RESERVE BANK
OF AUSTRALIA

Research Discussion Paper

Estimates of Uncertainty around the RBA's Forecasts

Peter Tulip and Stephanie Wallace

RDP 2012-07

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.

ISSN 1320-7229 (Print)

ISSN 1448-5109 (Online)

Estimates of Uncertainty around the RBA's Forecasts

Peter Tulip and Stephanie Wallace

Research Discussion Paper
2012-07

November 2012

Economic Research Department
Reserve Bank of Australia

We are grateful to many RBA colleagues for guidance on forecasting procedures and databases, in particular Christie Lewis, Patrick D'Arcy and Anna Park. We received helpful comments from Alex Heath, Jonathan Kearns, Christopher Kent, Mariano Kulish, Philip Lowe, David Norman, Adrian Pagan, Mick Plumb, Penny Smith and Tony Richards. 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.

Authors: tulipp and wallaces at domain rba.gov.au

Media Office: rbainfo@rba.gov.au

Abstract

We use past forecast errors to construct confidence intervals and other estimates of uncertainty around the Reserve Bank of Australia's forecasts of key macroeconomic variables. Our estimates suggest that uncertainty about forecasts is high. We find that the RBA's forecasts have substantial explanatory power for the inflation rate but not for GDP growth.

JEL Classification Numbers: E17, E27, E37
Keywords: forecast errors, confidence intervals

Table of Contents

Estimates of Uncertainty around the RBA's Forecasts

Peter Tulip and Stephanie Wallace

Consumers of forecasts should routinely be told something about the size of past errors. Stevens (2004)

1. Introduction

Each quarter, the RBA presents its forecasts for key macroeconomic variables in its *Statement on Monetary Policy* (*SMP*). Readers of the *SMP* may be interested in how informative those forecasts are. For example, how likely is it that inflation will be close to the forecast? How much weight should be placed upon different parts of the forecast? This paper addresses these questions by examining the historical properties of the RBA's forecasts. In particular, we use estimates of past forecast accuracy to construct confidence intervals for forecasts of CPI inflation, underlying inflation, real GDP growth and the unemployment rate.

Estimates of forecast uncertainty may also help to clarify communication. A difficulty policymakers face in discussing the economic outlook is that a single point estimate will have a very high probability of being incorrect. Even though the forecast is meant to be interpreted as the centre of a range of possible outcomes, misunderstandings are common. Perceptions that the central bank was wrong (whether well-founded or not) can undermine credibility and transparency. Our estimates enable the forecast to be considered as a range, which avoids many of these problems. Considering the forecast as a range rather than a point conveys additional information, is hopefully less susceptible to misunderstanding and highlights the considerable uncertainty attached to the outlook.

For these and other reasons, many central banks provide measures of uncertainty with their forecasts. We summarise these presentations in Appendix A. This paper begins by constructing similar measures of uncertainty for Australia, building on the overseas experience. We calculate forecast errors over the past two decades, measure their dispersion and hence construct confidence intervals. For example, 70 per cent of the RBA's forecasts for underlying inflation one year ahead have been within half a percentage point of actual outcomes. If future forecast errors are similar to those in the past, then there is a 70 per cent probability of actual underlying inflation falling within half a percentage point of the current forecast. We construct similar estimates for other variables at other horizons.

We then compare these estimates to some relevant benchmarks. Some of our findings are:

- Uncertainty about the forecasts is high. Confidence intervals span a wide range of outcomes.

- RBA one-year-ahead forecasts have substantial explanatory power for both the level and change in inflation. This contrasts with the experience of some foreign central banks.

- However, deviations of underlying inflation from the target at longer horizons are not predictable. For reasons discussed in Section 4.2, this is a desirable feature of an inflation-targeting framework.

- Underlying inflation is more predictable than headline CPI.

- Forecasting economic activity is more difficult. As has been found for many forecasters overseas, the RBA's forecasts of GDP growth lack explanatory power.

- Forecasts of the unemployment rate outperform a random walk only for a few quarters ahead.

- Relative to private sector forecasts, RBA forecasts of inflation have been marginally more accurate while forecasts of GDP growth have been less accurate. The differences are small.

- Uncertainty about some key variables does not increase with the forecast horizon. We know about as much about economic growth in the current quarter as we do about growth two years ahead.

The paper also discusses various properties of our confidence intervals, alternative measures of forecast uncertainty and some problems with using past errors as a gauge of forecast uncertainty.

Many of our results are qualitatively consistent with previous RBA work. For example, our confidence intervals, appropriately scaled, are similar to the model-based density forecasts of Gerard and Nimark (2008). We discuss the relationship between estimates of uncertainty derived from models and those derived from forecast errors in Section 6.1. More broadly, the RBA has regularly emphasised the difficulties of forecasting and the considerable uncertainty about the economic outlook. See, for example, Stevens (1999, 2004, 2011). In contrast to the approach of some foreign central banks, the RBA has responded to this uncertainty by placing relatively less emphasis on forecasts and more on analysis of current economic developments in its leading publications. In the *SMP*, forecasts of select variables are presented in a table using ranges beyond the near-term horizon to avoid an impression of excessive precision.

Any discussion of past forecast errors will raise questions about whether the forecasts might be improved. This is an important issue of ongoing research, but it is not our focus here. In this paper, we are primarily interested in how readers of the *SMP* should interpret uncertainty about a given forecast.

2. Data

Uncertainty about a forecast can be gauged by the performance of similar forecasts in the past. To this end, it would be desirable to have a record of forecasts resulting from a similar process applied to similar information sets. In practice, processes and datasets have evolved over time. For example, the *SMP* has only recently included numerical forecasts for inflation and GDP. In the earlier part of our sample, we use internal forecasts prepared by the RBA staff, which we assume have similar properties to those published in the *SMP*. At the risk of over-simplifying, we refer to all these forecasts as ‘RBA forecasts’ even though many of them have little official status.

We discuss our data in detail in Appendix B. However, a few important features are worth noting here. Our sample extends from 1993:Q1, when inflation targeting began, through to 2011:Q4. We try to measure actual outcomes with definitions close to those used in the forecast. For GDP growth, this means using near-real-time data, while for underlying inflation we use definitions used at the time of the corresponding forecast.

We show some illustrative data in Figure 1. The dark lines in the top half of each panel represent actual outcomes, measured with near-real-time data, for underlying inflation, the CPI, real GDP growth and the unemployment rate. The light lines in the same panels represent the forecasts of these variables from three quarters earlier. For series published with a one quarter lag, this horizon encompasses the first four quarters of data.[1]

One (of many) interesting features of Figure 1 is the differing explanatory power of forecasts for different variables. As can be seen in the top two panels, many of the variations in underlying and headline inflation were predicted in advance. In contrast, the relationship between the forecasts and GDP growth (third panel on left) is harder to see. We elaborate on this point below.

Figure 1 also presents 3-quarter-ahead forecast errors, measured as outcomes minus forecasts, shown in the bottom half of each panel. As might be hoped, the errors lack obvious patterns. With some exceptions, discussed below, they do not trend, they are centred on zero, they have little persistence (beyond that expected given the 3-quarter-ahead forecast horizon) and their variance does not change noticeably over time. Whereas many other countries experienced extreme adverse forecast errors during the recent global financial crisis, that did not happen in Australia.

Although the data are discussed in detail in Appendix B, one point worth noting here is that the forecasts are conditional on interest rate assumptions. For most of

1 These forecasts are for the current quarter and following three quarters, which we refer to as 'first year' or '3-quarter-ahead' forecasts. Because the timing of the forecast differs from the timing of the most recent data, the common '*h*-period-ahead' terminology is ambiguous. We measure horizons from the date of the forecast, which is convenient in dealing with variables with different release dates and publication frequency. A popular alternative convention is to measure horizons from the date of the latest data, in which case the forecasts and errors in Figure 1 would mainly be described as '4-quarter-ahead'.

Figure 1: Forecasts, Outcomes and Errors

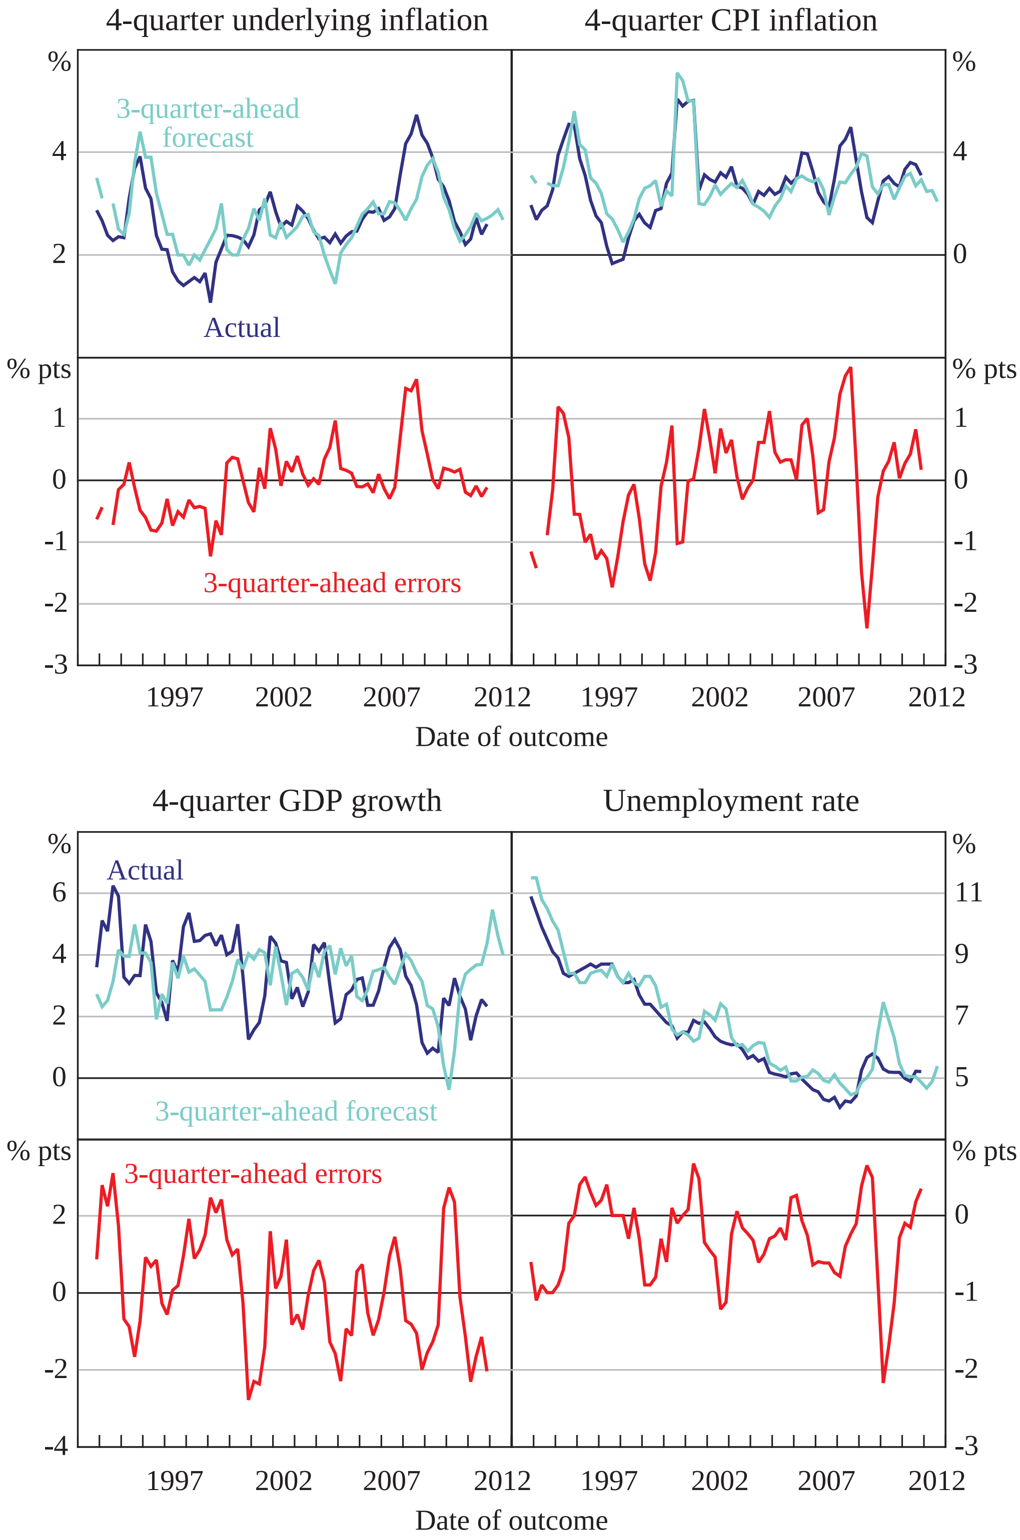

our sample, an unchanged cash rate was assumed. However, when this assumption seemed obviously unrealistic, as in the period following the global financial crisis, forecasts instead assumed a path broadly consistent with market expectations. Unless there is a change in procedures going forward, that assumption does not affect the construction of confidence intervals or other measures of uncertainty. Were this approach to change, it would probably have little effect on measures of forecast accuracy, as we discuss in more detail in Appendix B.

