distributed, there is no apparent bias, with the mean of the 1000 simulated values for γ being 0.0649 and the median 0.0639 – both almost identical to the "true" value used to generate the pseudo data for inflation in the simulations, $\gamma = 0.0642$. A similar result holds for the other parameters in the Phillips curve equation.

Finally, all the results reported here are for our final Phillips curve equation. The same tests were, however, carried out for a selection of other data vintages, with their corresponding Phillips curve specifications. In all cases the same general conclusions were found to hold.[21]

Appendix C: Detailed Phillips Curve Specifications

Specifications for the Preferred Phillips Curve Method

The five broad specification types for the preferred Phillips curve method are summarised in Table 1 in Section 2. Table C1 below provides a complete listing of the specifications used for each of our 121 data vintages.

The specifications over the period 1971:4 to 1973:3 are the only ones which do not contain any oil-price-inflation terms. This presumably reflects the relative lack of movement in oil prices before the OPEC I oil shock (see Figure 1). Our inability to identify a separate role for oil prices in these early data vintages may also reflect the fact that our measure of import prices before 1985 does not exclude oil (see Appendix D). Hence, oil-price effects may be captured in these early-vintage equations indirectly through lagged import-price-inflation terms, rather than showing up directly.

The first appearance of oil-price-inflation terms in our Phillips curve specifications occurs in 1973:4. Frequent re-specifications are required over the following two years of data vintages because of the extreme volatility in oil prices over this period. These re-specifications principally involve changes in the lags of oil-price-inflation terms with the inclusion, for 1974:4 and 1975:4 vintages, of that lag which includes the near quadrupling of Australian-dollar oil prices in 1974:1. Interestingly, despite the continued volatility of oil prices over the remainder of the 1970s, the 1975:4 specification continues to perform well until the mid-1980s – and when the Phillips curve specification is next changed in 1986:3, the required modifications are relatively minor.

Finally, the introduction of chain-linking in the National Accounts in 1998:3, and the associated switch from *System of National Accounting (SNA) 1968* to *SNA 1993* as the basis on which the accounts are prepared, results in significant revisions to the entire history of real GDP. As a consequence of these revisions,

[21] The main difference is that, for some early vintages, our iterative procedure fails to converge in some of the Monte Carlo simulations. This is especially true of vintages from around the mid-1970s, when convergence is not achieved for quite a high proportion of simulations. This reflects both the short length of these early data vintages (which makes equation estimation more difficult), and the higher fraction of these smaller data sets consisting of data from the period of relative economic turmoil in the first half of the 1970s.

Table C1: **Complete List of Specifications for Preferred Phillips Curve Method**

