FULL EMPLOYMENT AND OUTPUT GAP UPDATE - JUNE QUARTER 2024[1]

This document details the impact of new data and model updates on the NAIRU and output gap estimates. The model average NAIRU estimate, which is SAMM's preferred model-based estimate, increased from 4.60 per cent at the time of August SMP to 4.67 per cent. Around -1bps reflects the new flow of data and around +8bps reflects technical adjustments since last round. The model average model estimate of the output gap has been revised from 0.8 per cent to 0.9 per cent in the June quarter. Around 0.2 percentage points reflects the new flow of data while -0.1 percentage points reflects technical adjustments.

This note focuses on mechanical changes in the model estimates due to data updates and technical revisions; a more thorough update to the central estimates as they relate to the economic outlook will be provided at the start of the November forecast round.

Full employment

Our suite of models indicates that the unemployment gap remains negative, roughly between -1 to -¼ per cent, and has narrowed slightly relative to August SMP (Graph 1). The underutilisation gap also remains negative and has narrowed slightly. These estimates indicate that the labour market is continuing to gradually move towards full employment.

The model average NAIRU estimate, which is SAMM's preferred model-based estimate, increased from 4.60 per cent at the time of August SMP to 4.67 per cent (Graph 2). The increase in the model-based estimate of the NAIRU reflects:

- Data revisions following August SMP and prior to June Quarter national accounts: +4bps
- Model enhancements to full employment suite (see *Technical adjustments*): +8bps
- New June quarter data: -5bps

If EA were to maintain the -15bps judgement from the August SMP round, the implied NAIRU assumption would be 4.52 per cent.

Graph 1

Model Estimates of Spare Capacity*

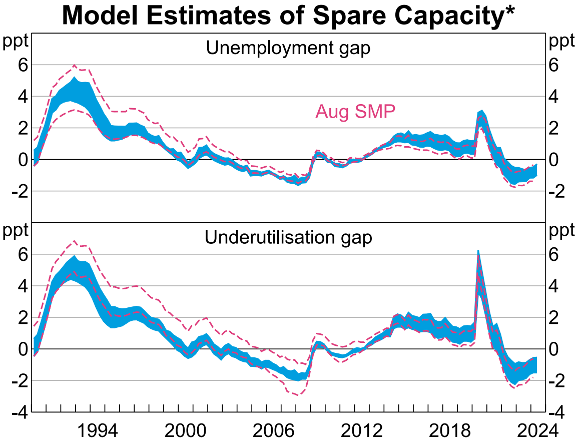

* Blue shaded region shows illustrative range of central gap estimates from a selection of models. Estimates are subject to substantial uncertainty, as well as revision due to data and model refinements.

Sources: ABS; RBA; Ruberl et al (2021).

Graph 2

Estimates of Spare Capacity*

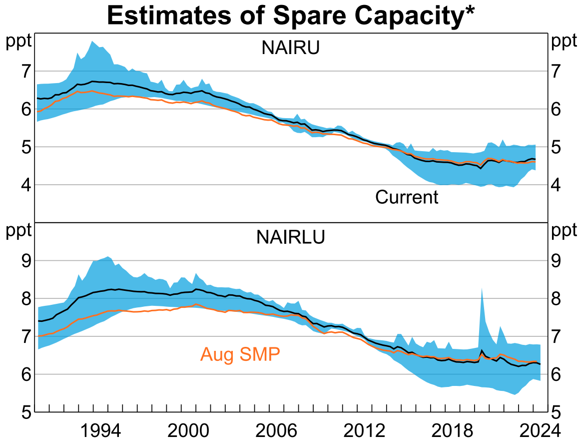

* Blue shaded region shows illustrative range of central estimates from a selection of models. Estimates are subject to substantial uncertainty, as well as revision due to data and model refinements.

Sources: ABS; OECD; RBA; Ruberl et al (2021).

Contributions of new data

New data in the June quarter contributed to a -5bps decrease in the NAIRU model average (Graph 3). The decrease was mostly driven by actual AENA outcomes coming in lower than predicted by the full employment models, contributing to a -7bps decrease in the NAIRU model average. This was slightly offset by a positive contribution from productivity.

1 We would like to thank and the rest of SAMM for help with the contents of this note. We'd also like to thank PWL and for support with the NAIRU technical adjustments.

The negative contribution of new data to the NAIRU model average in the June quarter follows three quarters of upward revisions (Graph 3). This is reflected in a broad-based increase in model estimates of the NAIRU over the past year (Graph 4). The exception is the Slack model, which has declined since early 2022. Unlike other models in the suite, the Slack model takes signal from alternative measures of labour market slack such as job ads, vacancies and labour utilisation.

Graph 3

Data Contributions to NAIRU Model Average

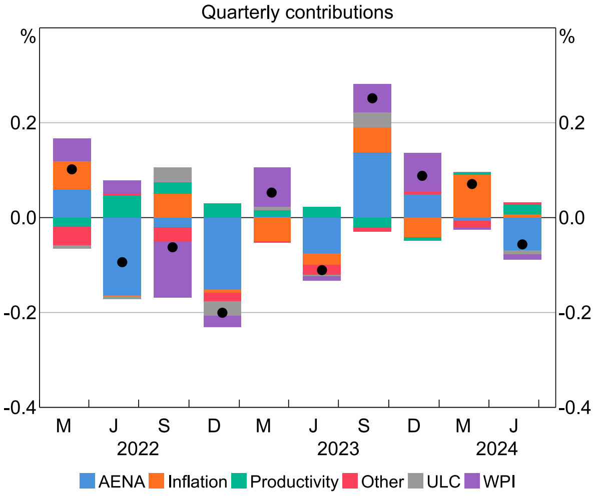

Source: RBA.