3. Estimates of Forecast Uncertainty

Forecast errors can be summarised in various ways. The root mean squared error (RMSE) is a standard measure of the 'typical' forecast error, with useful statistical properties. Figure 2 shows RMSEs for the four variables at different horizons.

Figure 2: Root Mean Squared Errors
1993:Q1–2011:Q4

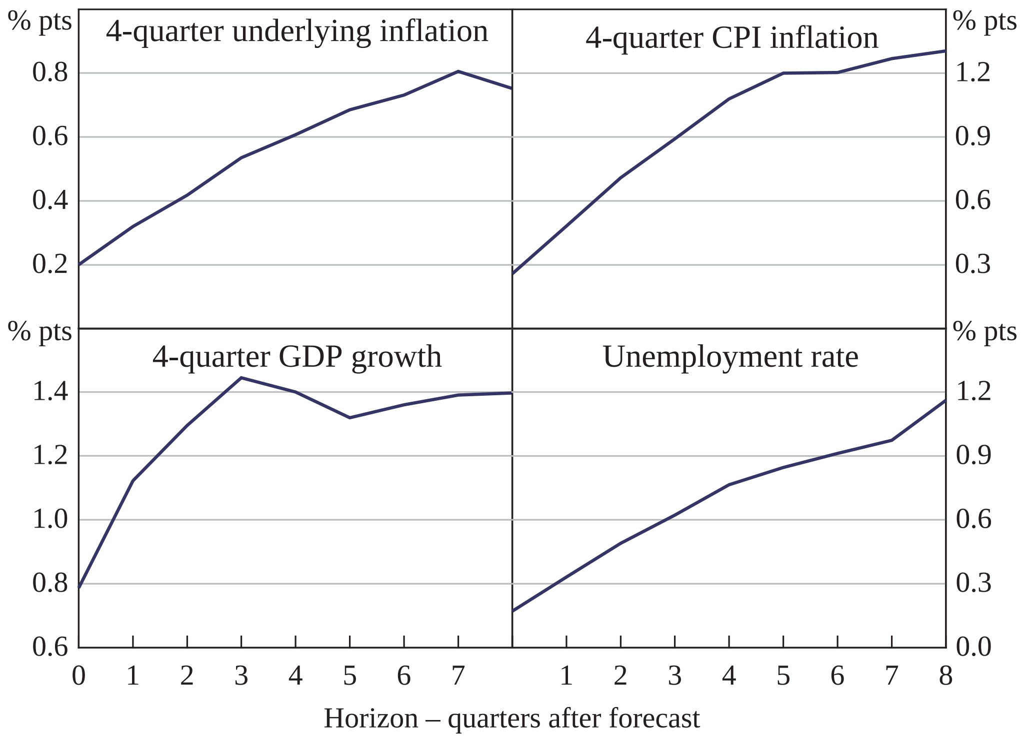

One popular way of presenting forecast uncertainty is through a confidence interval or fan chart, in which a margin of error is added to either side of the central forecast. RMSEs are often used for this purpose. If forecast errors are normally distributed, there is about a two-thirds probability that actual outcomes will fall within one RMSE of the forecast, and about a 95 per cent probability that they fall within two RMSEs of the forecast.

An alternative approach is to use quantiles (such as deciles or percentiles) of the distribution of historical forecast errors. Figure 3 shows 70 per cent and 90 per cent confidence intervals constructed by adding the 70$^{\text{th}}$ and 90$^{\text{th}}$ percentile of absolute forecast errors at each horizon on either side of the February 2012 *SMP* forecast. As a guide to interpretation, consider the 90 per cent interval for underlying inflation in the year-ended 2013:Q4, shown in the top left panel. This indicates that, if forecast errors are the same size as in the past then, at the time of the forecast, there was a 90 per cent probability that underlying inflation in the year ended 2013:Q4 would lie between 1.2 per cent and 4.0 per cent. In Appendix C, we present tables of the 70$^{\text{th}}$, 90$^{\text{th}}$ and other percentiles, so that confidence intervals can be constructed about future forecasts.

Figure 3: SMP Forecasts with 70 Per Cent and 90 Per Cent Confidence Intervals

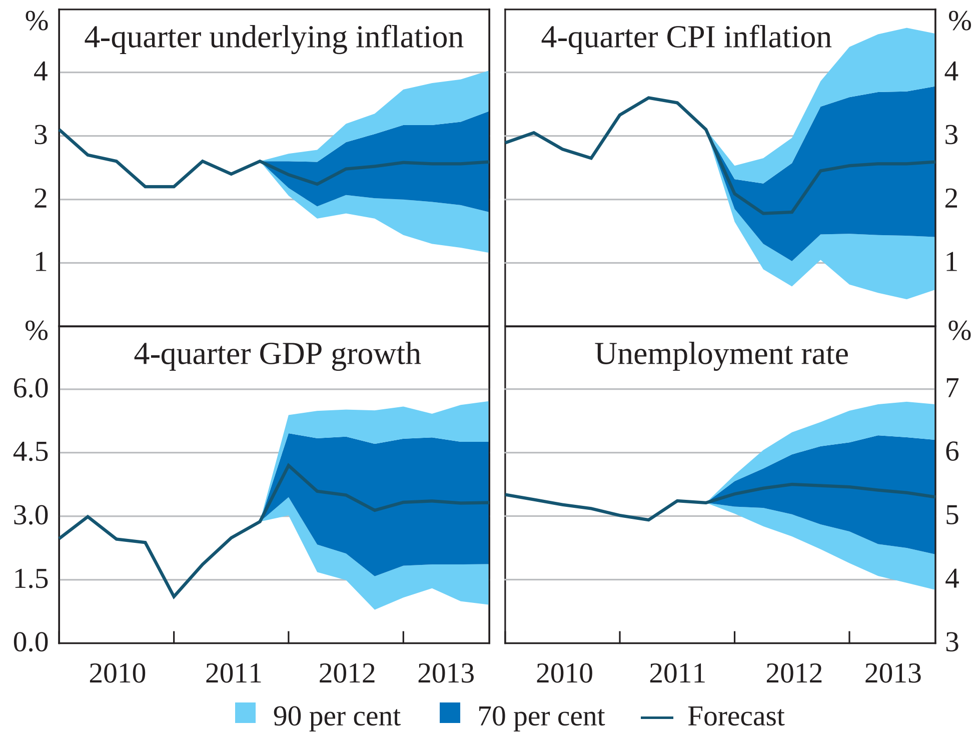

In Section 5.4 we compare RMSEs with empirical quantiles. For our dataset the two approaches give very similar estimates. For most purposes, the choice between them is essentially presentational, and we alternate between the measures in the following sections.

In constructing the confidence intervals in Figure 3 we have made many assumptions. For example, in using absolute errors, we have assumed that the confidence intervals are unbiased and symmetric. However, we have not assumed that the errors are normally distributed. We discuss these and other possible assumptions in Section 5. There are many other ways that forecast uncertainty can be measured and presented. This is illustrated in Appendix A, which describes measures of uncertainty presented by other central banks.

4. How Do These Estimates Compare?

The confidence intervals in Figure 3 strike many observers as wide, particularly for GDP growth. In other words, our estimates of uncertainty are surprisingly high.

Initial impressions presumably reflect comparisons with subjective estimates of uncertainty. Psychological studies find that subjective estimates of uncertainty are regularly too low, often by large margins. People have a systematic bias towards overconfidence.[2] Accordingly, in the absence of objective information, the general public may expect an unrealistically high standard of forecast accuracy.

However, the impression of high uncertainty is also consistent with comparisons to external benchmarks, to which we now turn. The intention in making these comparisons is not to run ‘horse races’ but to help interpret uncertainty about the forecasts.

4.1 Verbal Descriptions of the Forecast

The simplest benchmark is common, qualitative descriptions. The intervals in Figure 3 span outcomes that would be described very differently. For example, the 90 per cent confidence interval for GDP growth in the year ended 2013:Q4 extends from 0.9 per cent to 5.7 per cent. That is, although the central forecast is for growth to be moderate, it could easily turn out to be very strong, or quite weak. Similarly, while little change in the unemployment rate is expected, a large increase or

2 See Part VI, titled ‘Overconfidence’ in Kahneman, Slovic and Tversky (1982) or, for an accessible summary, the Wikipedia (2012) entry ‘Overconfidence Effect’. Contrary to what might be suspected, this bias is not easily overcome. Overconfidence is found among experts and among survey subjects who have been thoroughly warned about it.

decrease is possible. Although the most likely outcome for headline inflation is within the RBA's target range, it could easily be well outside. In comparison, we can be somewhat more confident about underlying inflation, which is likely to remain moderately close to the target range.

Verbal descriptions are simple and meaningful for many readers. But they are also subjective and imprecise. Accordingly, we turn to quantitative benchmarks.

4.2 Variation in the Data

A simple quantitative benchmark for assessing forecast uncertainty is the amount of variation in the data. This benchmark is useful for answering the question: How much does the forecast explain?

A simple measure of data variation is the standard deviation or variance of actual outcomes. This is explicit in some forecast comparisons (Campbell 2007; Vogel 2007; Edge and Gurkaynak 2011) and implicit in many more (as the denominator in the R^2 of popular Mincer-Zarnowitz regressions). However, a conventional or 'centred' standard deviation measures differences from the sample mean. The sample mean is not available at the time of the forecast and does not represent an uninformative alternative. So comparisons with the standard deviation can set an unreasonably high standard; they do not really measure whether the forecast has explanatory power. A more interesting (though very similar) benchmark is the RMSE of an 'uninformative' or 'null' forecast such as an assumption of no change. A forecast that is more accurate than this uninformative alternative can be said to explain some of the variation in the data. We focus on uninformative alternatives that lend themselves to simple interpretations. A forecast that outperforms a random walk can be said to explain changes. A forecast that outperforms the historic mean can be said to explain the level.

Table 1 compares the RBA's forecast errors with those of uninformative alternatives. We show results at horizons 3 quarters and 7 quarters after the forecast, a cut of the data that avoids the duplication arising from overlapping 4-quarter changes but still summarises most of the sample. We describe these horizons as the first year and second year of forecasts, recognising that the current quarter is covered by the 3-quarter-ahead forecast. Appendix D shows comparisons at other horizons.

The top row of Table 1 shows that the RMSE for underlying inflation in the first year of the forecast horizon is 0.54 percentage points (column (4)). This can be compared with forecasts that inflation will remain at its rate over the preceding four quarters. This 'no change' or 'random walk' forecast has an RMSE of 0.73 percentage points (column (5)).[3] The RMSE of a random walk forecast equals the (uncentred) standard deviation of changes.[4] An RMSE ratio (column (6)) less than one – 0.74 in this case – indicates that the forecast is able to explain some of the variation in changes in underlying inflation. This may sound a trivial accomplishment, but it is one that foreign central banks have often not achieved. For example, Atkeson and Ohanian (2001) find that CPI forecasts of the US Federal Reserve are less accurate than a random walk. Variations on this result using other sample periods and measures of inflation are reported by Reifschneider and Tulip (2007), Tulip (2009) and Edge and Gurkaynak (2011). Similarly, Goodhart (2004, p13) reports that the Bank of England 'does not appear to be able to provide any predictive guide at all to the fluctuations of output growth, or inflation, around its trend over a year in advance'.