Date of Vintage	Precise Equation Specification
1971:4 to 1972:4	$\pi_t = 0.5(\pi_{t-2}+\pi_{t-3})+\beta_2\text{bond}_{t-2}+\beta_3\text{bond}_{t-3}+\beta_4\text{bond}_{t-4}+$ $\gamma(y_t-y_t^*)+\delta_4\Delta(y_{t-4}-y_{t-4}^*)+$ $\xi_3\text{import}_{t-3}+\xi_4(\text{import}_{t-4}+\text{import}_{t-5}+\text{import}_{t-6})$
1973:1 to 1973:2	$\pi_t = 0.5(\pi_{t-2}+\pi_{t-3})+\beta_2\text{bond}_{t-2}+\beta_3\text{bond}_{t-3}+\beta_4\text{bond}_{t-4}+$ $\gamma(y_t-y_t^*)+$ $\xi_3\text{import}_{t-3}+\xi_4(\text{import}_{t-4}+\text{import}_{t-5}+\text{import}_{t-6})$
1973:3	$\pi_t = 0.5(\pi_{t-2}+\pi_{t-3})+\beta_2\text{bond}_{t-2}+\beta_3\text{bond}_{t-3}+\beta_4\text{bond}_{t-4}+$ $\gamma(y_t-y_t^*)+\xi_3\text{import}_{t-3}$
1973:4 to 1974:2	$\pi_t = 0.5(\pi_{t-2}+\pi_{t-3})+\beta_2\text{bond}_{t-2}+\beta_3\text{bond}_{t-3}$ $\gamma(y_t-y_t^*)+\eta_3\text{oil}_{t-3}$
1974:3	$\pi_t = 0.25(\pi_{t-2}+\pi_{t-3}+\pi_{t-4}+\pi_{t-5})+\beta_1\text{bond}_{t-1}+$ $\gamma(y_t-y_t^*)+\eta_2\text{oil}_{t-2}$
1974:4 to 1975:3	$\pi_t = 0.25(\pi_{t-2}+\pi_{t-3}+\pi_{t-4}+\pi_{t-5})+\beta_1\text{bond}_{t-1}+$ $\gamma(y_t-y_t^*)+\eta_2\text{oil}_{t-2}+\eta_3\text{oil}_{t-3}$
1975:4 to 1986:2	$\pi_t = 0.25(\pi_{t-2}+\pi_{t-3}+\pi_{t-4}+\pi_{t-5})+\beta_1\text{bond}_{t-1}+$ $\gamma(y_t-y_t^*)+\eta_2\text{oil}_{t-2}+\eta_3\text{oil}_{t-3}+\eta_7\text{oil}_{t-7}$
1986:3 to 1998:2	$\pi_t = 0.25(\pi_{t-2}+\pi_{t-3}+\pi_{t-4}+\pi_{t-5})+$ $\zeta_2(\pi_{t-2}-\pi_{t-6})+\zeta_3(\pi_{t-3}-\pi_{t-7})+$ $\beta_1\text{bond}_{t-1}+\beta_2\text{bond}_{t-2}+$ $\gamma(y_t-y_t^*)+\eta_2\text{oil}_{t-2}+\eta_3\text{oil}_{t-3}+\eta_7\text{oil}_{t-7}$
1998:3 to 2001:4	$\pi_t = 0.25(\pi_{t-1}+\pi_{t-2}+\pi_{t-3}+\pi_{t-4})+$ $\zeta_2(\pi_{t-2}-\pi_{t-6})+\beta_1\text{bond}_{t-1}+$ $\gamma(y_t-y_t^*)+\eta_2\text{oil}_{t-2}+\eta_7\text{oil}_{t-7}+$ $\xi_0\text{import}_t+\xi_1\text{import}_{t-1}$

Note: Start of sample for all regressions is 1961:2.

re-specification of our Phillips curve is required, with the coefficient on the output gap falling considerably for the 1998:3 specification relative to that for 1998:2, although part of this fall is retraced in subsequent vintages.

The optimal Phillips curve specification for the final data vintage (2001:4) is explicitly checked, since we regard the results from this vintage as giving us our best available estimate of the output gap over history. As it turned out, no further change in specification is required for this final vintage. Unlike in 1973 and 1974, the considerable rise in the oil price in the second half of 2000 seems not to have substantially affected the performance of the 1998:3-optimised specification, with the effects of this rise captured by the terms already included in this equation.

Specifications of Constant Potential-Output-Growth Phillips Curves

Here, potential output is assumed to follow a simple linear trend, $y_t^* = a + bt$. We then conduct specification searches for optimal Phillips curves for each data vintage, following the same approach as for the preferred Phillips curves. The resulting optimal Phillips curve specifications are set out in Table C2. The estimation results for the final-vintage Phillips curve are shown in Table C3. While the equation appears quite impressive in terms of goodness of fit, the derived output gap estimates appear very poor, as the results in Tables 3 and 4 make clear.