Graph 4

NAIRU Model Estimates

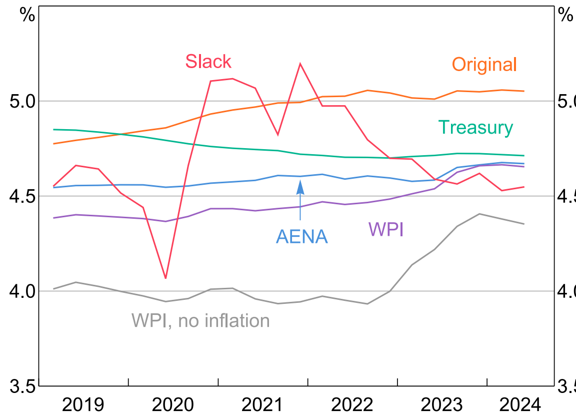

Source: RBA.

Technical adjustments

Following the August forecasting round, SAMM made several technical adjustments to the full employment suite (see [Technical Appendix](#) for further detail). This reflects that the full employment models are relatively new, and difficulties about how to treat volatile data due to the pandemic. Going forward, the frequency of technical adjustments is expected to decline.

The overall impact of technical adjustments was an increase in the model average NAIRU estimate by 8 basis points (Graph 5). The 'contributions' of each individual change to the model average NAIRU estimate depend on the order to which they are applied.

While some adjustments are technical in nature, others reflect changes in SAMM/PWL's assessment of the data. These are:

- FWC data: rather than take no signal from the September Quarter 2023 FWC decision, we are taking partial signal as per (2024). We are also taking partial signal from the September quarter 2022 FWC decision. This results in around a +11bps revision in the NAIRU model average.
- Serially correlated inflation errors: rather than reducing the signal that the NAIRU takes from inflation over the COVID period through a volatility break, we instead explicitly model positively correlated inflation errors. This reduces the signal taken from persistently high inflation over that period, reflecting our assessment that this mainly reflects cost-push shocks. This results in around a

Graph 5

Change in the NAIRU Assumption

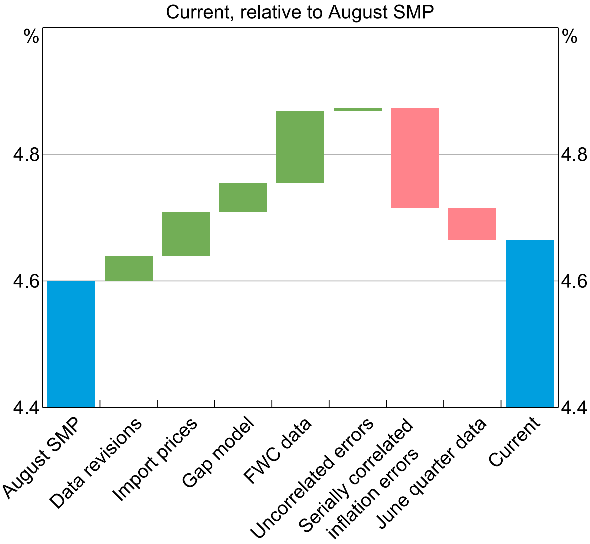

Source: RBA.

-15bps revision in the NAIRU model average.

Output gap

Our suite of model estimates suggest the output gap was still positive in June quarter 2024. The estimates range from about -0.4 per cent to 1.6 per cent of potential output (Graph 6). The range of estimates is wide, reflecting a large degree of uncertainty at this point in the cycle and different perspectives from different models in the suite.

The model average output gap estimate, which is SAMM's preferred model-based estimate, was revised 0.1 percentage points in June quarter 2024, from 0.8 per cent to 0.9 per cent (Graph 7). This revision largely reflects the effect of new data, while technical adjustments made to the models since August SMP had an offsetting effect.

Graph 6

Model Estimates of Output Gap*

Per cent of potential output

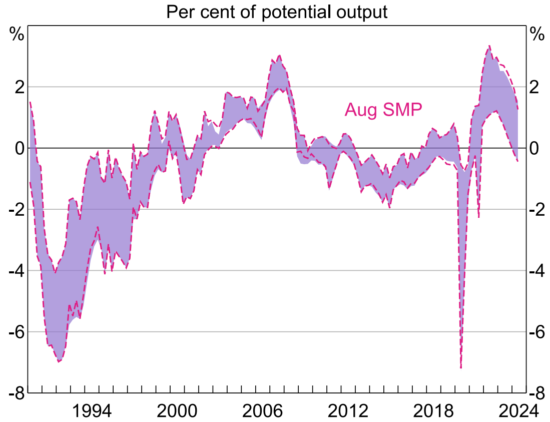

* Violet-shaded region shows illustrative range of central gap estimates from a selection of models encompassing different measures and definitions of the output gap; estimates are subject to substantial uncertainty, as well as revision due to data and model refinements.

Sources: OECD; RBA.

Graph 7

Revisions to Model Average*

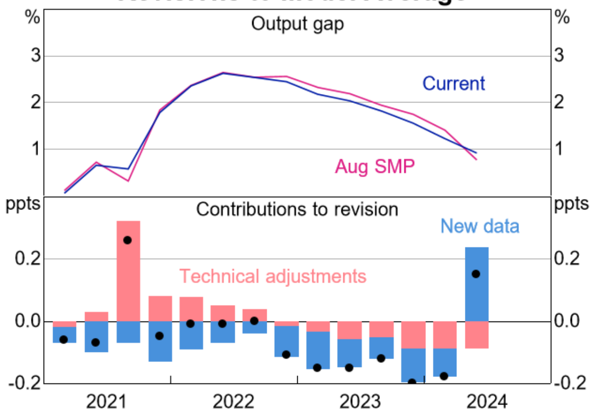

* Model average for August SMP reflects the average of the baseline specifications from the Production Function, SMOG and Joint-stars models. The current model average includes the addition of SMOG-PPE.

