



RESERVE BANK
OF AUSTRALIA

Research Discussion Paper

Okun's Law and Potential Output

David Lancaster and Peter Tulip

RDP 2015-14

The Discussion Paper series is intended to make the results of the current economic research within the Reserve Bank available to other economists. Its aim is to present preliminary results of research so as to encourage discussion and comment. Views expressed in this paper are those of the authors and not necessarily those of the Reserve Bank. Use of any results from this paper should clearly attribute the work to the authors and not to the Reserve Bank of Australia.

Enquiries:

Phone: +61 2 9551 9830

Facsimile: +61 2 9551 8033

Email: rbainfo@rba.gov.au

Website: <http://www.rba.gov.au>

The contents of this publication shall not be reproduced, sold or distributed without the prior consent of the Reserve Bank of Australia and, where applicable, the prior consent of the external source concerned. Requests for consent should be sent to the Head of Information Department at the email address shown above.

ISSN 1448-5109 (Online)

Okun's Law and Potential Output

David Lancaster and Peter Tulip

Research Discussion Paper
2015-14

December 2015

Economic Research Department
Reserve Bank of Australia

We would like to acknowledge earlier internal work at the RBA by Mark Baker and Alex Cooper, on which we have relied. We would also like to thank Adam Cagliarini, Alex Cooper, Kathryn Davis, David Jacobs, Christopher Kent, Michael Plumb, Matthew Read, Dan Rees and John Simon for helpful comments. The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Reserve Bank of Australia. The authors are solely responsible for any errors.

Author: tulipp at domain rba.gov.au

Media Office: rbainfo@rba.gov.au

Abstract

We find that Okun's law provides a simple and accurate means of understanding and predicting changes in the unemployment rate in Australia. Okun's law also implies a rate of output growth consistent with stable unemployment, called the growth of potential output. Our estimates of potential output growth are imprecise and fluctuate over time. A recent estimate is a bit below 3 per cent a year, with a \pm one standard error band covering the range $2\frac{1}{4}$ to $3\frac{3}{4}$ per cent. This is a percentage point or two below estimates from before the mid 1990s.

JEL Classification Numbers: E24, E27, J64

Keywords: Okun's law, potential output, unemployment

Table of Contents

1.	Introduction	1
2.	Comparisons with Previous Research	3
2.1	Measures of Potential Output	3
2.2	Okun's Law	5
3.	Models	7
3.1	Constant Coefficients Model	7
3.2	Time-varying Coefficients Model	8
4.	Estimation	8
4.1	Data	9
4.2	Results	9
4.3	Time-varying Parameters	12
5.	Forecasting	16
6.	Specification Issues	21
6.1	Asymmetry	21
6.2	Changes in the NAIRU	24
6.3	Levels of Potential Output and Equilibrium Unemployment	24
7.	Directions for Future Research	26
	References	27

Okun's Law and Potential Output

David Lancaster and Peter Tulip

1. Introduction

'Okun's law' refers to the strong correlation observed between the unemployment rate and real gross domestic product (GDP). It holds in changes and in simple transformations of the level. This paper uses Okun's law to predict changes in Australian unemployment. We also use Okun's law to estimate the growth rate of potential output as a time-varying parameter. We estimate that our definition of potential output growth has fallen from around 5 per cent in the 1970s to 2.9 per cent in 2015. An approximate 68 per cent confidence interval about the recent estimate spans 2.2 to 3.7 per cent.

Since it was popularised by Okun (1962), the term 'potential output' has been applied to a variety of different ideas. It can represent the permanent or smoothed component of GDP; the level of output that would obtain if prices and wages were flexible; the level of output at which resources are fully utilised; the level of output at which inflation is stable; the level of output at which some other criterion (like financial sustainability) is satisfied; and so on. Basu and Fernald (2009) and Kiley (2013) discuss the relationship between various definitions. In this paper we use a definition of potential output growth employed in the Okun's law literature: the rate of GDP growth at which the unemployment rate is stable. If the unemployment rate proxies for the overall level of resource utilisation or for the persistence of inflationary pressures, then our measure also indicates the rate of GDP growth at which these concepts are stable. We discuss our definition in Section 2.1.

It is not common for estimates of potential output to be formally evaluated. Accordingly, it can be unclear whether they are useful relative to alternatives and why they are designed the way they are. The lack of testing partly reflects the abstract, multiple or unspecified objectives of these measures, all of which make evaluation difficult. In contrast, we clearly specify the purpose of our measure, in a manner that facilitates evaluation. Our measure of potential is intended to (a) explain changes in unemployment when GDP is known, and (b) predict changes in unemployment given predictions of GDP growth. We evaluate our

measure on these criteria and show that it is useful. We also have broader objectives, though these are harder to evaluate. Our estimates imply a rate of longer-term GDP growth at which the unemployment rate is expected to be stable. And if stability of the unemployment rate is a policy objective, our estimates provide a benchmark for assessing GDP growth. Our understanding is that these purposes – especially forecasting unemployment – encompass the main ways that many central banks use estimates that they call ‘potential output’.¹

As background, over 2014, Australian GDP grew by 2½ per cent, which the Reserve Bank of Australia (RBA) (RBA 2015, p 64) described as ‘a bit below its trend rate’. Comparisons like this are common in RBA publications, particularly in discussions of the forecasts. Although several interpretations of this ‘trend’ are possible, one consistent with our approach is that it reflects the rate of GDP growth that normally would keep the unemployment rate constant. As discussed below, we estimate that potential GDP was growing 2.9 per cent per year in 2014. Of course, the assessments of the RBA reflect a wider range of information than the simple model presented in this paper. Recently, for example, population data have also been taken into account.

In addition to providing an estimate of ‘potential output growth’ (as defined above), our paper explores the negative relationship between output and unemployment, known as Okun’s law. Some of our key findings are:

- Okun’s law provides a useful description of changes in the Australian unemployment rate over the past 50 years.
- Forecasts of changes in the unemployment rate based on Okun’s law are about as accurate as RBA forecasts. In contrast to RBA forecasts, recent forecasts by Okun’s law have been unbiased.
- The short-run relationship between output and the unemployment rate has been stable over time.

¹ Some estimates of the output gap have been motivated by a desire to forecast inflation. However, conversations with inflation forecasters and our reading of the literature do not indicate that output gap estimates provide useful information about the behaviour of inflation beyond that contained in the unemployment gap. See, for example, Norman and Richards (2010).

- Allowing potential GDP growth to vary over time significantly improves simpler versions of Okun's law.
- The rate of GDP growth consistent with a stable unemployment rate has fallen from around 5 per cent in the 1970s to around 2.9 per cent in 2015.
- Changes in the unemployment rate have become increasingly persistent over time.
- Okun's law is able to explain most of the asymmetry in unemployment rate changes.
- Okun's law holds in changes. We do not find that relationships in levels provide information beyond that in changes.
- We do not find that estimates of the level or change in the NAIRU are helpful for explaining or predicting changes in unemployment.