Table C2: Complete List of Constant-Potential-Growth Phillips Curve Specifications

Date of Vintage	Precise Equation Specification
1971:4 to 1973:3	$\pi_t = 0.449\pi_{t-2} + 0.565\pi_{t-7} - 0.014\pi_{t-8} + \beta_7\text{bond}_{t-7} + \gamma(y_t - [a+bt]) + \delta_4\Delta(y_{t-4} - [a+bt]_{t-4}) + \eta_6\text{oil}_{t-6} + \eta_7\text{oil}_{t-7} + \xi_3\text{import}_{t-3}$
1973:4 to 1974:2	$\pi_t = 0.491\pi_{t-1} + 0.509\pi_{t-2} + \gamma(y_t - [a+bt]) + \eta_2\text{oil}_{t-2} + \eta_3\text{oil}_{t-3} + \eta_4\text{oil}_{t-4} + \eta_7\text{oil}_{t-7} + \eta_8\text{oil}_{t-8} + \xi_3\text{import}_{t-3} + \xi_5\text{import}_{t-5}$
1974:3	$\pi_t = 0.313\pi_{t-1} + 1.603\pi_{t-2} - 0.916\pi_{t-6} + \beta_1\text{bond}_{t-1} + \beta_2\text{bond}_{t-2} + \gamma(y_t - [a+bt]) + \eta_2\text{oil}_{t-2} + \xi_3\text{import}_{t-3}$
1974:4	$\pi_t = 0.313\pi_{t-1} + 1.603\pi_{t-2} - 0.916\pi_{t-6} + \beta_1\text{bond}_{t-1} + \beta_2\text{bond}_{t-2} + \gamma(y_t - [a+bt]) + \eta_2\text{oil}_{t-2} + \eta_3\text{oil}_{t-3} + \xi_3\text{import}_{t-3}$
1975:1 to 1975:3	$\pi_t = 0.313\pi_{t-1} + 1.603\pi_{t-2} - 0.916\pi_{t-6} + \beta_1\text{bond}_{t-1} + \beta_2\text{bond}_{t-2} + \gamma(y_t - [a+bt]) + \eta_2\text{oil}_{t-2} + \eta_3\text{oil}_{t-3} + \xi_3\text{import}_{t-3} + \xi_7\text{import}_{t-7}$
1975:4 to 1982:3	$\pi_t = 0.313\pi_{t-1} + 1.603\pi_{t-2} - 0.916\pi_{t-6} + \beta_1\text{bond}_{t-1} + \beta_2\text{bond}_{t-2} + \gamma(y_t - [a+bt]) + \eta_2\text{oil}_{t-2} + \eta_3\text{oil}_{t-3} + \eta_7\text{oil}_{t-7} + \xi_3\text{import}_{t-3} + \xi_7\text{import}_{t-7}$
1982:4 to 1986:2	$\pi_t = 0.245\pi_{t-1} + 0.854\pi_{t-2} + 0.333\pi_{t-3} - 0.432\pi_{t-6} + \beta_1\text{bond}_{t-1} + \beta_2\text{bond}_{t-2} + \gamma(y_t - [a+bt]) + \eta_2\text{oil}_{t-2} + \eta_3\text{oil}_{t-3} + \eta_7\text{oil}_{t-7}$
1986:3 to 1998:2	$\pi_t = 0.354\pi_{t-1} + 0.789\pi_{t-2} + 0.301\pi_{t-3} - 0.444\pi_{t-6} + \beta_1\text{bond}_{t-1} + \beta_2\text{bond}_{t-2} + \gamma(y_t - [a+bt]) + \eta_2\text{oil}_{t-2} + \eta_4\text{oil}_{t-4}$
1998:3 to 2001:3	$\pi_t = 0.364\pi_{t-1} + 0.243\pi_{t-2} + 0.302\pi_{t-3} + 0.091\pi_{t-4} + \zeta(\pi_{t-2} - \pi_{t-6}) + \beta_1\text{bond}_{t-1} + \beta_2\text{bond}_{t-2} + \gamma(y_t - [a+bt]) + \eta_2\text{oil}_{t-2} + \eta_7\text{oil}_{t-7}$
2001:4	$\pi_t = 0.365\pi_{t-1} + 0.241\pi_{t-2} + 0.303\pi_{t-3} + 0.091\pi_{t-4} + \zeta(\pi_{t-2} - \pi_{t-6}) + \beta_1\text{bond}_{t-1} + \beta_2\text{bond}_{t-2} + \gamma(y_t - [a+bt]) + \eta_2\text{oil}_{t-2} + \eta_7\text{oil}_{t-7}$

Note: Start of sample for all regressions is 1961:2. In the first quarter of each new equation specification, the coefficients on the lags of quarterly inflation are estimated, and these coefficient estimates are kept unchanged until a new equation specification is deemed appropriate. This approach is the same as that used for the preferred Phillips curve method.