Source: RBA.

Contributions of new data

New data contributed to a 0.2 percentage point increase to the output gap model average for June quarter 2024. This revision largely reflects changes to estimates of the supply side of the economy as model estimates replace the 'house view' assumption of potential output growth that is used over the forecast horizon (Graph 8). All our model estimates of potential output growth are lower than the house view (Graph 9). While trend labour productivity is assumed to contribute to potential output growth in the house view, the production function model (PF) suggests trend productivity continues to detract from growth while estimates from SMOG-PPE implies it has a negligible contribution to growth.

Graph 8

Output Gap Revisions

Contributions to data revision

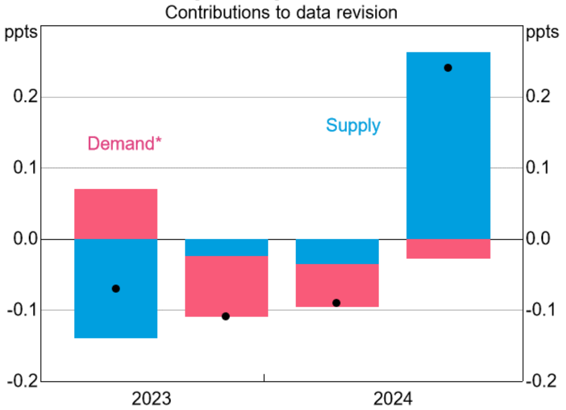

* We interpret non-farm GDP to reflect demand.

Source: RBA.

Graph 9

Contributions to Potential Output Growth

June Quarter 2024

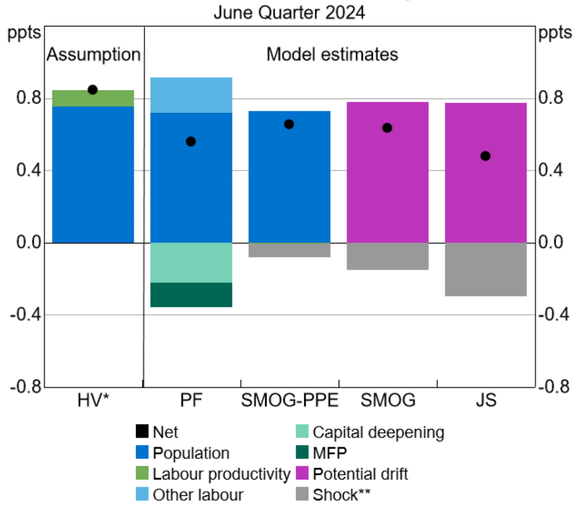

* To produce projections of the output gap we use a 'house view' of potential output growth. This view is largely based on EA's forecasting assumption for population and SAMM's assessment of trend labour productivity growth (with trend participation, trend average hours and changes in the NAIRU assumed to make a negligible contribution to potential growth in net terms).

** In the unobserved component models potential output in each quarter grows in a predictable way with the current potential drift, but is also affected by quarterly permanent shocks to the level of potential output.

Source: RBA.

Model adjustments

Following the August forecasting round, SAMM has made some technical adjustments to the model suite. In total, these changes result in a 0.1 percentage point downward revision to the model average output gap estimate in June quarter 2024. We have:

- Added SMOG-PPE (2024) to the suite and to the model average.
- Switched our baseline estimates in the Joint-Stars model to the 'ex-ante' specification of the model (we were previously using 'ex-post). The ex-ante specific deflates the nominal cash rate using forward-looking inflation expectations. This better captures how households and businesses make decisions (i.e. they are forward-looking).
- Removed forecasts from the production function model. Filtering over forecasts is a common strategy to try and improve the real-time reliability of filtered estimates. However recent analytical work (forthcoming) shows filtering over forecasts in the production function model does not improve the real-time reliability of our estimates.

SAMM
5 September 2024

A PRIMER ON THE RBA'S FULL EMPLOYMENT AND POTENTIAL OUTPUT MODELS[1]

Over the past year the Bank has overhauled its models for estimating full employment and potential output. We now have a suite of models for estimating both. Estimates from the suite are used to inform the Bank's assessment of spare capacity that is published in the *Statement on Monetary Policy* (Graphs 1 and 2). This is in line with the *Statement on the Conduct of Monetary Policy,* which requires the Bank to publish its assessment of full employment and potential output (Treasurer and Reserve Bank Board, 2023). Central estimates from the model suites also provide the foundations for the assumptions used in EA's forecasting framework and are particularly important for the inflation forecast.

Given recent model changes and the greater focus on full employment and potential output in the economic narrative, this note provides a high-level summary of the models that underpin our estimates and their limitations. The key takeaways from this note are:

1. All our models in the NAIRU and potential output suites are 'semi-structural' models that filter out cyclical variations and noise in the data from structural trends. Despite the name, **the models do not provide a structural explanation for changes in the estimate.** To help hone a coherent narrative, the estimates could be complemented with further evidence; this could involve routine analytical work, or more sophisticated approaches drawing on theory, other data or additional modelling.
2. Most of our models infer the NAIRU and potential output by using signal variables such as inflation and WPI. The models make predictions for these signal variables, and the errors (the difference between the actual outcome and prediction) inform how much to update estimates of the NAIRU or potential. These **model predictions should not be confused with past or present EA forecasts.** At times model prediction errors can be inconsistent with EA's forecast errors. This can present a communication challenge for the Bank.
3. **Estimates are uncertain and are subject to revision**. There is ongoing work to quantify the real-time reliability of our estimates. Nevertheless, the model estimates provide a useful foundation for forming the Bank's assessment of full employment and potential output.