2. Comparisons with Previous Research

2.1 Measures of Potential Output

A simple measure of potential output growth is the average rate of GDP growth. A linear trend of the logarithm of GDP is very similar. One difficulty with these measures is that average rates of GDP growth vary across time, and hence these estimates will be sensitive to the sample period. For example, real GDP growth averaged 5 per cent in the 1960s, around 3.2 per cent in the 1970s, 1980s and 1990s, but only 2.9 per cent since 2000. Assuming that more recent observations are more relevant, this instability could be accommodated by using a shorter moving average or a statistical filter, such as Hodrick-Prescott (henceforth HP). However, these approaches have two important problems. First, they involve an arbitrary judgement about either the length of the moving average or the degree of smoothness in the HP trend. Second, as more weight is placed on recent outcomes, the more will estimates of potential reflect temporary changes, such as the business cycle and noise.

A common motivation across most definitions of potential output is to remove temporary variations. In our assessment, the most important temporary factor associated with Australian GDP growth in recent decades has been changes in the unemployment rate. Cyclical variations in unemployment are well known but longer-term variations are also important. Unemployment averaged 8.8 per cent in the 1990s and 5.4 per cent over the past 10 years. On the assumption that this downward trend will not continue, the GDP growth that accompanied it is an over-estimate of growth going forward. Accordingly, it is desirable to remove this effect from an estimate of potential output growth.

A popular method for removing the effect of variations in unemployment is to construct ‘production function’ estimates of potential output. The non-accelerating inflation rate of unemployment (NAIRU) is typically estimated from a Phillips curve regression. Then the NAIRU-consistent level of employment is combined with separate contributions from capital and technology in a production function. Examples of this approach include OECD (2015), which draws on Johansson *et al* (2013); International Monetary Fund (2015), which draws on De Masi (1997) and references therein; and de Brouwer (1998, Section 2.5).

In principle, the Phillips curve/production function approach makes it possible to allow for other influences on potential GDP growth, such as expected demographic changes. In practice, typical production function estimates model many components as univariate processes, such as HP trends. So little extra information is gained and information on other variables and covariances between the components is lost. This approach makes potential GDP vary with the cycle, particularly near end points. Furthermore, this approach incurs costs of complexity and loss of transparency. We are not aware that these complications have been shown to provide a useful benefit. Indeed, forecasts that employ this approach have not performed well, as we discuss in Section 5.

Structural macroeconomic models provide several measures of potential output. Perhaps the most prominent measure, typically associated with dynamic stochastic general equilibrium models, is defined as GDP that is consistent with flexible wages and prices (e.g. Basu and Fernald 2009; Vetlov *et al* 2011). One limitation of such measures is that they are highly model-dependent. Changes to the specification of the economy’s underlying structure can yield quite different

estimates. Another limitation is that it is not clear that these measures satisfy other purposes of potential output.

The Kalman filter is commonly used to distinguish between trends, cycles, noise and other influences. Accordingly, it is an increasingly popular tool for estimating the growth rate of potential output. See, for example, Fleischman and Roberts (2011) and references therein. However, this research has typically involved large systems that draw signals from many variables. We also use the Kalman filter, but only draw signals from the unemployment rate, which makes our approach simpler to estimate, interpret and evaluate.

To be clear, our measure is not designed to serve every purpose of estimates of potential output. For example, we do not claim that it directly helps forecast inflation or short-term movements in GDP. Nor do we see our measure as a substitute for measures designed for these purposes. Different measures for different purposes can coexist. It would be nice to have one measure that did everything well, but that has yet to be found.

As noted above, potential output has been defined in many different ways. Those accustomed to other measures may find our usage confusing. This difficulty is regretted, but seems unavoidable if one is to be clear about what the term potential output means. We view our definition as reasonable and informative: it is simple, useful, standard in the literature on Okun's law, consistent with Okun's original usage and consistent with the main way that central banks use the concept. That said, little of substance depends on definitions. Were other measures to be in active use, more precise terminology would probably be necessary.

2.2 Okun's Law

Arthur Okun (1962) pointed out a negative relationship between unemployment and output, which became known as 'Okun's law'. The 'levels' form of the relationship can be written as:

$$U_t - U^* = \theta(Y_t - Y^*) \quad (1)$$

where U is the unemployment rate and Y is the logarithm of the level of real GDP. Okun called U^* ‘full employment’ and Y^* ‘potential output’. There are many estimates of relationships like this using Australian data, including Kalisch (1982), Nguyen and Siriwardana (1988), Peters and Petridis (1988) and Ball, Leigh and Loungani (2013). However, as we discuss in Section 6.2, we have not found estimates of U^* and Y^* to have predictive power either within sample or out of sample. Moreover, although measures of U^* and Y^* can be constructed so that Equation (1) fits the data, ‘explanations’ that rely on unobservable variables constructed after the event are not compelling.

Okun also pointed out that a similar relationship can be written in differences.

$$\Delta U_t = \beta(\Delta Y_t - \mu) \quad (2)$$

where ΔU is the change in the unemployment rate, ΔY is growth in log GDP (typically represented in terms of annualised percentage changes) and μ is the rate of GDP growth consistent with stable unemployment. Using quarterly US data from 1947 to 1960, Okun estimated $\mu = 4$ and $\beta = -0.3$.² Under some assumptions Equation (2) can be derived from Equation (1). However, we prefer to think of Equation (2) as a separate and simpler way of modelling the data that avoids some of the problems with Equation (1) noted above.

Variations on Equation (2) have been estimated for many different countries, at different frequencies and for different sample periods. Borland (2011) estimates $\mu = 3.2$ and $\beta = -0.41$ based on four-quarter changes from 1979 to 2011 in Australia. Earlier, similar Australian estimates include Watts and Mitchell (1991), INDECS Economics (1995) and Dixon and Thomson (2000). These estimates are broadly similar to those from Okun’s law in changes for other countries, such as those in Ball, Jalles and Loungani (2015).³ Ball *et al* also report that an equation like (2) describes the relationship between output and unemployment in professional forecasts in many countries.

2 We have rearranged coefficients from his ‘first differences’ specification to be in comparable units to those elsewhere in our paper. Okun’s preferred estimates, which are similar, averaged estimates from a range of specifications.

3 Ball *et al*’s estimates for $-\beta$ are, in ascending order, Germany (0.08), Japan (0.11), Italy (0.16), New Zealand (0.24), France (0.27), Canada (0.43), the United States (0.5) and Australia (0.5).