The superior accuracy of the RBA forecast over the random walk is statistically significant, with a p-value of 2 per cent (column (7)). These p-values, constructed from Diebold-Mariano (1995) tests,[5] represent the chance that we might see the differences between the forecasts' mean squared errors if the forecasts were equally accurate.

3 For consistency of comparisons, we only calculate errors for the alternative for those quarters for which there is a comparable forecast.

4 An uncentred standard deviation, variance or R^2 measures deviations about the population mean of zero instead of about the sample mean. In our context, centred and uncentred statistics have much the same interpretation and are empirically quite close.

5 We regress the difference in the squared errors on a constant and report the p-value from a t-test of the hypothesis that the constant is zero. We use Newey and West's (1987, 1994) autocorrelation-robust standard errors, calculated with their suggested lag truncation, which typically is three quarters. The reliability of Newey-West variances is not clear, given the small size of our samples, the non-normality of squared errors, and the moving-average (MA) structure of our data. We explored alternatives that address some of these issues, specifically alternative bandwidth selection rules, West's (1997) MA-robust standard errors, and a block-bootstrap. But none of these approaches address all the features of our data.

Table 1: RBA RMSEs Relative to Variation in the Data

Variable	Null alternative	Horizon	RMSE: RBA	RMSE: Alternative	RMSE: Ratio	Significance p-value	Uncentred R^2
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Underlying inflation; 4-quarter percentage change	Random walk	First year	0.54	0.73	0.74	.02	0.46
		Second year	0.80	1.08	0.74	.15	0.46
Underlying inflation; 4-quarter percentage change	Target	First year	0.54	0.74	0.72	.06	0.48
		Second year	0.80	0.78	1.03	.87	–0.05
CPI inflation; 4-quarter percentage change	Random walk	First year	0.89	1.90	0.47	.00	0.78
		Second year	1.27	2.19	0.58	.03	0.67
CPI inflation; 4-quarter percentage change	Target	First year	0.89	1.41	0.63	.04	0.60
		Second year	1.27	1.36	0.93	.78	0.13
GDP growth; 4-quarter percentage change	Historical mean	First year	1.44	1.28	1.13	.23	–0.28
		Second year	1.39	1.39	1.00	.94	–0.01
Unemployment rate; 4-quarter percentage change	Random walk	First year	0.62	0.67	0.92	.63	0.15
		Second year	0.97	0.89	1.10	.69	–0.20

A direct measure of the share of the variation in the data that is explained by the forecast is an uncentred R^2 statistic (Hayashi 2000, p20), defined as

$$R^2 = 1 - \frac{\sum_{t=1}^{n}\left(y_t - f_t\right)^2}{\sum_{t=1}^{n}\left(y_t - \hat{\mu}_t\right)^2} \tag{1}$$
$$= 1 - \frac{MSE}{Variance}$$

where: y_t is the variable being forecast; f_t is its forecast; and $\hat{\mu}_t$ is the uninformative forecast, which can often be interpreted as the population mean.[6] Weather forecasters refer to this measure as a 'skill score' (Murphy 1988). An R^2 of zero, meaning the forecast has no explanatory power, occurs when the forecast is as accurate as the uninformative alternative. When the alternative is a random walk, it is simple to think of the variable being forecast as being in changes, with a mean of zero.[7] The *MSE* and (uncentred) *Variance*, represent the square of the RMSEs shown in columns (4) and (5) respectively. The second line of Equation (1) follows from the first by dividing both numerator and denominator on the first line by n, the number of observations. The R^2 estimate of 0.46 shown in row 1, column (8) of Table 1 indicates that the RBA's forecasts account for about half the variance of changes in underlying inflation over the first forecast year.

Another benchmark is the midpoint of the RBA's target range for inflation, 2.5 per cent (henceforth, 'the target'). Comparisons between the forecast and this benchmark are shown in the next part of Table 1. The R^2 of 0.48 indicates that RBA first-year forecasts of underlying inflation account for about half the

---

6 The R^2 measure we present should not be confused with the R^2 from a hypothetical 'Mincer-Zarnowitz' regression of actual outcomes on the forecast. Conceptually, this hypothetical R^2 would equal ours if the coefficient on the forecast were constrained to equal one, the intercept was constrained to equal zero, and the dependent variable was measured as deviations from the uninformative alternative. Mincer-Zarnowitz regressions are popular. However, decision-makers need to form judgements about the explanatory power of the forecast, not the explanatory power of $\alpha + \beta Forecast$, where α and β are parameters that are estimated after outcomes are known.

7 The change in inflation, h quarters ahead, which is equal to the forecast error from the random walk forecast, is measured as $\pi_{t+h} - \pi_{t-1}$, where π_k is the percentage change in prices in the four quarters to k and t is the quarter in which the forecast is made.

deviations of underlying inflation from the target. However, over the second forecast year, the ratio of RMSEs is about one and the R^2 is about zero. So a forecast of 2.5 per cent was about as accurate a guide to underlying inflation as the second-year RBA forecast. This result is consistent with successful targeting of the inflation rate. At horizons over which monetary policy has a substantial influence, deviations of inflation from the target should generally be unpredictable. If there were predictable deviations, it would mean that the central bank was expecting that it would miss its target and was not acting to prevent this. See Edey and Stone (2004) for further discussion of forecast deviations of inflation from the target.

Results for CPI inflation are shown in the next two parts of the table. Again, forecasts have substantial explanatory power in both levels and changes. The first-year forecasts significantly outperform both a random walk and the target. One feature of the CPI estimates (which was less clear for underlying inflation) is that the target is more accurate than the random walk, reflecting rapid reversion of headline inflation to the mean. Reflecting this mean-reversion, the RBA's forecasts outperform a random walk more often than they beat the target. Put more simply, the forecasts can successfully predict changes in inflation, even when it is difficult to predict the level of inflation.

Two differences between the results for underlying inflation and the CPI are worth noting. First, as shown in column (4), forecast errors are considerably smaller for underlying inflation than for CPI inflation. That, of course, is one reason many economists like to focus on underlying inflation. We know more about it than we do about CPI inflation. The RBA has invested substantial resources in constructing measures of underlying inflation with higher signal/noise ratios (see Richards and Rosewall (2010) and references cited therein). The greater predictability of underlying inflation relative to the headline CPI is a reflection of that effort.

Second, the RBA's forecasts for headline inflation have had more explanatory power than those for underlying inflation, as measured by the R^2 estimates in column (8). This largely reflects the spike in the CPI in 2000:Q3, due to the introduction of the Goods and Services Tax (GST), which was factored in to the CPI forecasts from 1999:Q1. See Figure 1, top right panel. The GST had minimal direct effect on the measure of underlying inflation used at the time, which was the weighted median excluding interest and taxes.

For GDP growth, our uninformative alternative forecast is the historic (since 1959) mean.[8] For the first-year GDP forecast, this alternative is more accurate than the RBA's forecast. That is, forecasts have less explanatory power than the mean. Reflecting this, the RMSE ratio is greater than one and the R^2 is negative. For the second-year GDP forecast, the forecast is as accurate as the mean, so the R^2 is zero.

Low and even negative forecast R^2s are not unusual. They have been found by many researchers for many different kinds of macroeconomic forecasts. For example, Vogel (2007, Table 3) finds them for both Consensus Economics and OECD forecasts of GDP growth in the G7 economies. Atkeson and Ohanian (2001) implicitly find them for the US Federal Reserve's forecast of changes in the US CPI. Campbell (2007) finds them for the Survey of Professional Forecasters' forecasts of US GDP growth. Tulip (2009) finds them for Federal Reserve forecasts of US GDP growth and the GDP deflator. Goodhart (2004, Table 5) reports more dramatic results for the Bank of England's GDP forecasts (specifically outcomes are negatively correlated with forecasts).

That said, the low explanatory power of macroeconomic forecasts is a striking result, with important implications. For example, it affects how much weight should be placed upon forecasts of GDP in determining macroeconomic policy. More generally, it is relevant to debates as to whether policy should be 'backward looking' (as in some Taylor rules) or 'forward looking' (as in optimal control exercises).

Results for the unemployment rate are shown at the bottom of Table 1. We use the previous level as an alternative forecast, which is equivalent to examining whether unemployment forecasts outperform a random walk. The R^2 can be interpreted as measuring how much of the variance of *changes* in the unemployment rate is explained by the forecast. Short-horizon unemployment forecasts seem to have some explanatory power, accounting for 15 per cent of changes in unemployment

---

8 We measure the historic mean as average GDP growth from 1959 through to the quarter preceding the forecast, measured using real time data from Stone and Wardrop (2002), kindly updated for us by Tim Robinson. Similar data (more thoroughly documented) are now publicly available at the website of the Department of Economics at the University of Melbourne (http://www.economics.unimelb.edu.au/RTAustralianMacroDatabase/Database%20and%20Documentation.html).

over the first year. But at longer horizons, the forecasts have been less accurate than a random walk.

To summarise the results in this section, the forecasts have substantial explanatory power for both the level and change in inflation over the next year, but – consistent with successful inflation targeting – at longer horizons deviations in underlying inflation from the RBA's target seem to be unpredictable. Uncertainty about the forecasts for GDP growth and (beyond the immediate horizon) changes in unemployment is about the same as the variation in these variables. In other words, forecasts for these variables lack explanatory power.

The ability to predict short-term variations in inflation but not in activity might be interpreted in different ways. One possibility is that the two variables are unrelated: the Phillips curve is flat. However, empirical evidence of many forms from many countries is inconsistent with that view. Another interpretation is that GDP growth is a poor measure of inflationary pressures, perhaps because it reflects changes in supply conditions or because it is the *level* of activity (relative to potential) that affects inflation, rather than the growth rate. Related to this, it may be that the RBA's implicit forecasts of the output gap usefully inform the inflation forecast; though this signal is difficult to discern after the event due to supply shocks. A third possibility is that influences on inflation other than demand are important. Whatever the explanation, the different explanatory power of forecasts for different variables has clearer implications for the presentation of the outlook. Specifically, we can talk more confidently about the near-term outlook for inflation than we can about the outlook for GDP growth. That emphasis is reflected in the *SMP*.

4.3 Uncertainty about Others' Forecasts

A benchmark that is especially relevant to improving forecast performance is the accuracy of other forecasters. To this end, we examine the output and inflation forecasts provided by Consensus Economics, a regular survey of about two dozen private sector forecasters. We use the average forecasts of 4-quarter changes in the CPI and real GDP, which we have on a quarterly basis since December 1994. We focus on these forecasts, rather than those of year-average changes which Consensus publishes more frequently, to facilitate comparisons with the forecasts published in the *SMP*. Consensus forecasts are proprietary, available via subscription at www.consensuseconomics.com. Summary statistics are reported here with permission.

As shown in Table 2, RBA forecasts for CPI inflation have been slightly more accurate than those of Consensus at all horizons. The differences are small and not statistically significant.

Table 2: CPI Inflation Forecast RMSEs – RBA and Consensus
4-quarter change; 1994:Q4–2011:Q2

Horizon	**RMSE**			**Significance**
(quarters ahead)	RBA	Consensus	Ratio	p-value
0	0.27	0.31	0.86	.15
1	0.49	0.50	0.97	.74
2	0.71	0.71	0.99	.93
3	0.87	0.93	0.94	.48
4	1.07	1.14	0.94	.37
5	1.20	1.31	0.92	.21
6	1.21	1.35	0.90	.15
7	1.22	1.24	0.99	.90

As shown in Table 3, Consensus forecasts of GDP growth have been significantly more accurate than those of the RBA.

Table 3: GDP Growth Forecast RMSEs – RBA and Consensus
4-quarter change; 1994:Q4–2011:Q2

Horizon	**RMSE**			**Significance**
(quarters ahead)	RBA	Consensus	Ratio	p-value
0	0.80	0.78	1.03	.65
1	1.12	1.00	1.12	.03
2	1.28	1.16	1.11	.04
3	1.37	1.24	1.11	.02
4	1.33	1.22	1.10	.02
5	1.33	1.19	1.11	.07
6	1.35	1.18	1.14	.03
7	1.34	1.21	1.11	.08

There are several possible reasons for this result, though it is not clear that these fully account for the difference. For example, Consensus has a timing advantage for the last few years of the sample. Consensus conducts its survey in the last month of the quarter, after the publication of the national accounts. That is similar timing to RBA forecasts until 2008, after which its forecast for GDP was published in the middle month of the quarter. However, given that forecast accuracy does not vary much with the horizon (as can be seen in the table, and discussed further in Section 5.1), this advantage is not important other than for very short horizons. For horizons beyond one year, Consensus forecasts published the previous quarter outperform RBA forecasts by similar margins.