Graph 1

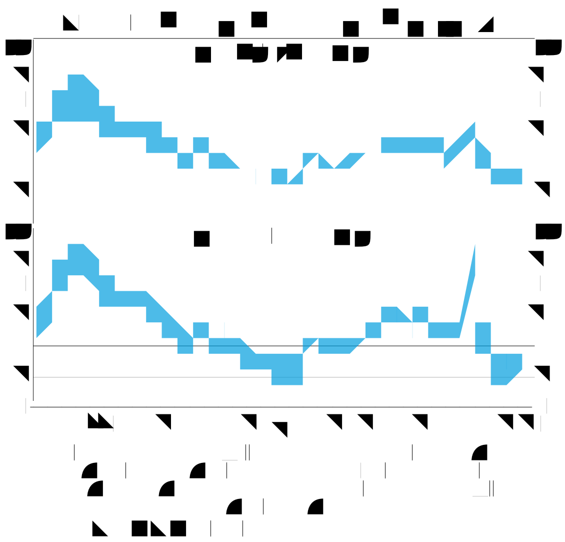

Graph 2

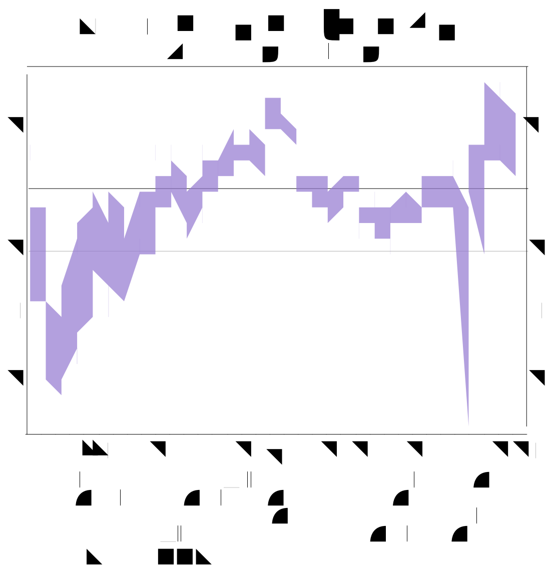

How do we think about the NAIRU and potential output?

There are multiple definitions of full employment and potential output. As an inflation-targeting central bank, we define full employment and potential as the levels of employment and output that are consistent with

---

1 This note draws on a large body of work undertaken by Economic Group over the past year to improve the Bank's estimate of full employment and potential output. Links to this body of work are in the Appendix. I'd like to thank

for their help on this note.

low and stable inflation in the medium term. This definition combines both prices and quantities, which means when the economy is operating at potential and full employment, there is both a balance between demand and supply in the labour and product markets, and inflation is also consistent with achieving the inflation target, at least in the medium to long-run. In the short run, the economy may be operating at potential but unemployment may not be at full employment or inflation may not be at target due to frictions in the goods and labour market and shocks in the economy.

How do we estimate the NAIRU and potential output?

The NAIRU and potential output are latent variables. They cannot be observed directly, but they can be inferred indirectly because they affect variables that are more easily observed. For example, if we observe high wages growth, we might infer the current unemployment rate is below the NAIRU. Models provide a formal framework for using observable data to infer the NAIRU and potential output. Most of the models in the potential output and NAIRU model suite are unobserved component models.[2] This type of model infers unobservable variables from observables variables by separating cyclical variations (i.e. the unemployment gap and output gap) and noise in the data from structural trends (i.e. the NAIRU and potential output).

Figure 1: Decomposition of Unemployment and GDP into Trend and Cycle

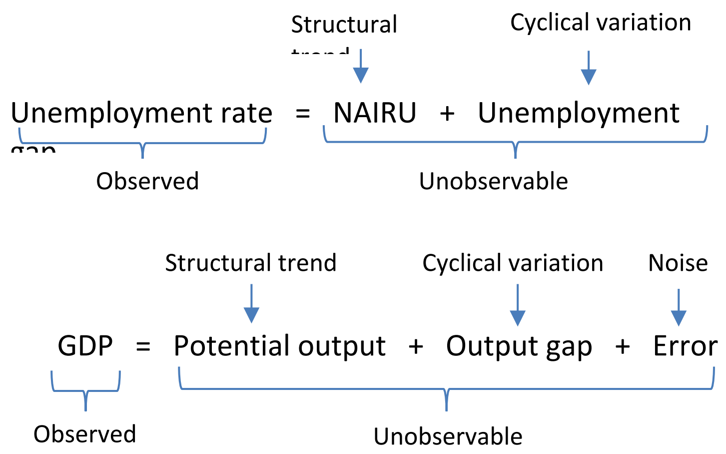

Observable variables (also called signal variables) differ between models. We choose signal variables based on a combination of economic theory and statistical properties. For example, the Phillips curve posits there is a relationship between inflation and spare capacity in the economy. This implies inflation may be a good observable variable to take signal from to infer cyclical variations in unemployment and output (i.e. the unemployment gap and output gap). Once we have the unemployment and output gap the model can then determine whether what is left over is noise or structural (persistent) trends in the data (i.e. the NAIRU and potential output). Signal variables for each model are summarised in Table 1.

Table 1: Signal variables

NAIRU models	Underlying inflation	ULC	AENA	WPI	Productivity	Underutil isation	Unemployment	Labour capacity utilisation	Jobs ads and Vacancies
Original	✓	✓							
AENA	✓		✓						
WPI	✓		✓	✓	✓				
WPI, no inflation			✓	✓	✓				
Treasury[(a)]	✓		✓						
Gap model	✓	✓		✓		✓	✓	✓	✓

2 We also have two production function models in the potential output suite. Like unobserved component models, production function models estimate unobserved or trend variables by filtering out the cyclical component from data. Rather than taking signal from price variables, production function models estimate potential output by focusing on the factors of production – that is labour, capital and productivity.