One of the main objectives of our paper is to re-examine Equation (2). We estimate over a long sample period (1960–2015) and pay close attention to its stability and specification. We corroborate previous researchers' finding that simple ordinary least squares (OLS) models of Okun's law are not stable. Watts and Mitchell (1991) attribute instability in the Okun's law relationship in Australia to a decline in trend GDP growth over time. Fortunately, we can model this instability, as we show in the next section.

More importantly, our paper differs from previous estimates of Okun's law in drawing out its implications. In particular, Okun's law provides (a) time-varying estimates of potential output and (b) forecasts of the unemployment rate that compare favourably with alternative forecasts.

3. Models

We present two models of Okun's law in changes.

3.1 Constant Coefficients Model

A simple benchmark with constant coefficients is estimated by OLS.

$$\Delta U_t = \alpha \Delta U_{t-1} + \beta (\Delta_2 Y_t - \mu) + \gamma \Delta_2 rulc_{t-2} + \varepsilon_t. \quad (3)$$

ΔU is the one-quarter change in the unemployment rate, $\Delta_2 Y$ is two-quarter annualised growth in log GDP and μ is the annualised growth rate of potential GDP. So $\Delta_2 Y - \mu$ is the annualised two-quarter change in the output gap. $\Delta_2 rulc$ is the annualised two-quarter change in log real unit labour costs.

A few features of our specification are not conventional. Our lag structure is chosen so as to help explain the data. We do not include the change in the NAIRU, or the levels of either the NAIRU or potential output, for reasons we explain in Section 6. And perhaps most unusual is our inclusion of real unit labour costs. We interpret this term as reflecting factor substitution. When the relative cost of labour increases, firms seek to increase labour productivity and lower their demand for labour for a given level of output, raising unemployment. This effect is frequently mentioned in discussions of the determinants of Australian unemployment, and we

are somewhat surprised that it does not seem to have been included in Okun's law specifications before. This is particularly so as the econometric evidence of its importance is strong, with a t -statistic of 4 in our preferred model. In principle, increases in real unit labour costs boost labour productivity and hence could increase the level of potential output. However, this interaction is small (especially if a reduction in capital productivity is assumed) and we ignore it below.

3.2 Time-varying Coefficients Model

Our preferred model allows potential GDP growth μ , and the coefficient on the lagged dependent variable α , to vary over time (hence we write them with a subscript t). We assume both of these parameters follow random walks, and we estimate them with the Kalman filter.

We specify a measurement equation:

$$\Delta U_t = \alpha_t \Delta U_{t-1} + \beta (\Delta_2 Y_t - \mu_t) + \gamma \Delta_2 rulc_{t-2} + \varepsilon_t \quad (4)$$

and state equations

$$\mu_t = \mu_{t-1} + v_t \quad (5)$$

$$\alpha_t = \alpha_{t-1} + w_t. \quad (6)$$

To initialise the Kalman filter, we assume potential GDP growth in 1960:Q2 was 4 per cent, with a variance of 5 percentage points, equal to the average rate and variance of annual GDP growth in the 1950s. We assume a zero starting value for the coefficient on our lagged dependent variable, with a variance of 1. Other priors produce different results in the first few years of estimation but converge by the late 1960s.

4. Estimation

We estimate both models over the period 1960:Q3 to 2015:Q1 using EViews default settings, except where noted.

4.1 Data

We use three series for estimation: real GDP, the unemployment rate and real unit labour costs, each of which is quarterly and seasonally adjusted. These data and estimation code are publicly available at <http://www.rba.gov.au/publications/rdp/2015/2015-14-data.html>. GDP data are sourced from the national accounts. Unemployment rate estimates are from the ABS Labour Force Survey, back to 1966:Q3. For earlier data, we use estimates for unemployment (number of persons) and the labour force from the ABS ‘Modellers’ Database, Jun Qtr 2014’ (ABS Cat No 1364.0.15.003).

Real unit labour costs are obtained from internal RBA databases. They are defined as the sum of compensation of employees (COE) and payroll taxes paid, less government employment subsidies received, divided by the product of real non-farm GDP and the implicit price deflator for gross national expenditure. The ratios of employment subsidies and payroll taxes to COE are assumed not to change until 1978:Q3 and 1972:Q3 respectively, when quarterly estimates begin.

4.2 Results

Table 1 presents estimated coefficients. We also report the long-run response of unemployment to changes in output, known as ‘Okun’s coefficient’, which equals:⁴

$$C = (4 * \beta) / (1 - \alpha). \quad (7)$$

Multiplication is by 4 because changes in output are annualised. Apart from the time-varying coefficients, the estimates for the two models are very similar. In particular, they have essentially the same coefficients on the change in the output gap and the change in real unit labour costs. As can be seen by comparing the standard error of the OLS equation (0.246) with its Kalman filter analogue, the standard deviation of measurement equation prediction errors (0.233), the model with time-varying coefficients fits the data somewhat better.

⁴ Our estimate holds real unit labour costs constant whereas other researchers often allow other variables to change. In practice, this difference is unimportant.

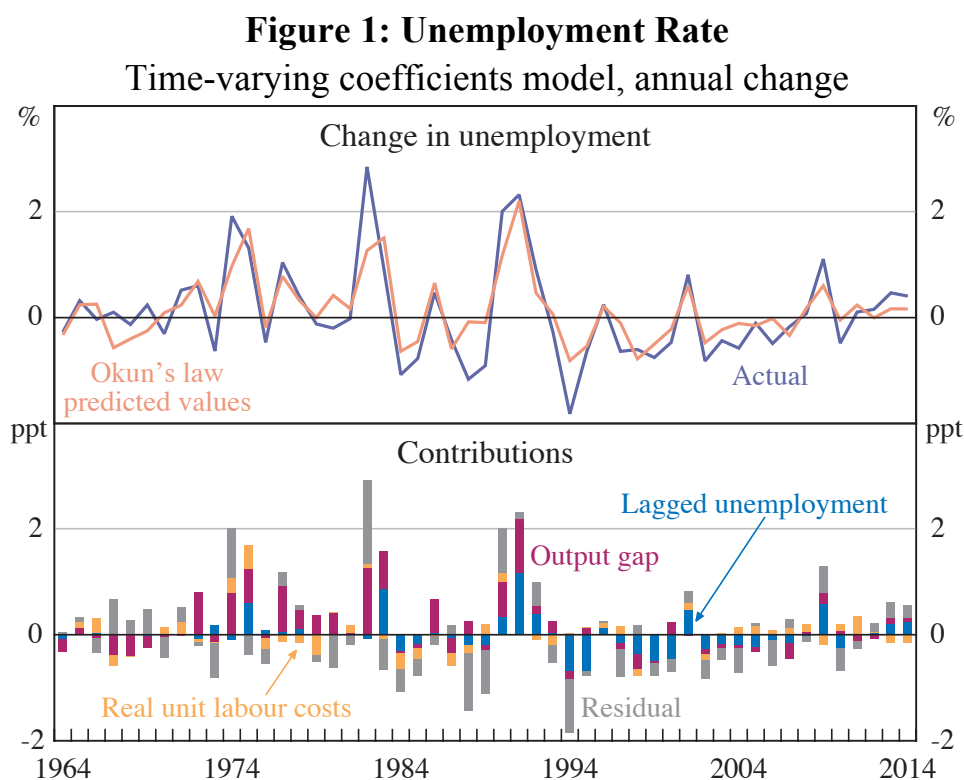
Table 1: Estimation Results
1960:Q3–2015:Q1