Another possible reason for the greater accuracy of Consensus is that their interest rate assumptions may be more realistic. As discussed in Appendix B, the RBA's GDP errors are significantly correlated with the slope of the yield curve at the time of the forecast, a measure of how much the market expects interest rates to change. However, estimates of this effect are small. When an estimate of the effect of the yield curve is removed from the RBA's GDP errors, the RMSE declines by 5 per cent at a 3-quarter-ahead horizon. This is still larger than the Consensus RMSE, though the difference is no longer statistically significant.[9]

The competitive (or 'horse race') aspect of these comparisons has curiosity value and is relevant to improvements in forecasting technique. For example, it implies the RBA might improve the accuracy of its GDP forecasts by placing greater weight on the Consensus average.[10] However, for our purposes, the similarity of the forecast errors may be more important than their differences. This is illustrated in Figure 4, which shows 3-quarter-ahead forecast errors for year-ended GDP growth for the RBA and Consensus. Differences between the errors are small relative to the variation in the data. At this 3-quarter-ahead horizon, the difference between the RMSE of the RBA (1.37 percentage points) and that of Consensus

---

9 Moreover, this comparison overstates the importance of differences in interest rate assumptions, given that we make no corresponding adjustment to Consensus and that the adjustment uses information after the event which was not available at the time of the forecasts.

10 It is possible that individual members of the Consensus panel might also be able to improve their forecasts by moving toward the mean. Of course, were many members of Consensus to do this, the behaviour of the average would noticeably change, becoming subject to herd dynamics.

(1.24 percentage points) is statistically significant (p = .02). However this difference is not obviously significant in economic terms, being close to rounding error.

Figure 4: RBA and Consensus GDP Forecast Errors

3-quarters-ahead

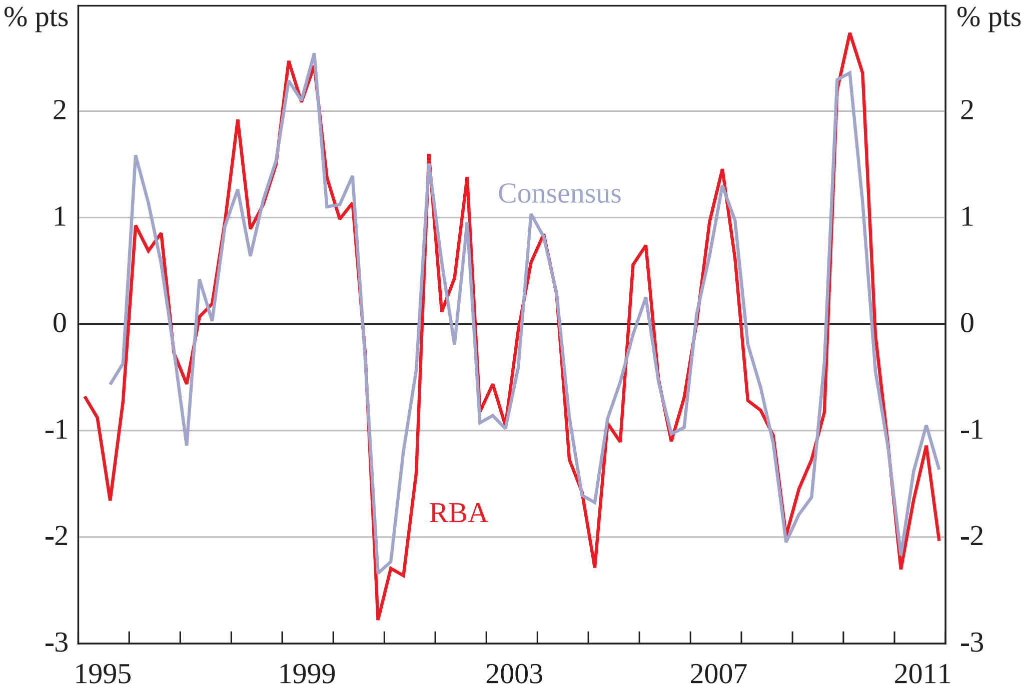

Sources: Consensus Economics; RBA; authors' calculations

Overall, differences in forecast accuracy seem small, with relative performance varying across different variables. So, in qualitative terms, uncertainty about the Consensus forecast seems to be about the same as uncertainty about the RBA forecast. This similarity in accuracy is often found in comparisons of macroeconomic forecasters (for example, Reifschneider and Tulip (2007)). One implication of that similarity is that inferences about uncertainty around the RBA's forecast can (cautiously) be based on the track record of private sector forecasts and vice versa.

4.4 Disagreements and Revisions

It is common for forecasters to argue over a disagreement of say half a percentage point in their GDP forecasts. Revisions of similar size are often described as a substantial change in the outlook. As Stevens (2011) discusses, the magnitude of forecast errors provides useful context for assessing these differences. Some 22 per cent of the RBA's forecasts of GDP growth over the following four quarters were accurate to within half a percentage point. The results in Section 4.3 suggest that private sector errors would be similar. So, even if one forecast were the most likely outcome, the likelihood of an outcome closer to the alternative would be high. That is, one cannot have confidence that one forecast is correct and that a similar forecast is not. Put slightly differently, if the 90 per cent confidence interval for one forecast spans a range of 1 to 6 per cent while that of another (or previous) forecast spans a range of 2 to 7 per cent, those forecasts should be seen as being in substantial agreement.

That said, when the costs of policy mistakes are symmetric, then policy decisions should be made on the basis of the central tendency of forecasts. Then differences in the central tendency can have important implications and explanations of revisions and disagreements can be useful.

5. Other Properties of the Confidence Intervals

The confidence intervals presented in Figure 3 are a blend of data and assumptions. In particular, we estimate the width of our confidence intervals, at each horizon, based on the empirical record. However, we assume that some other properties of past historical errors are unlikely to apply in the future. Specifically, our intervals are unbiased and symmetric. Whether normality should also be assumed depends on the purpose. We discuss each of these assumptions in turn.

5.1 Effect of the Horizon

One might expect that the further events are in the future, the harder they are to forecast. That is, fan charts should fan out. This may seem to occur in Figures 2 and 3. However, much of the fanning in those charts is an artefact of data construction. Near-term forecasts for 4-quarter changes are made within the period being forecast and include ABS estimates for the first few quarters. It is not surprising that uncertainty widens as the event being forecast includes less of the known past. This effect is larger for forecasts of year-average changes. When this effect is removed, there is surprisingly little effect of the horizon on uncertainty about growth rates.

The effect of 4-quarter changes on increasing uncertainty can be removed by examining *quarterly* changes. Figure 5 shows RMSEs by horizon for forecasts of quarterly GDP growth. The dark line denotes RMSEs using the same sample of forecasts as used above, for example in the lower left panel of Figure 2. These estimates are essentially unaffected by the horizon. That is, we seem to know about as much (or as little) about GDP in the current quarter as we do about GDP growth two years ahead.

Figure 5: RMSE by Horizon

Quarterly GDP change

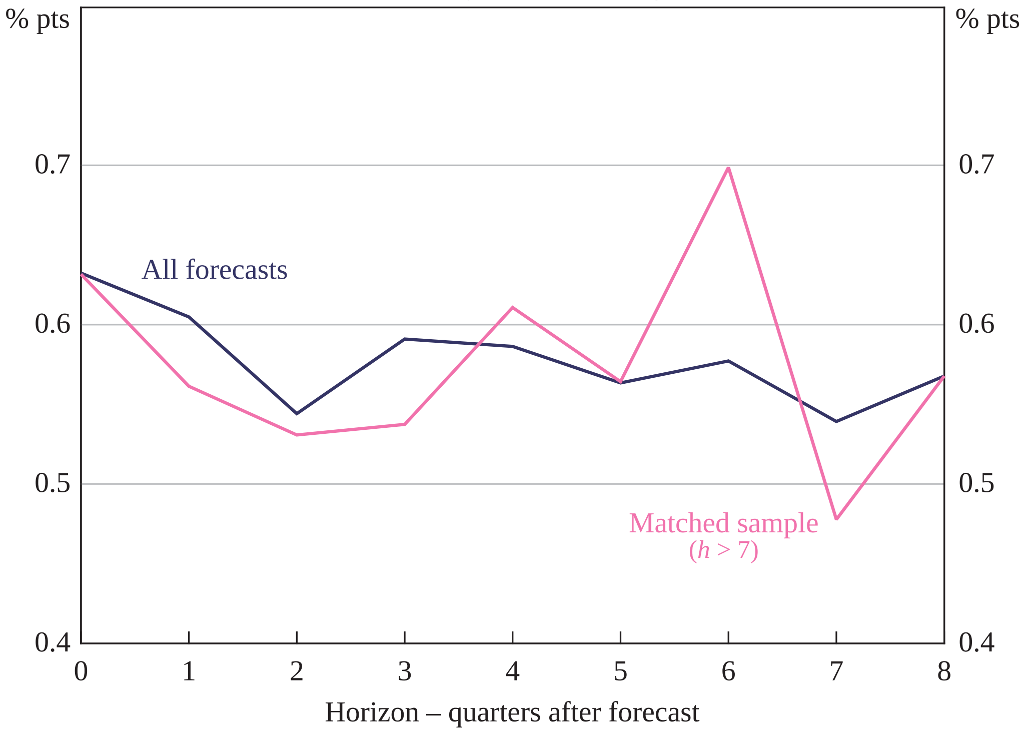

Another factor that distorts the effect of the horizon on uncertainty is changes in the sample. Because the forecast horizon has increased over time, our sample of long-horizon forecasts is smaller and more recent than our sample of short-horizon forecasts. So the GFC, for example, has a larger effect on our sample of long-horizon forecasts, than on our sample of short-horizon forecasts. To avoid differences like these affecting comparisons, we restrict the sample to the 24 forecasts in which the horizon extended at least 8 quarters ahead. The light line in Figure 5, labelled 'h > 7', shows RMSEs from this matched sample. The RMSEs shown by this line have the same initial conditions, with only the horizon changing. The matched sample estimates are more volatile. Still, uncertainty does not seem to increase with the horizon.

One implication of these results is that surprises to GDP growth are not persistent. That is, there is little inertia or momentum in the unpredictable component of GDP growth.[11] If there were substantial momentum, the surprises would accumulate and the fan chart would fan out.

Another implication is for comparisons between forecasts of GDP growth made at different times in a quarter. The absence of a substantial effect of the horizon on forecast errors means that one need not worry too much about the precise timing of forecasts or whether one forecast had an informational advantage over the other. Whether a forecast we record as being 3 quarters ahead is really 2 or 4 quarters ahead will make little difference to our results. Researchers in the United States (Romer and Romer 2000; Edge and Gurkaynak 2011) have precisely calculated the timing of economic forecasts relative to data releases and other forecasts. For analysis of Australian GDP forecasts, the precise timing of which does not seem to matter, that effort seems unlikely to be worthwhile.

Whether uncertainty about *inflation* increases with the horizon is harder to assess, given that much of our source data is in year-ended-change format. Internal RBA estimates suggest that current-quarter CPI forecasts benefit from high-frequency data on oil and food prices. Whether this advantage helps forecast inflation in following quarters is less clear. In Figures 2 and 3 the dispersion of inflation forecast errors widens with the horizon, even beyond a horizon of 3 quarters ahead. However, this 'fanning' could reflect changes in the sample: we have more forecasts with relatively short horizons and these forecasts could have been for periods when inflation did not behave unusually. That possibility is consistent with a matched sample of forecasts for year-ended underlying inflation for which the horizon extends at least 8 quarters, in which RMSEs are flat at horizons beyond 3 quarters. The limited data we have on quarterly underlying inflation also show a surprisingly small effect of the horizon on RMSEs, when initial conditions are held constant. However, given the small sample, these results are not strong.