Potential output models	Underlying inflation	ULC	GDP	Part rate	Productivity	Hours worked	Unemployment /Employment	Population	Capital
SMOG	✓	✓	✓						
SMOG-PPE	✓	✓	✓		✓			✓	
Joint-stars	✓		✓	✓	✓	✓	✓	✓	✓
Production function[(b)]			✓	✓	✓	✓	✓	✓	✓
OECD production function[(b)]			✓				✓	✓	

(a) Based on Ruberl et al (2021).
(b) Production function models are a different class of model compared to unobserved component models. Strictly speaking, variables in these models are not signal variables, but they do inform estimates of potential output and the output gap. These models are included in the table above to give a more complete picture of what variables inform our estimates of potential output.
Sources: RBA

More specifically, to infer estimates of unobserved variables unobserved component models exploit the historical relationship between observed variables. For example, if the unemployment rate declines and inflation does not increase by as much as historical relationships would suggest, then the model interprets this to mean there is more spare capacity in the labour market, which implies a lower NAIRU, all else equal. To infer estimates the model makes predictions for each signal variable using historical data, and the prediction error informs the model how much to update estimates of the NAIRU or potential[3]. For example, if a NAIRU model predicted that inflation would come it at 0.7 per cent in the quarter, but the actual outcome was 0.9 per cent, then to explain the higher inflation the model would prefer a larger (negative) gap, and so the NAIRU is revised upwards accordingly. Typically, the estimate of the unobserved variable only changes a small amount relative to the size of the prediction errors. The rest of the prediction error would be attributed to the residual in the signal equation. Model predictions are based on so-called signal equations (see Figure 2). For inflation, this typically takes the form of a Phillips Curve; for unemployment this typically takes the form of an Okun's law relationship.

Figure 2: Example of a Signal Equation

Unobserved variable of interest

Signal variable

Noise

$$\pi_t = \boldsymbol{BX}_t + \lambda gap_t + \epsilon_t^{\pi}$$

Vector of observable explanatory variables

What causes changes in the NAIRU and potential estimates?

When there is new data the model updates its estimate of the NAIRU and potential. How much the estimate is updated by is determined by the prediction error and model parameters, as discussed above. New data typically leads to past estimates being revised. In the case of smoothed estimates (or two-sided estimates), past estimates will be updated to take into account the latest data. The latest estimates of smoothed estimates are especially vulnerable to large updates as there is no future data to help refine the estimates. This problem is known as the end point problem. New data may also lead to different parameters in the model which can affect past estimates, but this is generally a minor source of revision forthcoming) . Finally, revisions to historical data can change model estimates and possibly model parameters.

Model predictions of signal variables like inflation should not be confused with past or present EA forecasts. They are unrelated to forecasts produced by EA and can at times give a different read to EA's forecasts. For

3 See (2023) for more information on state space models (of which unobserved component models are a specific type).

example, Graph 3 compares model predictions from the WPI NAIRU model from August SMP to PWL nowcasts (excluding judgement) from different SMP rounds. While PWL's nowcasts pointed to a gradual increase in WPI growth over 2023, the WPI NAIRU model prediction implied steady quarterly growth around 0.8 per cent.

Graph 3

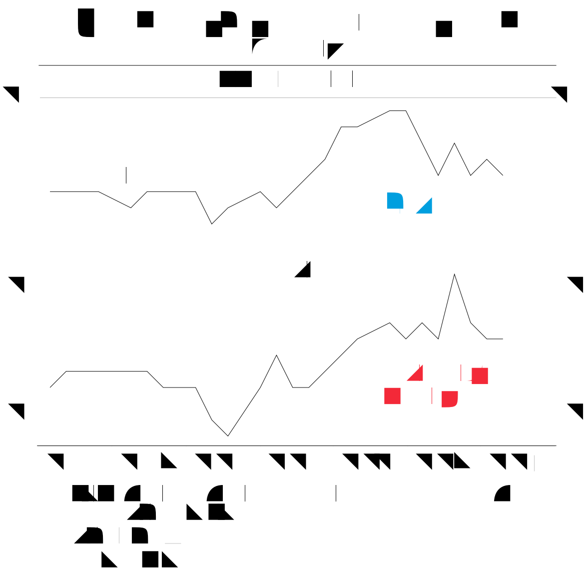

There are many reasons why model predictions diverge from EA's forecasts. First, EA's nowcasts are generally based on a suite of models which incorporate different methods and data to inform the nowcast. On the other hand, model predictions are from a single equation in the model. Second, even if the same modelling framework is used by EA and in the NAIRU model (e.g. a Phillips Curve), these models may have different specifications as the models are used for different purposes. Finally, EA's nowcasts incorporate judgement, while the model predictions do not. For example, in the September quarter 2023 PWL nowcasted strong wages growth in anticipation of an usually strong FWC decision on award wages, but the NAIRU models did not incorporate that information. Difference between model predictions and EA's forecasts can present a communication challenge for the Bank.

Limitations of the models

There are limitations to using unobserved component models to estimate the NAIRU and potential output. First, there is considerable uncertainty around the model estimates. Sources of uncertainty include uncertainty from the data (measurement error, statistical noise), parameter uncertainty, model uncertainty and filtering uncertainty which is inherent when trying to inferred unobservable variables through the movement of observable variables. Standard error bands capture the uncertainty around the model estimates (see for example Graph 2 in (2023)). Although our standard error bands are wide, they are not so wide that we can't say anything useful about the NAIRU or potential output. For example, we can provides probabilities of the labour market being above or below full employment, as in (2024).

Second, real time reliability of our model estimates can be a concern to the extent that new data leads to large changes in our real time estimates. This can pose a challenge for how the Bank communicates its inflation forecasts and its assessment of spare capacity. There is ongoing work to evaluate the real time reliability of our estimates. For example, work by (forthcoming) estimates the expected change in the current estimate of the NAIRU from the 'WPI' model in the NAIRU suite due to an additional quarter of data. It shows the expected change in the current estimate has a standard deviation of 15 basis points, and 95 per cent of revisions are expected to be smaller than 30 basis points in absolute value. These estimates account for both the arrival of new data and parameter re-estimation.