Parameter	Description	Constant coefficients	Time-varying coefficients
$\alpha, \alpha_{2015:Q1}$	Lagged dependent variable	0.310 (0.064)	0.443 ^(a) (0.251)
β	Change in output gap	-0.047 (0.008)	-0.049 (0.006)
$\mu, \mu_{2015:Q1}$	Potential output growth	3.69 (0.343)	2.94 ^(a) (0.781)
γ	Change in real unit labour costs	0.016 (0.005)	0.019 (0.004)
σ_ε	Standard deviation of measurement equation error		0.233
σ_v	Standard deviation of potential output equation error		0.124
σ_w	Standard deviation of inertia equation error		0.044
$(4 * \beta)/(1 - \alpha)$	Long-run effect of output on unemployment ('Okun's coefficient')	-0.27	-0.35 ^(a)
Equation standard error		0.246	
Adjusted R^2		0.43	
Estimation method		OLS, with White (1980) standard errors	Kalman filter

Notes: Standard errors reported in parentheses
(a) Time-varying parameter as at 2015:Q1

The upper panel of Figure 1 compares fitted values from the Kalman filter (on an annual change basis, to reduce clutter) with actual changes in the unemployment rate. The model seems to explain most of the variation in the data. Fitted values from the constant coefficients model are not shown but look very similar.

The lower panel of Figure 1 shows estimated contributions to changes in the unemployment rate. Most of the model's explanatory power comes from changes in the output gap. We discuss this effect, and its stability over time, in more detail in Section 4.3.



Sources: ABS; Authors' calculations

Real unit labour costs have had a clear and important influence on the Australian unemployment rate, with a coefficient that is highly statistically significant in both models. Long-run effects change over time due to interaction with the coefficient on the lagged dependent variable. As a guide, the largest response followed the 12 per cent (0.112 log points) increase in real unit labour costs over the five quarters to 1975:Q1. We estimate that this increase, if sustained, would have raised the unemployment rate by 1.1 percentage points in addition to its effects on the output gap. Since the late 1980s, real unit labour costs have been more stable and hence less important for explaining changes in unemployment.

Inclusion of real unit labour costs does not greatly affect other results. Without it, our residuals would be even more skewed – because some of the largest increases in unemployment, in 1975 and 1982, are partly explained by large increases in real unit labour costs. We discuss skewness in Section 6.1. Our estimate of the variance of changes in potential output growth would be slightly smaller, giving rise to a slightly flatter decline in potential over the past two decades, with a final (2015:Q1) estimate of 3.1 per cent. And, as mentioned in Section 5, recent forecasts would be slightly more accurate.

4.3 Time-varying Parameters

OLS estimates of Okun's law are unstable over time. This can be seen, for example, in Quandt-Andrews tests of parameter instability at unknown breakpoints, where the hypothesis of constant parameters is rejected with a p -value of 2 per cent.⁵ Tests of stability for individual coefficients, and rolling regressions, indicate that this instability is attributable to the intercept and to the coefficient on the lagged dependent variable. Accordingly, we model these as time-varying parameters.

As shown in Figure 2, we estimate that potential GDP growth fell from around 5 per cent in the 1970s to 2.9 per cent at 2015:Q1.⁶ Here, and henceforth in the text, we translate our regression estimates into approximate annualised percentage changes by multiplying coefficients on quarterly log differences by 400. Figure 2 also shows a one standard deviation range about the point estimates. Assuming normality, a 68 per cent confidence interval for the 2015:Q1 estimate spans 2.2 to 3.7 per cent.


Recent estimates of potential GDP growth have increased over recent years. This reflects an upward drift in the unemployment rate, averaging one quarter of a percentage point a year since 2011. Given that GDP growth was averaging about 2½ per cent a year, near its estimated potential rate, this surprised our model. The Kalman filter takes these surprises as a signal that its estimate of the output growth required to stabilise unemployment was too low, and revises it higher.

The sharpest changes in potential output growth tend to occur when unemployment rises suddenly, as in 1991 and 2009. This reflects the skewness of changes in unemployment. As we discuss in Section 6.1, our model is able to account for much, though not all, of this skewness.

5 Based on Exp-F tests with a 15 per cent trim and Newey-West covariance matrix; other tests give essentially the same result.

6 We estimate potential GDP growth was close to 5 per cent in the 1960s also. However, this reflects a significant weight on the initial prior value required for Kalman filter estimation, so we do not emphasise this as a finding.

Figure 2: Potential GDP Growth
Implied from Okun's law, one-sided estimates



An interesting feature of our estimates is that revisions to potential are positively correlated with surprises to unemployment, *given output*. These surprises usually involve increases in the growth rates of labour force participation or productivity, some of which the model estimates to be permanent. This countercyclical behaviour of potential is offset by procyclical revisions to potential when output is surprisingly strong, given unemployment. When GDP and unemployment move predictably together over the cycle, as will be true on average, potential will be unrevised.