Table C3: Estimation Results for the Final-Vintage Constant-Potential-Growth Phillips Curve

$$\pi_t = 0.365\pi_{t-1} + 0.241\pi_{t-2} + 0.303\pi_{t-3} + 0.091\pi_{t-4} + \zeta(\pi_{t-2} - \pi_{t-6}) + \beta_1 \text{bond}_{t-1} + \beta_2 \text{bond}_{t-2} + \gamma(y_t - [a + bt]) + \eta_2 \text{oil}_{t-2} + \eta_7 \text{oil}_{t-7}$$

Coefficient	Value	t-statistic
ζ	0.358	4.589
β_1	1.295	5.034
β_2	-1.026	-3.925
γ	0.015	2.017
a	10.60	N/A
b	0.009	N/A
η_2	0.008	3.781
η_7	0.006	2.603
Summary Statistics	**Value**	
R^2	0.819	
Adjusted R^2	0.811	
Standard error of the regression	0.004	
Breusch-Godfrey LM test for autocorrelation (p-value):		
First order	0.976	
First to fourth order	0.537	

Note: The sample is 1961:2 - 2001:4 ($n = 163$).

Appendix D: Data Sources and Definitions

Real-Time Real GDP

The Australian real-time GDP data are described in Stone and Wardrop (2002). For each vintage we use the hybrid series, which Stone and Wardrop (2002) consider represents what analysts at the time would have considered the best measure of GDP. Specifically, this corresponds to: GDP(I) for each vintage from 1971:4-1991:3; GDP(A) for each vintage from 1991:4-1998:2; and chain-volume GDP for each vintage from 1998:3-2001:4. The bulk of our real-time output data set was supplied to us as unofficial data by Peter Rossiter from the ABS. This data set was then supplemented by the manual entry of additional information from historical tables contained in original hard-copy National Accounts releases, where available.

For vintages where the data did not go back to 1959:3, these series were back-cast wherever possible based on the growth rates from the most recent preceding vintage for which these data were available.

Consumer Prices

The price series we use is a measure of core consumer prices, the weighted-median CPI, which is calculated by the RBA and is available back to 1976:3 from *Bulletin* (RBA) Table G1. The series is adjusted for the introduction of a Goods and Services Tax in 2000:3, assuming that the tax led to a 2.95 per cent rise in the weighted-median CPI in that quarter.

The weighted-median CPI is extended to cover the period from 1966:3 to 1976:2 by direct construction using data from *Consumer Price Index* (ABS Cat No 6401.0) [2001:4], *Consumer Price Index Particulars for Sub-groups and Special Groupings* (Commonwealth Bureau of Census and Statistics (CBS) Ref No 9.7) [1966:3-1971:2], *Labour Report* No 52, (CBS Ref No 6.7) [1965&1966], *The Australian Consumer Price Index, Concepts Sources and Methods* (ABS Cat No 6461.0), and additional data from the ABS.

Before 1966:3, the weighted-median series is back-cast using All Groups CPI excluding interest payments from the December quarter 2001 *Consumer Price Index* (ABS Cat No 6401.0).

Import Prices

From 1985:3 onwards, import prices are the implicit price deflator for merchandise imports, excluding fuels and lubricants, civil aircraft and RBA imports of gold, from *National Income, Expenditure and Product* (ABS Cat No 5206.0) [2001:4].

Before 1985:3, the series is back-cast using the import implicit price deflator from *National Income, Expenditure and Product* (ABS Cat No 5206.0).