Finally, the models do not provide a structural interpretation of changes in the NAIRU or potential output. They do not consider the underlying structural factors that determine the long-run productive capacity of the economy – for example how much labour households are willing to supply and how much labour businesses are willing to employ. Given the considerable uncertainty with estimating the NAIRU and potential output, additional evidence that provides a structural interpretation could help corroborate our estimates and communicate them to a wider audience. However, this would require a structural model (such as a SVAR or a DSGE) that is equipped to produce these insights, or analysis aimed at teasing out the underlying mechanisms. While the Bank has a DSGE model, it is not currently well suited to produce estimates of the NAIRU and potential output and would require considerable time and resourcing to develop this capability. Aside from structural interpretations, additional evidence can be brought to bear on how conditions in the labour market and broader economy currently compare to full employment and potential output, thereby indirectly providing corroborating evidence for our estimates.

Structural Analysis & Macroeconomic Modelling
Economic Analysis Department
9 October 2024

HOW MUCH SHOULD WE EXPECT OUR NAIRU ESTIMATES TO BE REVISED II: THE ROLE OF PARAMETER ESTIMATION[1]

This note quantifies the expected magnitude of quarterly changes in our NAIRU estimates, extending previous analysis to account for estimation of model parameters, like the slope of the Phillips curve. Using Monte Carlo simulation, I approximate the distribution of NAIRU estimates in 2024Q3 conditional on information available in 2024Q2 and hence construct a conditional distribution for the quarterly change in the estimate. Focusing on one model in the suite, the expected change in the current NAIRU estimate has a standard deviation of 15 basis points, and 95 per cent of changes are expected to be smaller than 30 basis points in absolute value. The bulk of this variation reflects the 'direct' influence of incoming data (i.e. it would occur absent parameter re-estimation). This suggests that the choice of how frequently to re-estimate the parameters of our NAIRU models would not have material implications for the variability of our NAIRU estimates. Also, it is unlikely that we would have to explain large parameter-driven changes in our estimates to decisionmakers or the public.

Introduction

Estimates of the non-accelerating inflation rate of unemployment (NAIRU) are important inputs into our assessment of full employment and inflation forecasts (e.g. Ballantyne, Sharma and Taylor 2024; 2024). Better understanding the properties of these estimates is important for managing their role in the policy process and is an ongoing area of work for EC. In (2024b), I quantified the expected magnitude of changes in our NAIRU estimates ('revisions'), which I referred to as 'revisability'.[2] I focused on revisions purely due to incoming data, holding model parameters fixed. By abstracting from other sources of revision, including parameter re-estimation and data revisions, I provided a lower bound on revisability.

In this note, I extend my earlier work to account for parameter re-estimation. I provide a more comprehensive indication about the likely size of NAIRU revisions and assess the relative importance of parameter re-estimation in driving revisions. This analysis can be useful for considering the implications of less-frequent re-estimation of our NAIRU models or when planning how to communicate about the NAIRU estimates; for example, it could be difficult to explain the reasons for changes in our NAIRU estimates to decisionmakers or the public if these are largely driven by changes in model parameters, whereas changes in estimates due to inflation prediction errors should be easier to explain.[3]

I use Monte Carlo simulation to approximate the distribution of the revision between 2024Q2 and 2024Q3, conditional on information available in 2024Q2.

Imagine that we have estimated one of our NAIRU models using data up to 2024Q2 and have computed an estimate of the NAIRU. When we receive data for 2024Q3, we will re-estimate the model and produce an estimate of the NAIRU in 2024Q3. The properties of the method that we use to estimate the NAIRU (the Kalman filter) imply that the expected value of the 2024Q3 estimate will be the same as the 2024Q2 estimate, so the revision will be zero in expectation. Of course, the actual revision will be non-zero. How different might the 2024Q3 estimate be to the 2024Q2 estimate?

To answer this question, I use Monte Carlo simulation to approximate the distribution of the input data (e.g. inflation) in 2024Q3, accounting for uncertainty about the parameters governing the data-generating process (DGP). At each simulated data sample, I re-estimate the NAIRU model and compute the 2024Q3 NAIRU estimate. This generates a distribution of revisions in the NAIRU estimate. These revisions can be thought of as comprising two components: 1) the 'direct' effect of incoming data on the NAIRU estimate given fixed model parameters, which is the concept considered in (2024b); and 2) the 'indirect' effect of new data on the NAIRU estimate via changes in the model's parameter estimates, which influence how the model infers movements in the NAIRU from the data. See the Appendix for details about the methodology.

---

1 Thanks to at the SAMM Wednesday Meeting and ER Thursday Coffee for useful input. Replication files here: D24/308747.

2 The previous analysis made a distinction between revisions and changes in estimates. If $U^*_{t|s}$ is the estimate of the NAIRU in time t given information available at time s, the revision was defined as $U^*_{t|t+1} - U^*_{t|t}$ and the change was defined as $U^*_{t+1|t+1} - U^*_{t|t}$. For ease of exposition, in the current note I refer to the latter concept as a revision.

3 Differences between model-implied prediction errors and SMP forecast errors could also raise issues, as discussed in (2024).

I focus on a single model in the NAIRU suite – the 'WPI' model from and (2023).[4] The exercise is hence narrower in scope than in (2024b), where I explored how revisability varies over time and across models, as well as revisability of the model-average estimate (*et al* 2024). I keep the scope narrow for a few reasons. First, the Monte Carlo exercise is computationally intensive. Second, I have encountered numerical convergence issues when estimating the model on earlier sample periods. Finally, it is not obvious what to assume about the DGP when interest is in the behaviour of a model-average estimate. Given the focus on a single model, the results below should be viewed as indicative.