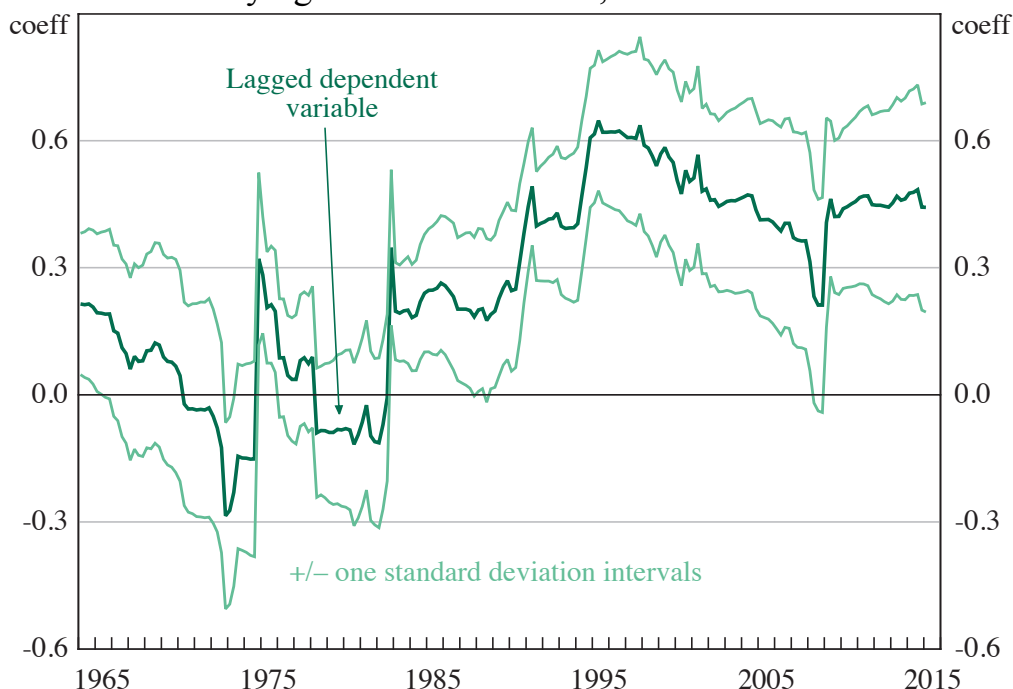
Our estimate of potential output growth in the 1970s of around 5 per cent is relatively high. In contrast, the 'production function' approach of the IMF and OECD tends to produce estimates for this period that are quite low. This largely reflects a conceptual difference. Our measure refers to expected (or *ex ante*) growth rates, the measure that is relevant for forecasting, whereas production function estimates refer to realised (or *ex post*) changes. The latter will include shifts in the *level* of potential output. An estimated increase in the NAIRU lowered *ex post* growth rates in the 1970s. But level shifts that are not expected to continue do not affect expected growth rates.

An advantage of production function approaches is that they can decompose changes in potential output growth – an issue on which our model is silent. Unfortunately, we are not aware of a recently published decomposition. However, our understanding of this approach is that it would also point to a substantial slowdown in potential output growth several decades ago, attributable to roughly similar reductions in the contributions from labour, capital and productivity.

Our other time-varying parameter α_t , shown in Figure 3, rarely attracts attention outside the world of unemployment forecasters. It captures the increasing persistence of movements in the unemployment rate. In the early part of our sample, movements in the unemployment rate were often transitory, but in recent decades they appear to have become quite persistent. We suspect that a large part of this change reflects measurement improvements. Advances in technique and sample size considerably reduce the noise evident in the early part of the sample. In particular, the introduction of rotation sampling in 1978 imposed a moving average term on the level of the unemployment rate.⁷

Figure 3: Coefficient on Lagged Dependent Variable


Time-varying coefficients model, one-sided estimates



⁷ Most of the increase in this coefficient occurred in the 1980s and pre-dates a substantial trend decline in the unemployment rate. Hence asymmetry in unemployment rate movements, discussed in the following section, is unlikely to be an explanation.

In contrast, the short-run response of changes in unemployment to changes in output has been surprisingly stable over time. To illustrate, the solid lines in Figure 4 show least squares estimates from three separate sub-samples of similar length. A percentage point decrease in the output gap is estimated to raise the unemployment rate within two quarters by about a quarter of a percentage point in each sub-sample.

Figure 4: Response of Unemployment to the Output Gap
1 percentage point decrease in the output gap



This stability has implications for views on structural change in the labour market. It is common to hear speculation that behaviour has changed over time: for example, because labour hoarding is more or less common or because workers transition between categories differently. The essentially unchanged short-run response of unemployment to GDP over several generations does not refute these hypotheses, but it does invite scepticism about their broader macroeconomic importance.

There is more variation in the longer-run effects, reflecting changes in estimates of persistence. For comparison, the dashed line in Figure 4 shows estimated responses from our time-varying coefficients model using parameter values at the end of our estimation sample. This model estimates a similar initial impact to the three least

squares estimates. Reflecting the high level of persistence at the end of the sample, the long-run effect (the ‘Okun coefficient’) is relatively large, at -0.36 percentage points.

The stability of the coefficient on real unit labour costs is difficult to assess. Most of the variation in this variable occurred in the 1970s and 1980s, with more recent data being uninformative.

5. Forecasting

An important problem confronting macroeconomic policymakers – and a motivation for this paper – is that unemployment forecasts have been biased. For example, since 2000 the unemployment rate has, on average, been 0.22 percentage points lower than the RBA staff forecast it to be one year earlier. This is difficult to attribute to chance; the hypothesis that the forecasts are unbiased can be rejected with a p -value of 4 per cent.⁸ This bias has been a recurring feature of the RBA forecasts, also being statistically and economically significant in data from the 1990s. Tulip and Wallace (2012) discuss this further.

The bias in the unemployment forecast is not due to low GDP forecasts. Since 2000, the RBA’s forecasts of real GDP growth have actually been too high (the wrong sign to explain unemployment bias), whereas over the 1993–2011 sample reported by Tulip and Wallace (2012, Table 4), they have been unbiased.

One possible explanation for this discrepancy is that forecasters are doing a reasonable job in forecasting actual output growth, but when they translate that into unemployment, they use an out-of-date estimate of potential output growth. They may not realise that larger declines in unemployment are likely to accompany a given growth rate of GDP than in the past.

To assess this, we examine the pseudo real-time forecasting performance of our constant coefficients and time-varying coefficients models and compare this with the actual real-time forecasts of the unemployment rate by the RBA. For each quarter t , we estimate our models up to quarter $t-1$, from which we take

⁸ Hypothesis tests mentioned in this section use Newey-West standard errors with EViews default settings.

coefficients and the latest estimates of state variables.⁹ We combine these with real-time RBA forecasts of GDP growth for quarters t , $t + 1$, ..., $t + h$ to project right-hand side variables, and then forecasts of the unemployment rate through to $t + h$. An important assumption underpinning this exercise is that the RBA forecasts of output growth are independent of changes in our estimate of potential GDP growth. We then repeat for quarter $t + 1$ and so on, giving us sequences of forecasts at each horizon.


We attempt to estimate our models on the data available to forecasters at the time. Real-time data and forecasts for the unemployment rate and GDP growth are described in Tulip and Wallace (2012), which we have updated. We do not have access to real-time vintages of the change in real unit labour costs, so we use the 2015 vintage of data for quarters before the forecast and assume zero change in real unit labour costs for quarters t , $t + 1$, ..., $t + h$. The use of revised data may give rise to concerns that we may overstate our model's forecasting performance at short horizons. However, including real unit labour costs actually reduces the model's forecasting performance over the sample we use (the explanatory power of this variable comes from the 1970s and 1980s).