For other variables, there is a stronger effect of the horizon. This is clearest for variables measured in levels, such as the unemployment rate or the level of GDP. Although surprises to growth rates are not persistent, those to levels are.

---

11 Results in Section 4.2 suggest that virtually all variations in GDP growth are unpredictable. So there is little momentum in total GDP growth.

5.2 Bias

We have assumed that confidence intervals are centred on the forecast. An alternative assumption would be to centre the intervals on the forecast plus the mean or median error. Whether this matters depends on whether average errors have differed from zero. To assess this, we regress past forecast errors for each variable at each horizon on a constant. Results are reported in Table 4. The coefficient on the constant represents the average amount by which outcomes exceed forecasts. This is reported in the columns labelled 'bias'. Whether this bias is large relative to the noise in the data can be gauged by t-tests for the hypothesis that the constant is zero. P-values for these tests, calculated using autocorrelation-robust standard errors, are also reported in the table.

Table 4: Forecast Bias

Horizon (quarters ahead)	**Underlying inflation**		**CPI inflation**		**GDP growth**		**Unemployment rate**	
	Bias	p-value	Bias	p-value	Bias	p-value	Bias	p-value
0	–0.01	.78	–0.04	.13	0.18	.12	–0.07	.00
1	–0.00	.99	–0.04	.56	0.18	.30	–0.16	.00
2	–0.00	.96	–0.02	.83	0.17	.45	–0.25	.00
3	–0.03	.81	–0.06	.71	0.07	.80	–0.31	.00
4	–0.04	.74	–0.07	.76	–0.13	.60	–0.36	.01
5	–0.07	.61	–0.11	.67	–0.29	.22	–0.37	.01
6	–0.06	.71	–0.09	.73	–0.44	.10	–0.40	.02
7	–0.02	.91	–0.16	.58	–0.59	.04	–0.39	.07
8	0.06	.73	0.06	.86	–0.71	.02	–0.61	.06

As can be seen in the left half of the table, bias in the inflation forecasts is approximately zero over this sample period. So centring confidence intervals for inflation on the forecast is in line with past experience. In contrast, GDP forecasts were too low at short horizons and too high at longer horizons, though the results are generally not significant. However, outcomes for the unemployment rate have, on average, been significantly below expectations. For example, the unemployment rate 3 quarters after the forecast averaged 0.3 percentage points below its prediction ($p = 0.004$). As can be seen in Figure 1 (third panel on right), the downtrend in unemployment was consistently underestimated during our sample period.[12]

A finding of *ex post* forecast bias is not unusual. It commonly occurs when there is a persistent shock to the economy that forecasters learn about gradually. For example, forecasters in many OECD countries persistently understated the rate of inflation in the 1960s and 1970s, then persistently overstated it in the 1980s and 1990s. Over the past decade, rising oil prices have repeatedly pushed headline inflation above expectations. Over the same period, Kearns and Lowe (2011, Figure 10) show that the RBA consistently underestimated the strength of the terms of trade.

However, it is doubtful whether the bias in these small samples is representative of the population. The errors may be random or systematic. If the latter, forecasters can be expected to learn and adjust their forecast. Neither case is likely to persist. So even when the errors do not have a zero mean (*ex post*, or after the event), we would still centre the confidence intervals about the forecast. That is, we assume past bias will not continue.

5.3 Symmetry

It is often suggested that the distribution about a particular forecast is skewed. For example, the November 2011 *SMP* described the risks to the central projection of global activity as skewed to the downside because of financial problems in the euro area. In practice, empirical estimates of skewness are difficult to interpret when the sample mean is not zero. Furthermore, they often reflect large outliers, which are observed infrequently. Neither small-sample bias nor outliers are a reliable guide to the future.

12 Three observations might be interesting to note. First, the bias in the unemployment forecasts was accompanied by zero bias in underlying inflation. This suggests that the bias in the unemployment forecasts was offset by similar bias in the NAIRU and/or unanticipated appreciation of the exchange rate. Second, although unemployment fell over the sample, it is stationary over long time periods. It may be that unemployment would be systematically underpredicted when it trends upward. If so, the problem would be persistence in errors, rather than bias. Third, much (though not all) of the bias reflects predictions of rising unemployment at times of the Asian financial crisis, the global slowdown of the early 2000s and the global financial crisis that did not come to pass.

For simplicity, the confidence intervals in Figure 3 assume that errors are symmetric. We are not arguing that the distribution about future forecasts should be assumed to be symmetric. Rather, judgements about skewness are likely to be based on information other than the skewness in the historical data.

5.4 Normality

In Figure 3, we reported quantiles of the empirical distribution of forecast errors. A more common approach to estimating confidence intervals is to assume that errors follow a known distribution, the parameters of which can be estimated or imposed. In particular, many foreign central banks assume that the errors are normally distributed with a zero mean and a standard deviation equal to that of a sample of past errors (see Appendix A).

In practice, the two approaches provide similar estimates. This is shown in Figure 6, which compares RMSEs with the 68^{th} percentile of the empirical distribution of absolute errors for our four variables. Were the errors distributed normally, these estimates would be equal. In practice, they differ slightly, with the differences typically being within rounding error.[13]

The similarity of the estimates in Figure 6 suggests that normality is a reasonable description of the data. To be more precise, a confidence interval equal to the forecast plus and minus one RMSE can realistically be described as approximately ‘a two-thirds confidence interval’. Whether normality is a reasonable description as one extends into the tails is harder to assess. In a sample of 19 years, we have very few ‘once-in-a-generation’ surprises. However, the experience of other countries, noted below, suggests that large surprises may be more frequent than implied by a normal distribution. Haldane (2012) argues that macroeconomic surprises often have fat tails.

---

13 Alternative approaches to this issue include a Jarque-Bera test or comparing histograms with a normal distribution. These alternatives test for normality at places in the distribution, such as near the centre, which are less interesting for the purpose of constructing confidence intervals. A Jarque-Bera test (which would need to be adjusted for serial correlation, presumably by Monte Carlo) gauges statistical significance, whereas Figure 6 indicates the magnitude (or ‘economic significance’) of departures from normality.

Figure 6: Alternative Measures of Error Dispersion
1993:Q1–2011:Q4

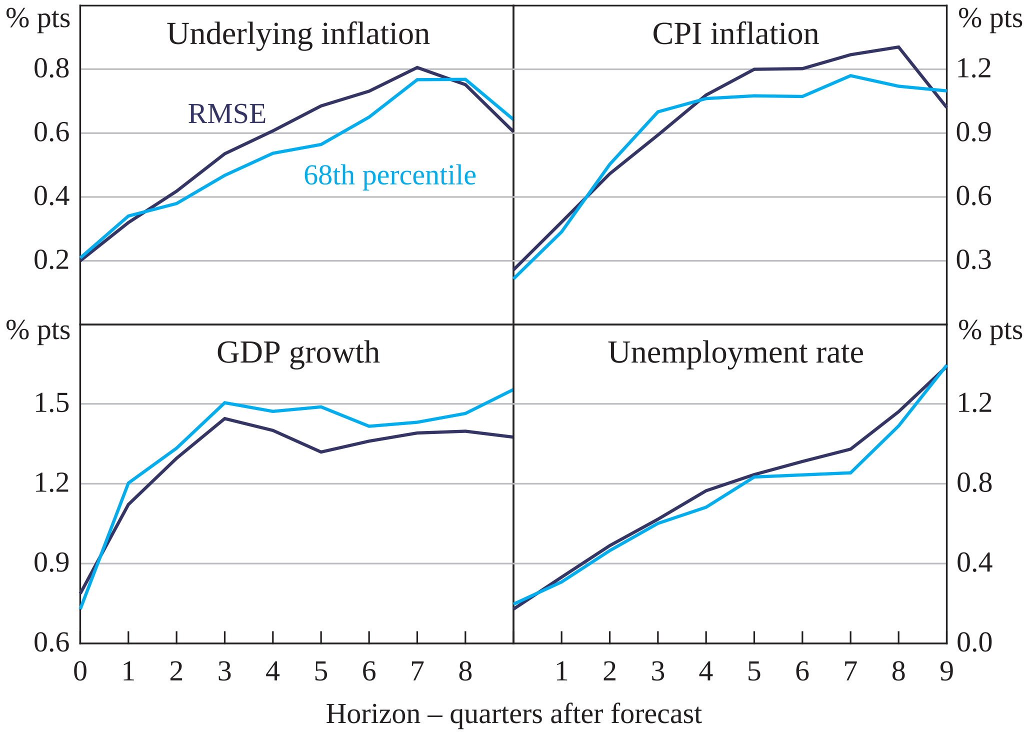

As a measure of dispersion, quantiles and RMSEs have different advantages. If the purpose is to construct confidence intervals (especially other than two-thirds), then quantiles are simple, direct and do not require a questionable assumption about normality. However, for a summary measure of uncertainty or forecast comparisons, RMSEs are more comprehensive, with useful statistical properties. For example, they map into analysis of variance, they are easily scaleable, and they have less sampling variability: in small samples, RMSEs jump around less unpredictably than quantiles.[14]

14 For example, we construct 100 000 Monte Carlo simulations of artificial data with similar ARMA properties to our 3-quarter-ahead underlying inflation errors. The mean RMSE and 68th percentile (precisely, the 0.6827 quantile) are both 0.5 percentage points. The standard deviation of the RMSE is 0.070 percentage points, while that of the 68th percentile is 0.085 percentage points, about one-fifth larger.

6. Alternatives and Limitations

6.1 Forecast Versus Model Errors

Past forecast errors are one possible method of gauging forecast uncertainty. Another possible approach is to use an economic model. Model-based estimates of forecast uncertainty are considerably more flexible than estimates based on historical forecast errors. Model-based estimates can accommodate new variables, longer horizons, new information sets or forecasting techniques, and so on. Model-based estimates are attractive whenever there is a substantial change in economic structure, the forecasting framework, or in how the forecast is presented.

The disadvantage of model-based confidence intervals is that they can be unrealistic, though the direction of overall bias will vary. Models can overstate uncertainty because they use too little information. They tend to be considerably simpler than the approach used by forecasters, who often pool together multiple data sets, forecasting techniques and qualitative data. Reflecting this, no one model provides a good guide to the Reserve Bank forecast, in which judgement plays a large role.

However, models can also understate uncertainty because, in a different sense, they use too much information. Most obviously, they are typically estimated using revised data that was only available after the forecast. In principle, recursive estimation using real-time data ('pseudo-real-time forecasting') can remove this advantage, and mimic some of the conditions of actual forecasting. But perhaps more important are issues of model specification. After the event, it may seem obvious that some variables trended while others were stationary, and that some influences were important while others were irrelevant. But in real time, these issues are often unclear. Unless one can use models actually used at the time, it is very difficult to deprive a model specification of the benefit of hindsight.

To illustrate, we compare our estimates of uncertainty with those of a representative Bayesian VAR, specifically, the 'BVAR2' model of Gerard and Nimark (2008, Section 2.1).[15] This comprises two lags of domestic GDP growth, underlying inflation, the cash rate and the exchange rate, together with foreign

---

15 Thanks to Penny Smith for constructing these estimates.

output growth, inflation and interest rates. Gerard and Nimark provide further details. We depart from their specification by including CPI inflation and the unemployment rate, using the OLS estimate of the shock covariance, and estimating over 1993:Q2–2011:Q3.

Figure 7 compares the 70th percentile of absolute errors from the two approaches, a standard model-based metric. This measure, approximately one standard deviation, is added on either side of the central forecast to construct the confidence intervals in Figure 3. It is often referred to as a 'half-70 per cent confidence interval'.[16] The model suggests more certainty about the outlook for unemployment and for GDP growth than do past forecast errors, but less certainty about CPI inflation. These differences presumably reflect different information sets. For example, the model knows about the downtrend in unemployment over this sample, though this was not known before the event. In contrast, the model did not know in advance about the introduction of the GST. More could be said about these alternative estimates. For example, they are presumably sensitive to changes in sample period or model specification. However, the key point is that the estimates of uncertainty are broadly similar. Hence model estimates could be used in place of, or to augment, estimates based on past forecast errors.

16 Ordinarily, one might expect the 68th percentile of the distribution of forecast errors (shown in Figure 6) and 70th percentile (shown in Figure 7) to be extremely close. But with small samples, as we have at far horizons, sizeable differences can arise. This is an example of the sampling variability discussed in footnote 13.