The revision has a standard deviation of 15 basis points and 95 per cent of revisions are expected to be smaller than 30 basis points in absolute value.

Graph 1 plots the conditional distribution of the NAIRU revision between 2024Q2 and 2024Q3 when holding the parameters fixed at their 2024Q2 values and when re-estimating them. When re-estimating the parameters, the standard deviation of the revision is 15 basis points and 95 per cent of revisions are smaller than 30 basis points in absolute value. As expected, the size of revisions is larger than implied by the analysis in (2024b), which abstracted from parameter estimation.[5]

Graph 1

Conditional Distribution of Revision in NAIRU Estimate*

Between 2024Q2 and 2024Q3

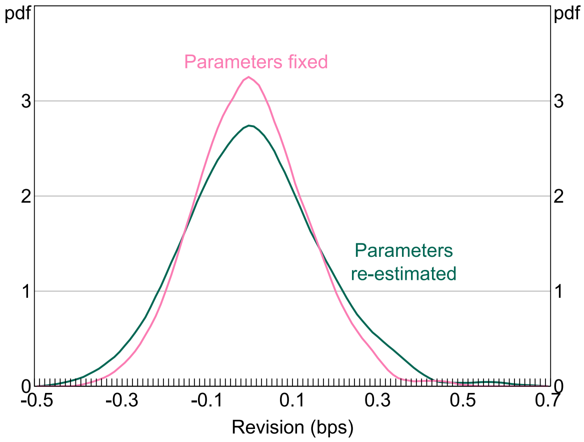

* Conditional on data available in 2024Q2; kernel density estimate based on Epanechnikov kernel.

Graph 2

NAIRU Estimates – Uncertainty and Revisability*

Per cent of labour force, quarterly

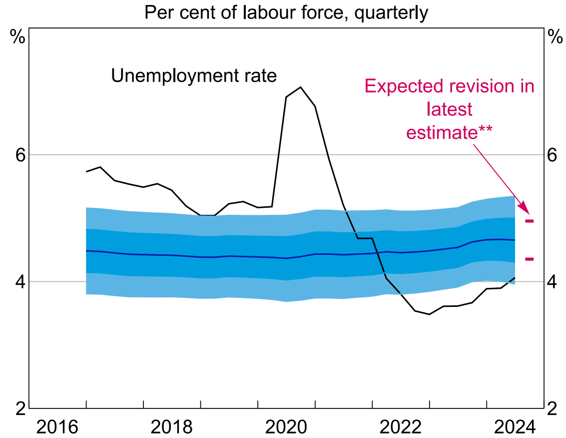

* Shaded areas represent one and two standard error bands around smoothed estimates, accounting for filtering and parameter uncertainty; based on 'WPI' model.

** Two standard deviation band for the one-quarter revision in latest estimate.

Sources: ABS; author's calculations.

Graph 2 illustrates the degree of revisability relative to statistical uncertainty, as measured by standard error bands.[6] A two standard error band around the latest NAIRU estimate is about 1.4 percentage points wide, while a two standard deviation band for the expected revision in the NAIRU estimate (a 'revisability band') is 60 basis points wide. The difference in widths reflects differences in the nature of the bands; the standard error bands represent uncertainty about the *true value* of the NAIRU, whereas the revisability bands represent how much we might expect the *estimate* of the NAIRU to change. One way to provide some economic context for the size of revisions is to consider what a given revision in the NAIRU implies about forecasts of inflation; in MARTIN a 30 basis point increase in the NAIRU increases inflation by about 23 basis points after two years.

Most of the expected variation in revisions reflects the direct effect of incoming data.

The bulk of the expected variation in the revision is due to the 'direct' effect of input data, rather than parameter re-estimation; while there is greater dispersion in the distribution of revisions when re-estimating the parameters compared with when the parameters are fixed, the difference in dispersion is not large (see Graph 1). More precisely, the conditional variance of the revision can be decomposed into three components: the 'direct' effect of incoming data, holding parameters fixed; the 'indirect' effect of parameter

4 The model incorporates correlated measurement errors (2024a) and the adjustments in (2024).

5 The results are not directly comparable to the results in (2024b), since the model used here incorporates recent adjustments made by SAMM. Applying the approach in (2024b) to the current model yields a standard deviation of 11.7 basis points.

6 The standard error bands capture both 'filtering uncertainty' and 'parameter uncertainty', as in (2023a).

re-estimation; and a covariance term. The 'direct' effect accounts for about 70 per cent of the conditional variance or, alternatively, 12 basis points of the total 15 basis point standard deviation.[7]

The small contribution of parameter re-estimation suggests that the choice of how frequently to re-estimate the parameters of our NAIRU models would not have material implications for the revisability of our NAIRU estimates. In contrast, as shown in (2024b), the frequency at which we update the NAIRU *estimates* themselves can have large implications for revisability; updating the estimates at an annual frequency roughly doubles the standard deviation of revisions relative to quarterly updating, making large data-driven revisions more likely. Additionally, the small contribution of parameter re-estimation means it is unlikely that we would have to explain large parameter-driven revisions in our estimates to decisionmakers or the public (which could be challenging).

Declining parameter uncertainty over time appears to have contributed to a decline in revisability.

This exercise captures potential time-variation in revisability arising from two sources. First, explicit time-variation in some model parameters (e.g. structural breaks in variances) can change how incoming data are expected to influence the NAIRU estimate. Second, assuming a stable DGP, model parameters should become more precisely estimated over time. Consequently, as the sample size grows, incoming data should have a smaller effect on parameter estimates and thus a smaller 'indirect' effect on the NAIRU estimates.