We focus on a subset of the available data that we think is most informative. Our evaluation period begins in 2000, given the availability of forecast information. As discussed below, some of our conclusions are sensitive to this choice. We show results for horizons out to six quarters ahead. We have a smaller sample of longer-term forecasts, which is heavily influenced by the global financial crisis, and we suspect is unrepresentative. Moreover, it is less plausible to take longer-term output growth forecasts as given when varying potential output assumptions.

Forecast comparisons are shown in the figures below. Figure 5 shows root mean squared errors, a standard measure of forecast accuracy. By this metric, all three forecasts perform similarly. Differences are not economically or statistically significant (judging by Diebold-Mariano tests, not shown).

⁹ Since 2008, we estimate our models up to quarter $t - 2$ to allow for comparisons with RBA forecasts, which are produced prior to the release of quarter $t - 1$ national accounts data.

Figure 5: Unemployment Rate Forecast Errors
 Root mean squared errors, 2000:Q1–2015:Q1



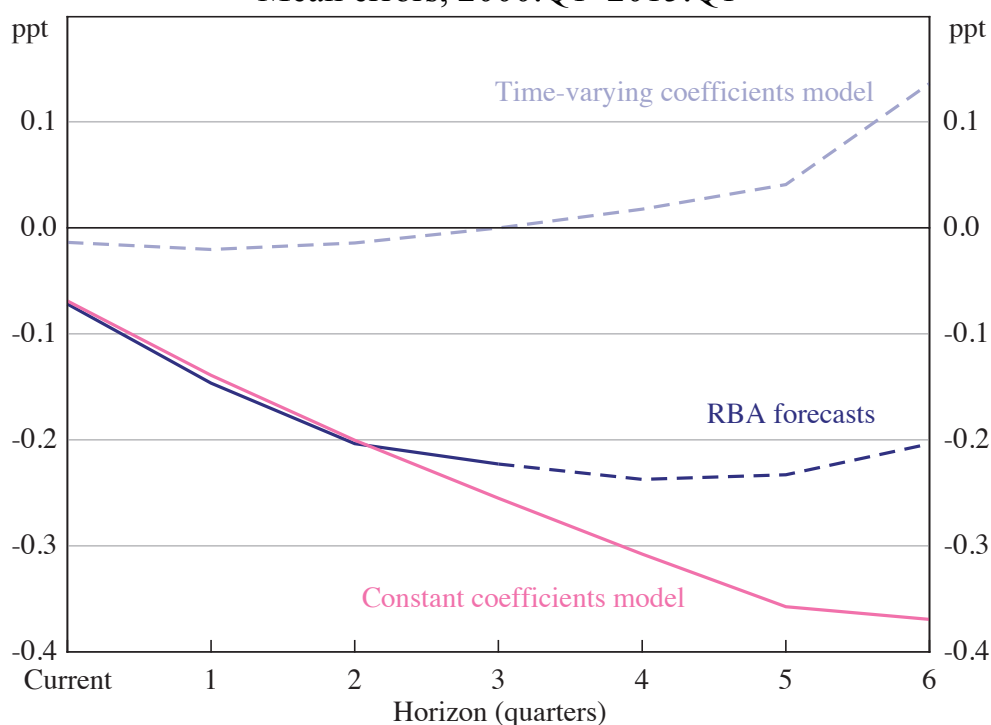
Sources: Authors' calculations; RBA

Figure 6 shows mean errors, a measure of bias. Solid lines are significantly different from zero at a 10 per cent level. The precision (though not the magnitude) of estimates tends to decline as the horizon increases, as we have fewer independent observations. Whereas the RBA and constant coefficients forecasts are noticeably biased, the bias in the time-varying coefficients forecasts is insignificant, in either statistical or economic terms.

The favourable performance of the time-varying coefficients forecast is surprising and important. The relationship between output and unemployment is often described in complicated structural terms (Layard, Nickell and Jackman 1991; Debelle and Vickery 1998). But Figures 5 and 6 suggest that all that is needed to forecast the unemployment rate well is a simple reduced-form equation with very few variables.¹⁰

¹⁰ This result was suggested to us by Alex Cooper.

Figure 6: Unemployment Rate Forecast Bias
Mean errors, 2000:Q1–2015:Q1



Notes: Actual unemployment minus forecast unemployment; dashed lines represent insignificance at the 10 per cent level

Sources: Authors' calculations; RBA

The ranking of forecasts in Figures 5 and 6 would be different if we used earlier data. RBA forecasts outperformed the time-varying coefficients model in the 1990s. Given that structural change is a concern, and that forecast procedures and information sets have changed, we view recent comparisons as more relevant. However, ranking alternative approaches is not our objective. The important point is that a relatively simple model forecasts about as well as more complicated alternatives. This result is not especially sensitive to the sample period.

As noted above, one of our hopes in undertaking this project was that the Kalman filter would remove the bias in the RBA's unemployment forecasts. Although the evidence from Figure 6 seems to show this, evidence from earlier samples is less encouraging. The Kalman filter forecasts from the 1990s were upwardly biased – indeed, by more than the RBA forecasts. We interpret this as a 'learning' effect. As can be seen in Figure 2, the Kalman filter estimates that potential output growth declined substantially over the 1990s. But the model did not have this information before the event and so forecast that the moderate GDP growth of this period

would be accompanied by little change in the unemployment rate. In this respect, the Kalman filter suffers from the same lack of hindsight as other forecasters: an unexpected structural break will be followed by persistent forecast errors.

However, having made overpredictions, the filter adjusts down its estimate of potential growth in response. That is, the Kalman filter ‘learns’ from its mistakes, with subsequent forecasts being unbiased. As shown in Figure 6, neither OLS forecasts, nor the RBA staff, responded in a similar manner. The responsiveness of the Kalman filter to its own forecast errors makes it relatively robust to structural breaks, in contrast to RBA staff procedures or OLS.


An alternative approach to the instability in Okun’s law might be to estimate our constant coefficients model over a short sample period – say, over the past 10 or 20 years. Such a model would have forecast the unemployment rate over the past decade about as well as our Kalman filter. If all one was interested in was forecasting unemployment, that approach would be simple and easy. Whether it would be reliable is less certain. One difficulty with this approach is that the growth of potential output has changed in the past and can be expected to change again in the future. As discussed above, least squares estimates are not robust to structural change. A second difficulty is that the choice of the sample period, and hence the responsiveness of parameter estimates to new data, is arbitrary. The n -period average chosen for one dataset may work poorly elsewhere. In contrast, the Kalman filter ‘gain’ is estimated so as to best describe the data. Third, short-sample least squares estimates allow all parameters to change substantially in response to unusual observations. In contrast, our Kalman filter model constrains parameters that have been stable over long periods – such as the short-run response of unemployment to output – to be insensitive to blips in the data.