The series we use is adjusted for tariffs (but not for the Balassa-Samuelson effect) using the approach described in Appendix C of Beechey *et al* (2000). For 1969:3-2001:4 the tariff rate is customs duty receipts divided by the value of merchandise imports (excluding fuels and lubricants, civil aircraft and RBA imports of gold), seasonally adjusted. Tariff revenue is from the Australian

Customs Service. For 1959:3 -1969:2 customs duties receipts are sourced from *Overseas Trade* (CBS) [1959-60 - 1968-69]. These annual figures are linearly interpolated to obtain a quarterly series.

Oil Prices

The US$ oil price is the average-quarter value of the price per barrel of West Texas Intermediate crude. For 1982:1 onward, this is sourced from the nearest contract price on Bloomberg, CL1 CMDTY. For the period before this it is back-cast using the average quarterly spot price for crude from the International Monetary Fund (IMF) *International Financial Statistics* (IFS) database (Datastream code WDI76AAZA).

The US$ oil price is converted to A$ using a quarter-average AUD/USD bilateral exchange rate. For 1970:1 onwards this is obtained from the IMF IFS database (Datastream code AUI..RF.). For the period before 1970:1, quarterly data are generated by linear interpolation of annual figures for the AUD/USD bilateral rate from Foster (1997), Table 1.19a, (available at <*http://www.rba.gov.au/Statistics/op8_index.html*>).

Inflation Expectations From the Bond Market

The bond market inflation expectations series is derived by splicing together two different series. For the period from 1993:1 onwards we use the difference in the yield between a 10 year government bond and an indexed bond of the same duration. For the period prior to 1993, for the bulk of which indexed bonds were not issued by the Commonwealth, we use a variant of the approach used by Debelle and Laxton (1997) to generate inflation expectations estimates.

Following Debelle and Laxton (1997), we construct an Australian equilibrium real 10-year government bond rate series, r_t^*, based on the ratio of the stock of OECD net public debt to GDP, $debt_t$:

$$r_t^* = C + \beta * debt_t \tag{D1}$$

and use Debelle and Laxton's value of 0.07 for β, which was based on the work of Tanzi and Fanizza (1995), so that a 1 percentage point increase in $debt_t$ increases r_t^* by 7 basis points.

The constant C in equation (D1) is a sum of the world real interest rate when OECD net public debt is zero and an Australia-specific risk premium, which we assume is constant. We choose a reference quarter, 1959:3, and assume that inflation expectations in that quarter were equal to average year-ended inflation for the preceding two years, since it was a period of quiescent inflation. Then setting the nominal Australian bond rate, i_t, equal to the sum of the Australian equilibrium real government bond rate, r_t^*, and inflation expectations, π_t^b, in that quarter, gives a value for C of -0.186.

These choices for β and C yield a series for inflation expectations for the period 1959:3 to 2001:4, based on the assumption:

$$\pi_t^b = i_t - r_t^* . \tag{D2}$$

This series is then spliced together with our inflation expectations series from indexed bond data, with the latter replacing the former in 1993:1, when the two measures differ by only 15 basis points.

As regards data sources, we use the end-quarter Australian 10-year Government bond yield from *Bulletin* (RBA) Table F2. Australian Treasury capital indexed bond yields are from Bloomberg (screen: ILB).

An annual series for the OECD net public debt to GDP ratio is sourced from OECD Online Information Services (OLISnet). This series starts in 1970, and over the 1970s and 1980s moves quite closely with US General Government Net Financial Liabilities as a ratio of Nominal GDP from *OECD Economic Outlook Database*, Annex Table 34 (Datastream code USOCFNF%). In light of this co-movement, for the period 1960-1970 we back-cast our OECD debt to GDP series based on changes in this US-debt-to-GDP series. For the few years in our database before 1960, we assume that the OECD net public debt to GDP ratio is constant. To obtain a quarterly series, the annual series thus constructed is linearly interpolated.