Figure 7: Forecast Versus Model Errors

Half-70 per cent confidence intervals

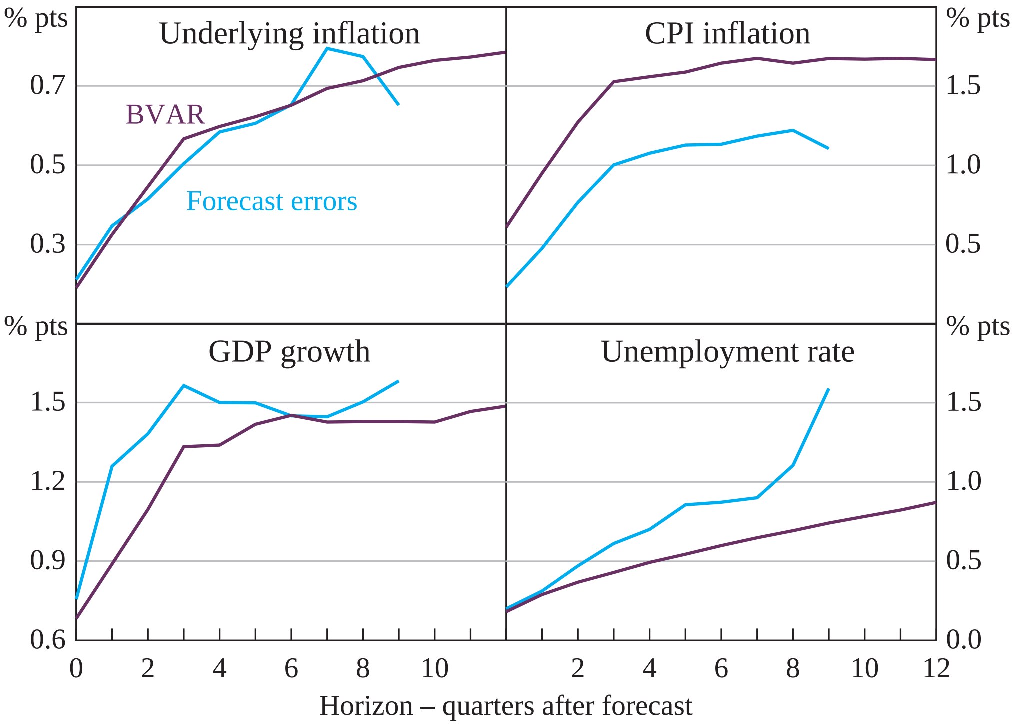

6.2 Past Forecast Errors are an Unreliable Guide to the Future

Forecastability will vary with economic conditions. For example, forecast errors are likely to be greater than usual following large macroeconomic shocks. In principle, it might be tempting to address these changes through models of conditional heteroskedasticity. In practice, the macroeconomic shocks that give rise to the greatest uncertainty are often unusual and difficult to quantify, so data-based modelling is unreliable.

Foreign central banks have dealt with this issue in different ways. The Bank of England calculate average historical forecast errors, then adjust these judgementally, so as to reflect the uncertainty about the present forecast more accurately. The US Federal Reserve presents estimates of average forecast errors over the previous twenty years, coupled with a qualitative assessment of the degree to which current uncertainty may differ from normal. For references, see Appendix A.

A more fundamental problem is that average levels of uncertainty seem to be unstable. For example, there was a large reduction in the variability and predictability of macroeconomic conditions in many OECD countries in the early 1980s, the so-called 'Great Moderation'. There have been large increases in unpredictability (for some variables and countries) over the past few years associated with the global financial crisis.

The difficulties this instability poses are illustrated by the experience of the US Federal Reserve.[17] The Federal Open Market Committee (FOMC) started releasing estimates of uncertainty with its forecasts in November 2007. (These should not be confused with the measures of disagreement among FOMC participants that the Fed also releases.) At that time the midpoint of FOMC projections of the unemployment rate in 2010:Q4 was 4.8 per cent. This projection was associated with a RMSE of 1.1 percentage points, which was the average over the previous twenty years (FOMC 2007, pp 10 and 12). In the event, the unemployment rate rose to 9.6 per cent, representing a forecast error of 4.8 per cent, or 4.4 RMSEs. If forecast errors were normally distributed, with a constant mean and variance, then such an error should occur once every 80 000 forecasts. More likely explanations are that the variance is not actually constant or (perhaps more plausibly, given that subsequent errors have been modest) that the distribution is not normal.

A corollary of this instability is that our estimates of uncertainty are sensitive to the sample period we have used. Were we to include forecast errors from the recession and disinflation of the early 1990s, our RMSEs would increase. Were we to include the volatile 1970s and 1980s, our RMSEs would presumably increase further.

A closely related problem is that there have been many substantive, presentational, and other changes in the forecasts over the past two decades. Some of these are noted in Appendix B. Perhaps more important, the economy we are trying to predict continues to evolve. So past performance is only a rough guide to the future.

---

17 Full disclosure: one of us was involved in producing the Fed's estimates, see Reifschneider and Tulip (2007).

Although this instability means that future uncertainty may differ from the past, it does not indicate in which direction. Historical estimates are not obviously biased. This is in contrast to, for example, subjective estimates of uncertainty, which numerous studies in other areas have found to be overconfident (see footnote 2).

7. Conclusion

Using past forecast errors, we construct confidence intervals about the *SMP* forecasts of underlying inflation, the CPI, real GDP growth and the unemployment rate. These intervals, shown in Figure 3, span a wide range of outcomes. Our estimates indicate that uncertainty about the forecast is high, particularly for GDP growth.

These estimates can be compared to various benchmarks, such as the amount of variation in the data. We find that RBA forecasts have substantial explanatory power for inflation over the first forecast year. However, like many other forecasters, the RBA forecasts explain very little of the variations in GDP growth, medium-term changes in unemployment, or the medium-term deviations of underlying inflation from the target.

Uncertainty about RBA forecasts is similar to that about private sector forecasts. More precisely, RBA forecasts of inflation have been marginally more accurate than the average of private sector forecasts, while RBA forecasts of GDP growth have been less accurate. However, the differences are not large.

This paper is part of a broader process of review. The RBA continuously examines its forecasting performance with a view to understanding the economy and improving forecast accuracy. Changing conditions and ongoing research lead to new techniques and information sets being regularly adopted. But even as the forecasts evolve, considerable uncertainty will remain.

Appendix A: Measures of Uncertainty Presented by Foreign Central Banks

Institution	Measures	Method of construction	For which variables	References
European Central Bank	Range with no point estimate	Twice the mean absolute error made in the past, with outliers excluded (‘consistent with a 57.5 per cent confidence interval’).	Real GDP and its components; Harmonised Index of Consumer Prices	ECB (2009, 2011)
US Federal Reserve	RMSEs	Average of RMSEs for past 20 years of six leading forecasters. Accompanied by a qualitative description of how uncertainty may be unusual.	Real GDP growth; unemployment rate; total consumer prices	FOMC (2007); Reifschneider and Tulip (2007)
Bank of England	10, 20, 30, ... 90 per cent confidence intervals, and more	Based on forecast errors over past 10 years, assuming a normal distribution. The dispersion and skewness are then judgementally adjusted.	Level and growth of real GDP (including revisions); CPI	Bank of England (2011); Elder *et al* (2005)
Bank of Japan	Histograms	Average of probability distributions assumed by individual Board members	Real GDP growth; CPI excluding fresh food	Bank of Japan (2008, 2011)
Bank of Canada	50 and 90 per cent confidence intervals	Combination of historical forecast errors and model errors	Core CPI; total CPI	Bank of Canada (2009, 2011)
Sveriges Riksbank	50, 75 and 90 per cent confidence intervals	RMSEs of past forecast errors by Riksbank and implied forward rates (adjusted for risk premia), assuming normality	Real GDP growth; CPI; core inflation; repo rate	Sveriges Riksbank (2007)
Norges Bank	30, 50, 70 and 90 per cent confidence intervals	Macroeconomic model, calibrated to the experience of past 12 years, assuming a normal distribution, constrained by zero lower bound	Policy interest rate; Output gap; CPI; core CPI	Norges Bank (2005, 2011) Alstadheim *et al* (2010)
Reserve Bank of New Zealand	Point estimates only			Reserve Bank of New Zealand (2011)
Peoples Bank of China	No quantitative measures. The presentation of the outlook is verbal.			PBOC (2011)

Appendix B: Data

The data used in this paper are available at www.rba.gov.au/publications/rdp/2012/2012-07.html, except for the proprietary forecasts from Consensus Economics, which are available via subscription at consensuseconomics.com.

B.1 Sample

Our main results use forecasts beginning in 1993:Q1, when the RBA began targeting an inflation rate of 2 to 3 per cent. Errors before this period were larger, and are arguably unrepresentative of those likely to be encountered under the existing policy framework. Furthermore, the forecasts were less detailed before 1993 – for example, horizons were shorter – which makes comparisons difficult. The latest quarter for which we calculate forecast errors is 2011:Q4.

B.2 Forecasts

Our dataset of forecasts has been put together over the years by a long series of RBA staff, including Dan Andrews, Andrea Brischetto, Adam Cagliarini, David Norman, Anna Park and Ivan Roberts. That data collection represented an enormous effort without which this paper would not have been possible. Previous public uses of these data include Stevens (2004, 2011) and Edey and Stone (2004).

We have spot-checked these data against original and other data sources, but have not sought to rebuild it. Previous compilers of the data made many choices regarding what to include and we largely follow their judgement. One consequence of that approach is that our dataset includes forecasts for different variables made at different times in the quarter. That inconsistency may matter for some questions but does not seem important for our purposes.

The RBA produces several different forecasts throughout a quarter, of which we use one. For the past few years, we use the detailed forecasts summarised in the *SMP*. Before these were available, our choices largely follow those made in previous internal RBA research, summarised in Table B1. The table lists main data sources, however, there are many exceptions, for example when forecasts are missing, when they are superseded by more authoritative sources, or when an alternative source has a longer horizon. Forecasts sometimes combined the general contour from the *SMP* with detail from other sources.

Table B1: Main Data Sources for Forecasts

Forecast date	Underlying inflation	CPI inflation	GDP	Unemployment
1991:Q1–2000:Q1	JEFG, *SMP* text and Board papers	JEFG and *SMP* text	JEFG	JEFG
2000:Q2–2004Q2	PDG and *SMP* text	PDG	JEFG	PDG
2004:Q3–2007:Q4	*SMP*	*SMP*	JEFG and Board papers	PDG
2008:Q1–present	*SMP*	*SMP*	*SMP*	*SMP*

Notes: 'JEFG' represents the forecast taken to the Joint Economic Forecasting Group meeting.

'*SMP* text' refers to the verbal description of the inflation outlook in the *Statement on Monetary Policy*.

'PDG' is the forecast prepared for the internal Policy Discussion Group in the middle month of the quarter.

'*SMP*' represents the detailed quarterly forecasts prepared for the *Statement on Monetary Policy.* The forecasts actually presented in the Statement, typically for year-ended growth rates, have less quarterly detail and precision.

'Board papers' represents the forecast prepared in the third month of the quarter for the next Board meeting.

The various data sources differ in terms of detail, intended audience and in other ways, but perhaps their most important difference concerns timing. The forecasts prepared for the Joint Economic Forecast Group (JEFG) were prepared toward the end of the quarter, following the release of the national accounts. Forecasts for the *Statement on Monetary Policy* (*SMP*) and the internal Policy Discussion Group (PDG) were prepared in the middle of the quarter, between the release of the CPI data and the national accounts.