6. Specification Issues

6.1 Asymmetry


A striking feature of the unemployment rate is that the distribution of changes is asymmetric. As shown in Figure 7, the unemployment rate tends to go up quickly, and down slowly. (The gradual increase over 2013 and 2014 is an exception.) Figure 8 shows the distribution of four-quarter changes. The most common outcome is a small decline. There are several very large increases of over 3 percentage points, but no decreases of more than 2 percentage points. This skewness is also evident in changes in other frequencies and in other countries.

Figure 7: Unemployment Rate



Source: ABS

Figure 8: Distribution of Changes in Unemployment Rate
Four-quarter change



Sources: ABS; Authors' calculations

Table 2 shows a standard measure of asymmetry, the skewness statistic. We exclude the first few years from the calculations, to allow the diffuse priors about our state variables to tighten. For this sample, the skewness of quarterly changes in the unemployment rate is 1.60. For a sample of the same size, a normally distributed variable would exhibit skewness this large less than 0.01 per cent of the time.¹¹

¹¹ All p -values in this section are generated by Monte Carlo. We take 100 000 draws of 202 observations of a normal random variable to create the distribution of skewness statistics.

Table 2: Skewness Statistics

1964:Q4–2015:Q1

	Skewness statistic	p -value (%); $H_0 = 0$
Unemployment rate, quarterly change	1.60	< 0.01
Output gap, two-quarter change (constant coefficients model)	−0.32	3.29
Output gap, two-quarter change (time-varying coefficients model)	−0.84	< 0.01
Okun’s law model with time-varying coefficients		
Predicted change in unemployment rate	1.53	< 0.01
Measurement equation residuals	0.94	< 0.01

The asymmetry of changes in the unemployment rate is largely, though not entirely, explained by our model. Changes in the output gap are also skewed, especially when potential output growth varies over time (skewness = -0.8 , $p < 0.01$ per cent). Moreover, the interaction of changes in the output gap with the other regressors results in predicted values of our model that are similarly skewed as changes in the unemployment rate. There remains some unexplained skewness in the residuals (skewness = 0.9 , $p < 0.01$ per cent), but it is much smaller than the skewness in the fitted values (skewness = 1.5 , $p < 0.01$ per cent). That is, we explain most, though not all, of the skewness in the data.

The unemployment rate rises quickly in recessions and falls slowly in expansions. Although our model explains most of this asymmetry, an alternative explanation is that the response of unemployment to changes in the output gap is nonlinear. To assess this, we interact the change in the output gap with a ‘boom’ dummy variable that equals one when the output gap is increasing and zero otherwise. This variable is not significantly different to zero ($p = 39$ per cent). Ball *et al* (2013) conduct a similar test on US data and find similar results.

We consider other specifications. Some of these variations were marginally significant but difficult to interpret (for example, implying discontinuous relationships). Others were insignificant. Having experimented with a range of functional forms, we have not found clear evidence that the relationship between output growth and changes in unemployment is nonlinear.

6.2 Changes in the NAIRU

Okun's law equations like Equations (3) and (4) are often estimated with the dependent variable being the change in the unemployment *gap* (the deviation of the unemployment rate from the NAIRU) rather than the change in the unemployment *rate*. From a forecasting perspective, these variables are typically the same, as the NAIRU is conventionally modelled as a random walk. But for explaining history, we find that this specification fits the data poorly. We update the domestic demand deflator version of the NAIRU estimated by Ballantyne, De Voss and Jacobs (2014) and include the contemporaneous quarterly change in the one-sided NAIRU in Equation (4). The coefficient is incorrectly signed and is not significantly different from zero ($p = 6$ per cent). In this specification, the expected growth rate of potential output is about half a percentage point higher in the 1970s and about half a percentage point lower in the 2000s than the estimates shown in Figure 2, though the 2015:Q1 estimate is unchanged at 2.9 per cent. The hypothesis that the dependent variable is the change in the gap, which implies constraining the coefficient on the change in the NAIRU to equal one, is very strongly rejected ($p < 0.1$ per cent). The hypothesis that the coefficient equals $1 - \alpha_t$ (allowing for interaction with the lagged dependent variable) also has $p < 0.1$ per cent. Ballantyne *et al*'s weighted median version of the NAIRU provided an even poorer fit to the data.

6.3 Levels of Potential Output and Equilibrium Unemployment

In contrast to most previous research on potential output and Okun's law, we do not present estimates of levels of potential output, the output gap or the equilibrium rate of unemployment. We have not been able to construct estimates that are useful for explaining unemployment. For example, the lagged level of the unemployment rate and a constant (allowing for reversion to the mean or to the NAIRU) are jointly insignificant with a p -value of 17 per cent when included in a least squares version of our equation. (We constrain time-varying coefficients to equal their Kalman filter estimates, so as to avoid collinearity between the constant and potential output growth.) Similarly, the lagged output gap, constructed as the deviation of log GDP from its one-sided HP trend, has a coefficient that is insignificantly different from zero ($p = 25$ per cent). More sophisticated specifications, including cyclically adjusted estimates and one-sided estimates of

the NAIRU, do not do much better. This is not primarily because our time-varying coefficients soak up the effect of the levels variables. When the lagged deviations of log GDP and the unemployment rate from their one-sided HP trends are included in our constant coefficients specification (as shown in Table 1) they are individually and jointly insignificant ($p = 43$ per cent).

To be clear, we are not saying that the relationship between unemployment and output does not hold in levels. Rather, variables in levels do not seem to provide information about changes in unemployment beyond the information provided by variables in changes.

The insignificance of level terms may seem surprising given the stationarity (absence of trend) of the unemployment rate. In the past, when the unemployment rate was far from its average, it returned toward normal. Our estimates suggest that this mean reversion arises through output growing faster or slower than normal. Once one allows for that, extra reversion to the mean is unimportant.

Exceptions to our findings above are specifications that include two-sided trends (whether HP, kinked linear or Kalman) of both output and the unemployment rate. Like Ball *et al* (2013) we find these terms to be highly statistically significant. But this correlation seems to reflect reverse causation. In particular, disturbances to the unemployment rate result, by construction, in positively correlated changes in the lagged HP trend.

To illustrate, we construct artificial data resembling our unemployment rate and Ball *et al*'s (2013) NAIRU, in which the null hypothesis, that the NAIRU has no effect on the unemployment rate, is true by construction. Specifically, we simulate 50 000 unemployment rate series, each of 219 observations, taking the right-hand side variables and coefficients of the time-varying model as given and random draws of disturbances from a normal distribution, where $\hat{\varepsilon} \sim N(0, 0.233^2)$. We then construct a HP trend of the cumulative sum of each series using the same procedure as Ball *et al* (2013) (with a smoothing coefficient $\lambda = 1\,600$). Although the artificial unemployment data are constructed independently of the artificial NAIRU series, when included in our Okun's law model, the resulting lagged level of the unemployment gap usually appears to be highly significant, with an average t -statistic of -4.2 . At a 5 per cent level of significance, a t -test on the lagged

deviation of the unemployment rate from its HP trend incorrectly rejects the null hypothesis of a zero effect 99.8 per cent of the time.¹² We conclude that the finding that the unemployment rate moves towards a two-sided trend of itself is spurious. That is, Ball *et al* (2013) overstate the evidence that Okun's law holds in levels. This conclusion applies in essentially the same way to kinked linear trends and two-sided Kalman trends. We should add that, in contrast, our other results strongly corroborate Ball *et al*'s findings relating to Okun's law in changes.