Real-time Issues

There are a few real-time issues relating to these data. First, we use the final (2001:4) vintage of consumer price data throughout. While Australian CPI data are not subject to revision, periodic re-basing of the index may, for early periods, have resulted in very minor differences between inflation rates reported in real-time and those reported in 2001:4, resulting from the ABS's practise of rounding the index to one decimal place. These differences, however, are negligible.

Secondly, the 2001:4 vintage of import prices is also used throughout. Inspection of several vintages of the import price data suggests that, while revisions do occur, they tend to be small, especially relative to real GDP revisions, and are therefore unlikely to materially affect our results.

Finally, the construction of the bond market inflation expectations series requires an estimate of β, the assumed sensitivity of the equilibrium real 10 year bond rate to increases in OECD public debt. We use the value $\beta = 0.07$, based on Tanzi and Fanizza (1995), which was clearly unavailable before that time. Likewise, our use of a single, current series for the ratio of OECD public debt to GDP in our calculation of pre-1993 bond market inflation expectations is also, strictly speaking, subject to a real-time problem, as it neglects revisions over time to estimates of OECD GDP.

We have, however, examined the sensitivity of our results to the chosen value of β. Assuming a value twice as large, $\beta = 0.14$ (which also requires an adjustment to the constant, C, assuming that inflation expectations in the reference quarter, 1959:3, are equal to average year-ended inflation for the preceding two years, as before), leads to real-time output gap estimates with a root-mean-square difference of 0.5 ppt from real-time estimates assuming our standard value, $\beta = 0.07$, over the period, 1971:4 to 2001:4. The corresponding RMSD between the final output gap estimates over the same period is 0.8 ppt. We conclude that even large changes in the value of β lead to only small changes in our estimated output gaps.

References

Beechey M, N Bharucha, A Cagliarini, D Gruen and C Thompson (2000), 'A Small Model of the Australian Macroeconomy', Reserve Bank of Australia Research Discussion Paper No 2000-05.

Debelle G and D Laxton (1997), 'Is the Phillips Curve Really a Curve? Some Evidence for Canada, the United Kingdom, and the United States', *IMF Staff Papers*, 44(2), pp 249–282.

Foster R (1997), *Australian Economic Statistics 1949-50 to 1994-95*, 2nd edn, Reserve Bank of Australia, Sydney.

Laxton D and R Tetlow (1992), 'A Simple Multivariate Filter for the Measurement of Potential Output', Bank of Canada Technical Report No 59.

McCallum B (1995), 'New Zealand's Monetary Policy Arrangements: Some Critical Issues', Victoria University of Wellington Working paper.

McCallum B (2001), 'Should Monetary Policy Respond Strongly to Output Gaps?', *American Economic Review*, 91(2), pp 258–262.

Nelson E and K Nikolov (2001), 'UK Inflation in the 1970s and 1980s: The Role of Output Gap Mismeasurement', Bank of England Working Paper No 148.

Orphanides A (2000), 'The Quest for Prosperity Without Inflation', European Central Bank Working Paper No 15.

Orphanides A and S van Norden (2001), 'The Unreliability of Output Gap Estimates in Real-time', Centre for Interuniversity Research and Analysis of Organisations Working Paper No 2001s-57.

Sargent TJ (1971), 'A Note on the "Accelerationist" Controversy', *Journal of Money, Credit and Banking*, 3(3), pp 721–725.

Stone A and S Wardrop (2002), 'Real-Time National Accounts Data'.

Tanzi V and D Fanizza (1995), 'Fiscal Deficit and Public Debt in Industrial Countries, 1970-1994', International Monetary Fund Working Paper No 95/49.

Taylor J (1993), ‘Discretion versus Policy Rules in Practice’, *Carnegie-Rochester Conference Series on Public Policy*, 39, pp 195–214.

Taylor J (1999), *Monetary Policy Rules*, University of Chicago Press, Chicago.

Taylor J (2000), ‘Comments on Athanasios Orphanides’ The Quest for Prosperity Without Inflation’, Powerpoint presentation slides, available at < *http* : //*www.stanford.edu/ johntayl/PptLectures* >.