At the beginning of our sample the forecast horizon varied between 3 and 6 quarters ahead. It has been gradually extended since then, recently varying between 9 and 11 quarters ahead. Because short-horizon forecasting began earlier, and because those forecasts overlap less, we have many more independent observations of short-horizon errors than we have at longer horizons. So we can talk more confidently about uncertainty regarding the next few quarters than we can about uncertainty regarding the next few years. Indeed, we believe we have too few errors at horizons beyond eight quarters for a reliable sample and we do not include these in our formal analysis.[18]

As mentioned in Section 2, the forecasts have often been conditioned on an assumption of unchanged interest rates. Alternative assumptions, such as choosing a path in line with market expectations, might give more accurate forecasts, however the potential improvement seems likely to be very small. That assessment is partly based on internal post-mortems on specific forecast errors, which have concluded that the constant interest rate assumption was not important. More generally, we regress RBA GDP errors on a measure of the yield curve (the difference between one-year and overnight Treasuries) at the time of the forecast. The coefficient is highly statistically significant ($p = 0.001$ at a 3-quarter-ahead horizon) and correctly signed, suggesting market assumptions on the path of interest rates could potentially reduce RBA errors. However, the effect on forecast accuracy is tiny: subtracting predicted values from the errors lowers the RMSE by only 5 per cent at a 3-quarter-ahead horizon or by 2 per cent 7 quarters ahead. This difference is barely discernible in charts, and would not qualitatively affect most of the comparisons we make in this paper. Even then, it overestimates the effect, given that it maximises accuracy after the event rather than using information available in real time.

The unimportance of interest rate assumptions is partly because the yield curve predicts short-term interest rates only slightly better than a random walk (see Guidolin and Thornton (2010), and references therein). The unimportance of interest rate assumptions also reflects a large part of the 'transmission mechanism' being projected separately. Market expectations are already priced in to the exchange rate, asset prices, and longer-term interest rates; a change in the short-term interest rate assumption need not affect the anticipated path of these variables. Similarly, many of the models used to construct the forecast (most obviously, univariate time series) implicitly embody historical interest rate behaviour. Goodhart (2009) discusses the role of interest rate conditioning assumptions.

---

18 For underlying inflation, we have 57 six-quarter-ahead errors, 36 seven-quarter-ahead errors, 22 eight-quarter-ahead errors and 10 nine-quarter-ahead errors.

As we discuss in the text, there have been many other changes to the forecasts over this period. For example, the forecasts are informed by models that evolve in response to new data and research. The RBA is continually learning, including from examination of past forecast errors.

Finally, before 2000, the source data for GDP forecasts are paper records showing quarterly changes. These records provide quite a limited history; sometimes only two or three quarters. This is insufficient to calculate near-term forecasts of year-ended or year-average changes. Accordingly, we splice the forecast data with real-time estimates from the Stone and Wardrop (2002) database. This is not an issue for inflation forecasts, where the source data are 4-quarter changes.

B.3 Outcomes

Defining actual outcomes or 'truth' involves judgement. The most recently published estimates are easily available, and reflect more information and better methods than earlier estimates. In that sense, they may be closer to ultimate 'truth'. However, they often do not correspond to the series being forecast because definitions have changed.

This is most importantly a problem for underlying inflation, which has changed definition several times, as discussed below. In practice, redefinitions of other variables have not been empirically important in our sample, though two examples may illustrate problems that could occur in the future. First, forecasts up to 2009 assumed that GDP included research and development as an intermediate input. However, the data published since then treat this expenditure as final output of capital. GDP under the later definition (and other changes introduced at the same time) is about 4 per cent larger than GDP measured according to the earlier definition, though average growth rates were little affected. As a second example, until the adoption of chain-weighting in 1998, the ABS periodically used to update the base period from which constant-price estimates were calculated, reducing the weight of goods with declining relative prices, such as computers. This gave rise to predictable revisions to both growth rates and levels. Even though these and other revisions were predictable, our data sources do not include forecasts of revisions to published data. Implicitly, the variable being forecast is an early version, not the 'final'.

Revisions that arise from changes in definitions are difficult to classify as a forecast error, or as an example of economic uncertainty. One obvious example is when the redefinition is known in advance, but not incorporated in the forecast or in ‘backcasts’. Another obvious example is when multiple forecasts are generated for different definitions, as occurs for inflation. Changing views on the merits of each definition should not be confused with the accuracy of each forecast. This problem can be reduced by using multiple measures of outcomes or by measuring outcomes with real-time data – that is, data published soon after the event.

Using real-time measures of outcomes has other advantages. First, if one can take the data available at the time of the forecast as given, then forecast errors apply both to changes and levels of the variable being forecast. This is a substantial simplification, especially for unemployment. Second, to perform statistical tests, it is helpful if the forecast errors are independent of each other, but that is not the case when subsequent redefinitions or benchmarking impose serial correlation.

However, there are also substantial costs to using real-time data. First, the initial estimates published by the ABS reflect a partial inclusion of source data, combined with various interpolations and extrapolations. Errors defined using these preliminary estimates may reflect skill in mimicking ABS internal procedures rather than an understanding of macroeconomic behaviour. Second, real-time data for some variables can be difficult to obtain.

The literature on forecast errors has generally assumed that the problems from changing definitions outweigh those from incomplete incorporation of source data. So outcomes are typically measured with near-real-time data. Examples include forecast evaluations conducted by the OECD (Vogel 2007), the IMF (Timmerman 2007), the US Federal Reserve (Reifschneider and Tulip 2007), and the ECB (ECB 2009). For a discussion see Robertson and Tallman (1998).

The different timing of data revisions in Australia leads us to a balance that is similar in principle, though slightly different in practice. For GDP, we use the fourth-published estimate; released four quarters after the relevant event. That permits inclusion of most source data, including one round of annual data, while minimising the effect of data redefinitions. For the unemployment rate, we use the estimate as of the forecast one quarter after the event.

For inflation, we have not judged the benefits of compiling a real-time dataset as being worth the costs and instead we use recent estimates. For the total CPI, that is unimportant, given that the data are not revised (forecasts are on a not seasonally adjusted basis). Instead of the headline CPI, some previous researchers have used the CPI excluding interest charges, in an attempt to correct for the constant interest rate assumption, but we did not find the rationale for this complication compelling. For underlying inflation we use recent estimates of various measures, and match these with the definition used at the time of the forecast. For recent forecasts, which are seasonally adjusted, 'truth' is the 15th series CPI through 2011:Q2, and 16th series estimates for 2011:Q3 and 2011:Q4. (Because the distribution of price changes is skewed, changes in seasonal adjustment have a noticeable effect on estimates of year-ended underlying inflation). With the exception of the change in seasonal adjustment just mentioned, our forecast errors do not reflect changes in the definition of underlying inflation, though they will include subsequent revisions to each measure. Table B2 shows the series we use for both forecasts and actual outcomes.

Croushore (2006) shows that the definition of truth can make a substantial difference to how forecasts are evaluated. However, for the questions raised in this paper, the definition of truth matters only slightly. As one might expect, using a definition closer to that used at the time of the forecast results in smaller errors. For example, as shown in Table 1, the RMSE of RBA 3-quarter-ahead forecasts of underlying inflation is 0.54 percentage points when outcomes are measured using the definitions used at the time of the forecast. These forecasts were more accurate than forecasts using the midpoint of the target at marginal significance levels (p = .06). However, if we measure these errors using the most recent data, the RMSE increases to 0.64, which is no longer statistically different from that of the target (p = .60). Similarly, early estimates of GDP growth have tended to be revised toward the historical mean (but not toward the forecast), so measuring actual GDP growth using recent data would result in a further deterioration in the explanatory power of the GDP forecasts.

Table B2: Measures Of Underlying Inflation

Date of forecast[a]	Measure
1991:Q1–1995:Q1	CPI excluding interest charges, fresh fruit and vegetables and automotive fuel.
1995:Q2–1998:Q2	Treasury's underlying rate.
1998:Q3–2005:Q3	Weighted median CPI, excluding interest and tax, not seasonally adjusted.
2005:Q4–2006:Q4	Trimmed mean. Outcomes are seasonally adjusted using 15th series seasonal factors.
2007:Q1–2009:Q2	Average of trimmed mean and weighted median. Outcomes are seasonally adjusted using 15th series seasonal factors.
2009:Q3–2011:Q4	Trimmed mean. Outcomes through 2011:Q2 are seasonally adjusted using 15th series seasonal factors. Outcomes for 2011:Q3 and 2011:Q4 use 16th series seasonal factors.

Notes: (a) Definitions vary with the date of the forecast, not the date of the event. For quarters following a change in definition, we carry two measures of truth: errors are measured using the old measure of truth for long-horizon forecasts and the new measure of truth for short-horizon forecasts.

Appendix C: Percentiles of Forecast Errors

The following tables show select percentiles of absolute forecast errors at different horizons for forecasts from 1993 to 2011. The estimates for the 70th and 90th percentiles are those shown in Figure 3; other columns are constructed in the same manner. Being a 19-year average, these estimates are unlikely to change quickly over time. So readers can easily construct confidence intervals about future forecasts.

Table C1: Underlying Inflation – Quantiles of Absolute Error Distribution
4-quarter change; 1993:Q1–2011:Q4

Horizon	**Percentile**				
	10%	30%	50%	70%	90%
0	0.03	0.09	0.15	0.21	0.33
1	0.02	0.10	0.22	0.34	0.54
2	0.04	0.14	0.27	0.42	0.70
3	0.08	0.14	0.30	0.50	0.83
4	0.05	0.16	0.29	0.58	1.14
5	0.05	0.14	0.34	0.61	1.27
6	0.09	0.19	0.36	0.65	1.32
7	0.07	0.24	0.48	0.79	1.43
8	0.11	0.29	0.47	0.77	1.19

Table C2: CPI Inflation – Quantiles of Absolute Error Distribution
4-quarter change; 1993:Q1–2011:Q4

Horizon	**Percentile**				
	10%	30%	50%	70%	90%
0	0.02	0.09	0.15	0.23	0.44
1	0.06	0.17	0.30	0.48	0.88
2	0.06	0.27	0.49	0.77	1.17
3	0.09	0.31	0.62	1.00	1.40
4	0.23	0.49	0.79	1.08	1.87
5	0.19	0.51	0.80	1.13	2.03
6	0.16	0.43	0.84	1.13	2.13
7	0.13	0.43	0.91	1.18	2.02
8	0.11	0.46	0.73	1.22	2.16

Table C3: GDP Growth – Quantiles of Absolute Error Distribution
4-quarter change; 1993:Q1–2011:Q4

Horizon	Percentile				
	10%	30%	50%	70%	90%
0	0.11	0.31	0.55	0.76	1.19
1	0.14	0.46	0.80	1.26	1.91
2	0.25	0.58	0.95	1.38	2.02
3	0.27	0.77	1.05	1.56	2.36
4	0.36	0.72	1.19	1.50	2.25
5	0.38	0.65	0.95	1.50	2.06
6	0.15	0.54	0.93	1.45	2.32
7	0.22	0.59	0.93	1.45	2.40
8	0.15	0.68	0.95	1.50	2.46

Table C4: Unemployment Rate – Quantiles of Absolute Error Distribution
1993:Q1–2011:Q4

Horizon	Percentile				
	10%	30%	50%	70%	90%
0	0.00	0.04	0.10	0.20	0.30
1	0.01	0.1	0.20	0.31	0.60
2	0.08	0.17	0.30	0.47	0.82
3	0.07	0.24	0.35	0.61	1.00
4	0.09	0.29	0.50	0.70	1.20
5	0.15	0.39	0.58	0.86	1.35
6	0.09	0.33	0.67	0.87	1.43
7	0.03	0.27	0.70	0.90	1.46
8	0.13	0.35	0.75	1.10	2.07

Appendix D: Comparisons with Errors from Null Alternatives

Table D1: Underlying Inflation Forecast RMSEs – RBA and Random Walk
4-quarter change; 1993:Q1–2011:Q4

Horizon (quarters ahead)	**RMSE** RBA	Random walk	Ratio	**Significance** p-value	R^2
0	0.20	0.28	0.72	.01	0.48
1	0.32	0.45	0.70	.01	0.51
2	0.42	0.60	0.69	.00	0.52
3	0.54	0.73	0.74	.02	0.46
4	0.61	0.80	0.76	.04	0.42
5	0.69	0.86	0.80	.10	0.36
6	0.73	0.94	0.78	.12	0.39
7	0.80	1.08	0.74	.15	0.46
8	0.75	1.23	0.61	.08	0.63

Table D2: Underlying Inflation Forecast RMSEs – RBA and Target Midpoint
4-quarter change; 1993:Q1–2011:Q4