7. Directions for Future Research

Our conclusions are stated in the introduction and we do not repeat them. But two possible extensions are worth noting.

Okun's law is a reduced form that does not distinguish between shocks to output arising from demand and those arising from supply. It is possible that the correlations arising from different underlying shocks would differ. The stability of the short-term responses shown in Figure 4 suggests either that responses to different kinds of shocks are similar or that the constellation of shocks has not changed much over our sample. Nevertheless, a deeper understanding of the structure would be interesting. In particular, our model ignores information about other labour market variables, such as productivity, the participation rate and so on. One might hope that taking this information into account would improve estimates of potential output and forecasts of unemployment. The forecast comparisons in Section 5 suggest that this extra information has been difficult to usefully apply. However, that may reflect the simplifications that are often used in 'structural' estimates. We suspect multivariate estimation with time-varying parameters would improve performance.

Estimation of our models on annual data provided estimates that forecast surprisingly well out of sample. This partly involved estimates of potential output growth in the 1990s that were relatively low, with unbiased forecasts. We suspect that longer lag structures may have helped in this, though we were unable to capture those with quarterly data. Again further exploration here might be fruitful.

¹² These results are not materially different if we draw disturbances from the empirical distribution of ε , vary the standard deviation of the residuals or use an alternative value of λ .

References

Ball L, JT Jalles and P Loungani (2015), ‘Do Forecasters Believe in Okun’s Law? An Assessment of Unemployment and Output Forecasts’, *International Journal of Forecasting*, 31(1), pp 176–184.

Ball L, D Leigh and P Loungani (2013), ‘Okun’s Law: Fit at 50?’, IMF Working Paper No WP/13/10.

Ballantyne A, D De Voss and D Jacobs (2014), ‘Unemployment and Spare Capacity in the Labour Market’, RBA *Bulletin*, September, pp 7–20.

Basu S and JG Fernald (2009), ‘What Do We Know (and Not Know) about Potential Output?’, Federal Reserve Bank of St. Louis *Review*, 91(4), pp 187–213.

Borland J (2011), ‘The Australian Labour Market in the 2000s: The Quiet Decade’, in H Gerard and J Kearns (eds), *The Australian Economy in the 2000s*, Proceedings of a Conference, Reserve Bank of Australia, Sydney, pp 165–218.

Debelle G and J Vickery (1998), ‘The Macroeconomics of Australian Unemployment’, in G Debelle and J Borland (eds), *Unemployment and the Australian Labour Market*, Proceedings of a Conference, Reserve Bank of Australia, Sydney, pp 235–265.

de Brouwer G (1998), ‘Estimating Output Gaps’, RBA Research Discussion Paper No 98-09.

De Masi PR (1997), ‘IMF Estimates of Potential Output: Theory and Practice’, IMF Working Paper No WP/97/177.

Dixon R and J Thomson (2000), ‘Okun’s Law and Movements over Time in the Unemployment Rate in Australia’, The University of Melbourne Department of Economics Research Paper No 741.

Fleischman CA and JM Roberts (2011), ‘From Many Series, One Cycle: Improved Estimates of the Business Cycle from a Multivariate Unobserved Components Model’, Board of Governors of the Federal Reserve System Finance and Economics Discussion Series No 2011-46.

INDECS Economics (1995), *State of Play 8: The Australian Economic Policy Debate*, Allen & Unwin Australia, Sydney.

International Monetary Fund (2015), ‘World Economic Outlook Database’, April, accessed June 2015.

Johansson Å, Y Guillemette, F Murtin, D Turner, G Nicoletti, C de la Maisonneuve, P Bagnoli, G Bousquet and F Spinelli (2013), ‘Long-Term Growth Scenarios’, OECD Economics Department Working Paper No 1000.

Kalisch DW (1982), ‘The Output Loss from Unemployment: Okun’s Law Revisited’, Bureau of Labour Market Research Working Paper No 10.

Kiley MT (2013), ‘Output Gaps’, *Journal of Macroeconomics*, 37, pp 1–18.

Layard R, S Nickell and R Jackman (1991), *Unemployment: Macroeconomic Performance and the Labour Market*, Oxford University Press, Oxford.

Nguyen DT and AM Siriwardana (1988), ‘The Relationship between Output Growth and Unemployment: A Re-examination of Okun’s Law in Australia’, *The Australian Economic Review*, 21(1), pp 16–27.

Norman D and A Richards (2010), ‘Modelling Inflation in Australia’, RBA Research Discussion Paper No 2010-03.

OECD (Organisation for Economic Co-operation and Development) (2015), ‘Economic Outlook, Analysis and Forecasts: EO Sources - Notes to Statistical Annex Tables 20-22: Key Supply-Side Data: Annex Table 21 – Potential GDP, Employment and Capital Stock’, OECD.org site, accessed 20 April 2015. Available at <<http://www.oecd.org/eco/outlook/eosources-notestostatisticalannextables20-22keysupply-sidedata.htm>>.

Okun AM (1962), ‘Potential GNP: Its Measurement and Significance’, *1962 Proceedings of the Business and Economic Statistics Section American Statistical Association*, pp 98–104.

Peters RW and A Petridis (1988), ‘Potential Output and Okun’s Law for Australia since 1974 - Implications for Unemployment’, The University of Western Australia Department of Economics Discussion Paper 88.10.

RBA (Reserve Bank of Australia) (2015), ‘Economic Outlook’, *Statement on Monetary Policy*, May, pp 63–69.

Tulip P and S Wallace (2012), ‘Estimates of Uncertainty around the RBA’s Forecasts’, RBA Research Discussion Paper No 2012-07.

Vetlov I, T Hlédik, M Jonsson, H Kucsera and M Pisani (2011), ‘Potential Output in DSGE Models’, ECB Working Paper No 1351.

Watts M and W Mitchell (1991), ‘Alleged Instability of the Okun’s Law Relationship in Australia: An Empirical Analysis’, *Applied Economics*, 23(12), pp 1829–1838.

White H (1980), ‘A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity’, *Econometrica*, 48(4), pp 817–838.



RESERVE BANK OF AUSTRALIA