Horizon (quarters ahead)	**RMSE** RBA	Target midpoint	Ratio	**Significance** p-value	R^2
0	0.20	0.67	0.30	.03	0.91
1	0.32	0.69	0.46	.05	0.79
2	0.42	0.71	0.59	.06	0.65
3	0.54	0.74	0.72	.06	0.48
4	0.61	0.77	0.79	.07	0.38
5	0.69	0.77	0.88	.29	0.22
6	0.73	0.81	0.91	.41	0.18
7	0.80	0.78	1.03	.87	–0.05
8	0.75	0.50	1.49	.08	–1.22

Table D3: CPI Inflation Forecast RMSEs – RBA and Random Walk

4-quarter change; 1993:Q1–2011:Q4

Horizon (quarters ahead)	**RMSE** RBA	Random walk	Ratio	**Significance** p-value	R^2
0	0.26	0.79	0.32	.00	0.89
1	0.48	1.24	0.39	.00	0.85
2	0.71	1.62	0.44	.00	0.81
3	0.89	1.90	0.47	.00	0.78
4	1.08	2.02	0.53	.00	0.71
5	1.20	2.09	0.57	.01	0.67
6	1.20	2.06	0.58	.02	0.66
7	1.27	2.19	0.58	.03	0.67
8	1.30	2.47	0.53	.03	0.72

Table D4: CPI Inflation Forecast RMSEs – RBA and Target Midpoint

4-quarter change; 1993:Q1–2011:Q4

Horizon (quarters ahead)	**RMSE** RBA	Target midpoint	Ratio	**Significance** p-value	R^2
0	0.26	1.39	0.19	.00	0.97
1	0.48	1.39	0.35	.00	0.88
2	0.71	1.39	0.51	.01	0.74
3	0.89	1.41	0.63	.04	0.60
4	1.08	1.42	0.76	.15	0.42
5	1.20	1.41	0.85	.38	0.27
6	1.20	1.39	0.87	.48	0.25
7	1.27	1.36	0.93	.78	0.13
8	1.30	1.40	0.93	.83	0.13

Table D5: GDP Growth Forecast RMSEs – RBA and Historical Mean

4-quarter change; 1993:Q1–2011:Q4

Horizon (quarters ahead)	RMSE RBA	RMSE Historical mean	RMSE Ratio	Significance p-value	R^2
0	0.79	1.25	0.63	.00	0.61
1	1.12	1.26	0.89	.27	0.21
2	1.30	1.27	1.02	.86	–0.04
3	1.44	1.28	1.13	.23	–0.28
4	1.40	1.26	1.11	.18	–0.23
5	1.32	1.26	1.05	.41	–0.10
6	1.36	1.32	1.03	.59	–0.06
7	1.39	1.39	1.00	.94	–0.01
8	1.4	1.51	0.92	.12	0.15

Table D6: Unemployment Rate Forecast RMSEs – RBA and Random Walk

1993:Q1–2011:Q4

Horizon (quarters ahead)	RMSE RBA	RMSE Random walk	RMSE Ratio	Significance p-value	R^2
0	0.17	0.24	0.71	.03	0.50
1	0.33	0.4	0.82	.18	0.32
2	0.49	0.55	0.9	.46	0.20
3	0.62	0.67	0.92	.63	0.15
4	0.76	0.78	0.98	.89	0.05
5	0.85	0.85	1.00	1.0	0.00
6	0.91	0.90	1.02	.92	–0.03
7	0.97	0.89	1.10	.69	–0.20
8	1.16	0.94	1.24	.53	–0.53

References

Alstadheim R, IW Bache, A Holmsen, J Maih and Ø Røisland (2010), ‘Monetary Policy Analysis in Practice’, Norges Bank *Staff Memo No 11*.

Atkeson A and LE Ohanian (2001), ‘Are Phillips Curves Useful for Forecasting Inflation?’, Federal Reserve Bank of Minneapolis *Quarterly Review*, 25(1) Winter, pp 2–11.

Bank of Canada (2009), ‘Methodology Used to Construct Fan Charts in the April 2009 *Monetary Policy Report*’, background note.

Bank of Canada (2011), *Monetary Policy Report*, October.

Bank of England (2011), *Inflation Report*, November.

Bank of Japan (2008), ‘Risk Balance Charts’, *Outlook for Economic Activity and Prices*, April, p 9.

Bank of Japan (2011), *Outlook for Economic Activity and Prices*, October.

Campbell SD (2007), ‘Macroeconomic Volatility, Predictability and Uncertainty in the Great Moderation: Evidence from the Survey of Professional Forecasters’, *Journal of Business & Economic Statistics*, 25(2), pp 191–200.

Croushore D (2006), ‘Forecasting with Real-Time Macroeconomic Data’, in G Elliott, C Granger and A Timmermann (eds), *Handbook of Economic Forecasting*, Vol 1, Handbooks in Economics, Elsevier, Amsterdam, pp 961–982.

Diebold FX and RS Mariano (1995), ‘Comparing Predictive Accuracy’, *Journal of Business & Economic Statistics*, 13(3), pp 253–263.

ECB (European Central Bank) (2009), ‘New Procedure for Constructing Eurosystem and ECB Staff Projection Ranges’. Available at <http://www.ecb.int/pub/pdf/other/newprocedureforprojections200912en.pdf>.

ECB (2011), 'ECB Staff Macroeconomic Projections for the Euro Area', March. Available at <http://www.ecb.int/pub/pdf/other/ecbstaffprojections201103en.pdf>.

Edey M and A Stone (2004), 'A Perspective on Monetary Policy Transparency and Communication', in C Kent and S Guttmann (eds), *The Future of Inflation Targeting*, Proceedings of a Conference, Reserve Bank of Australia, Sydney, pp 73–100.

Edge RM and RS Gurkaynak (2011), 'How Useful are Estimated DSGE Forecasts', Board of Governors of the Federal Reserve System Finance and Economics Discussion Series No 2011-11.

Elder R, G Kapetanios, T Taylor and T Yates (2005), 'Assessing the MPC's Fan Charts', Bank of England *Quarterly Bulletin*, 45(3) Autumn, pp 326–348.

FOMC (Federal Open Market Committee) (2007), 'Minutes of the Federal Open Market Committee, October 30-31 2007', Board of Governors of the Federal Reserve System, Washington DC.

Gerard H and K Nimark (2008), 'Combining Multivariate Density Forecasts Using Predictive Criteria', RBA Research Discussion Paper No 2008-02.

Goodhart CAE (2004), 'The Interaction between the Bank of England's Forecasts and Policy, and the Outturn', Financial Markets Group, London School of Economics and Political Science, Discussion Paper No 496.

Goodhart CAE (2009), 'The Interest Rate Conditioning Assumption', *International Journal of Central Banking*, 5(2), pp 85–108.

Guidolin M and DL Thornton (2010), 'Predictions of Short-Term Rates and the Expectations Hypothesis', Federal Reserve Bank of St. Louis Working Paper No 2010-013B, revised January 2011.

Haldane A (2012), 'Tails of the Unexpected', Speech presented at the School of Social and Political Science and Business School at the University of Edinburgh 'The Credit Crisis Five Years On/Unpacking the Crisis' Workshop, Edinburgh, 8–9 June.

Hayashi F (2000), *Econometrics*, Princeton University Press, Princeton.

Kahneman D, P Slovic and A Tversky (eds) (1982), *Judgment under Uncertainty: Heuristics and Biases*, Cambridge University Press, Cambridge.

Kearns J and P Lowe (2011), 'Australia's Prosperous 2000s: Housing and the Mining Boom', RBA Research Discussion Paper No 2011-07.

Murphy AH (1988), 'Skill Scores Based on the Mean Square Error and Their Relationships to the Correlation Coefficient', *Monthly Weather Review*, 116(12), pp 2417–2424.

Newey WK and KD West (1987), 'A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix', *Econometrica*, 55(3), pp 703–708.

Newey WK and KD West (1994), 'Automatic Lag Selection in Covariance Matrix Estimation', *The Review of Economic Studies*, 61(4), pp 631–653.

Norges Bank (2005), 'Uncertainty Surrounding Future Interest Rate Developments', *Inflation Report*, 3/2005, pp 19–21.

Norges Bank (2011), *Monetary Policy Report*, 3/2011.

PBOC (People's Bank of China) (2011), *China Monetary Policy Report*, Quarter Two.

Reifschneider D and P Tulip (2007), 'Gauging the Uncertainty of the Economic Outlook from Historical Forecasting Errors', Board of Governors of the Federal Reserve System Finance and Economics Discussion Series No 2007-60.

Reserve Bank of New Zealand (2011), *Monetary Policy Statement,* December.

Richards T and T Rosewall (2010), 'Measures of Underlying Inflation', RBA *Bulletin*, March, pp 7–12.

Robertson JC and EW Tallman (1998), 'Data Vintages and Measuring Forecast Model Performance', Federal Reserve Bank of Atlanta *Economic Review*, 83(4) Fourth Quarter, pp 4–20.

Romer CD and DH Romer (2000), 'Federal Reserve Information and the Behavior of Interest Rates', *The American Economic Review*, 90(3), pp 429–457.

Stevens G (1999), 'Economic Forecasting and Its Role in Making Monetary Policy', Address to the Economic Society of Australia Forecasting Conference, Melbourne, 19 August.

Stevens G (2004), 'Better than a Coin Toss? The Thankless Task of Economic Forecasting', Address to the Economic Society of Victoria and the Australian Industry Group 'Economic Focus – Australia's Prospects', Melbourne, 17 August.

Stevens G (2011), 'On the Use of Forecasts', Address to Australian Business Economists Annual Dinner, Sydney, 24 November.

Stone A and S Wardrop (2002), 'Real-Time National Accounts Data', RBA Research Discussion Paper No 2002-05.

Sveriges Riksbank (2007), 'Calculation Method for Uncertainty Bands', *Monetary Policy Report*, 2007:1, p 22.

Timmerman A (2007), 'An Evaluation of the *World Economic Outlook* Forecasts', *IMF Staff Papers*, 54(1), pp 1–33.

Tulip P (2009), 'Has the Economy Become More Predictable? Changes in Greenbook Forecast Accuracy', *Journal of Money, Credit and Banking*, 41(6), pp 1217–1231.

Vogel L (2007), 'How do the OECD Growth Projections for the G7 Economies Perform?: A Post-Mortem', OECD Economics Department Working Paper No 573.

West KD (1997), 'Another Heteroskedasticity- and Autocorrelation-Consistent Covariance Matrix Estimator', *Journal of Econometrics*, 76(1–2), pp 171–191.

Wikipedia (2012), 'Overconfidence Effect'. Available at <http://en.wikipedia.org/wiki/Overconfidence_effect>, accessed 10 January 2012.

RESEARCH DISCUSSION PAPERS

These papers can be downloaded from the Bank's website or a hard copy may be obtained by writing to:

Printing Administrator
Information Department
Reserve Bank of Australia
GPO Box 3947
SYDNEY NSW 2001

Enquiries:

Phone: +61 2 9551 9830
Facsimile: +61 2 9551 8033
Email: rbainfo@rba.gov.au
Website: http://www.rba.gov.au

2011-03	Urban Structure and Housing Prices: Some Evidence from Australian Cities	Mariano Kulish Anthony Richards Christian Gillitzer
2011-04	Assessing Some Models of the Impact of Financial Stress upon Business Cycles	Adrian Pagan Tim Robinson
2011-05	Terms of Trade Shocks: What are They and What Do They Do?	Jarkko Jääskelä Penelope Smith
2011-06	Does Equity Mispricing Influence Household and Firm Decisions?	James Hansen
2011-07	Australia's Prosperous 2000s: Housing and the Mining Boom	Jonathan Kearns Philip Lowe
2011-08	The Mining Industry: From Bust to Boom	Ellis Connolly David Orsmond
2012-01	Co-movement in Inflation	Hugo Gerard
2012-02	The Role of Credit Supply in the Australian Economy	David Jacobs Vanessa Rayner
2012-03	ATM Fees, Pricing and Consumer Behaviour: An Analysis of ATM Network Reform in Australia	Clare Noone
2012-04	Chinese Urban Residential Construction to 2040	Leon Berkelmans Hao Wang
2012-05	Payment System Design and Participant Operational Disruptions	Ashwin Clarke Jennifer Hancock
2012-06	The Impact of Payment System Design on Tiering Incentives	Robert Arculus Jennifer Hancock Greg Moran

RESERVE BANK OF AUSTRALIA