To briefly explore how revisability has varied over time, I re-run the Monte Carlo exercise for an earlier period. Specifically, I approximate the distribution of the revision between 2017Q4 and 2018Q1, conditional on information available in 2017Q4 (based on the current vintage of data).[8] The conditional standard deviation of the revision in the earlier period was 22 basis points, which is larger than currently (15 basis points). This partly reflects greater uncertainty about model parameters in the earlier period; to give a crude sense of this, around two-thirds of parameters have a smaller standard error now than in 2017Q4. These results tentatively suggest that a real-time estimation exercise, where the model is recursively re-estimated to generate a sequence of realised revisions, could overstate how susceptible our current NAIRU estimates are to revision.

Conclusion

Variability in our NAIRU estimates of the degree estimated here may pose challenges for how we communicate about the role of these estimates in forecasting inflation and assessing full employment. Exploring strategies for navigating these challenges is an ongoing area of work for EC. For example, will assess options for updating the NAIRU assumption. On the other hand, the results in this note suggest that parameter re-estimation is unlikely to drive material revisions in our NAIRU estimates, so it is unlikely that we would have to explain large parameter-driven revisions in our estimates to decisionmakers or the public (which could be difficult).

I have focused on a single model in the suite, though EA's current 'preferred' estimate is a model average. The results should therefore be viewed as indicative; the model-average estimate would probably be less prone to revision.[9] On the other hand, I have abstracted from data revisions, which can be substantial for some input series. Capturing the role of data revisions would require real-time estimation exercises.[10]

Economic Research Department
15 October 2024

7 Part of the 'direct' data effect reflects uncertainty about the parameters governing the DGP, which increases dispersion in the input data relative to the case where the parameters are fixed at the MLE (the assumption in (2024b)). In the current exercise, the 'direct' data effect is 12.4 basis points, compared with 11.7 basis points when ignoring parameter uncertainty.

8 The choice of period is somewhat arbitrary; I encountered numerical problems when considering earlier periods.

9 (2024b) examined the behaviour of quarterly changes in the model-average estimate when recursively estimating the NAIRU but without re-estimating model parameters. The sample standard deviation of quarterly changes in the model-average estimate was about three-quarters that of the model in this note. As a back-of-the-envelope calculation, scaling down the conditional standard deviation of revisions obtained in the current exercise (15 basis points) suggests that the conditional standard deviation of revisions in the model-average estimate would be about 11 basis points.

10 In ongoing work PWL and SAMM have been investigating the properties of 'star variable' estimates using recursive estimation exercises with real-time data. This also produces 'vintages' of star variable estimates, which are useful in forecast evaluations.

Appendix

Let $U^*_{t|t}(\widehat{\boldsymbol{\theta}}(\boldsymbol{Y}^t)) = E(U^*_t|\boldsymbol{Y}^t; \widehat{\boldsymbol{\theta}}(\boldsymbol{Y}^t))$ be the 'filtered' estimate of the NAIRU at time t, U^*_t, which is the estimate conditional on the data available up to time t, $\boldsymbol{Y}^t = (\boldsymbol{y}'_1, \dots, \boldsymbol{y}'_t)'$. Given the data and a value of the model parameters $\boldsymbol{\theta}$, the filtered estimate is computed using the Kalman filter.[11] I write $\widehat{\boldsymbol{\theta}}(\boldsymbol{Y}^t)$ to represent the maximum likelihood estimate (MLE) of $\boldsymbol{\theta}$ conditional on data up to time t. So $U^*_{t|t}(\widehat{\boldsymbol{\theta}}(\boldsymbol{Y}^t))$ is the filtered estimate of the NAIRU when the model parameters are set equal to the MLE given time-t information.

I quantify the expected magnitude of the quarterly change in the filtered estimate, which I refer to as 'revisions'. Taking into account that new data will change the MLE of the parameters, the revision is

$$\Delta_{t:t+1} = U^*_{t+1|t+1}(\widehat{\boldsymbol{\theta}}(\boldsymbol{Y}^{t+1})) - U^*_{t|t}\left(\widehat{\boldsymbol{\theta}}(\boldsymbol{Y}^t)\right).$$

Consider decomposing the revision as:

$$\Delta_{t:t+1} = \left[U^*_{t+1|t+1}\left(\widehat{\boldsymbol{\theta}}(\boldsymbol{Y}^{t+1})\right) - U^*_{t+1|t+1}\left(\widehat{\boldsymbol{\theta}}(\boldsymbol{Y}^t)\right)\right] + \left[U^*_{t+1|t+1}\left(\widehat{\boldsymbol{\theta}}(\boldsymbol{Y}^t)\right) - U^*_{t|t}\left(\widehat{\boldsymbol{\theta}}(\boldsymbol{Y}^t)\right)\right].$$

The first term captures how re-estimation of model parameters changes the NAIRU estimate in period $t+1$, while the second term captures how incoming data change the estimate between period t and $t+1$, holding parameters fixed.

The conditional variance of the revision, $\mathrm{Var}_t(\Delta_{t:t+1})$, is

$$\begin{aligned}\mathrm{Var}_t&\left[U^*_{t+1|t+1}(\widehat{\boldsymbol{\theta}}(\boldsymbol{Y}^{t+1})) - U^*_{t+1|t+1}\left(\widehat{\boldsymbol{\theta}}(\boldsymbol{Y}^t)\right)\right] + \mathrm{Var}_t\left[U^*_{t+1|t+1}\left(\widehat{\boldsymbol{\theta}}(\boldsymbol{Y}^t)\right) - U^*_{t|t}\left(\widehat{\boldsymbol{\theta}}(\boldsymbol{Y}^t)\right)\right] \\ &+ 2\mathrm{Cov}_t\left[\left[U^*_{t+1|t+1}(\widehat{\boldsymbol{\theta}}(\boldsymbol{Y}^{t+1})) - U^*_{t+1|t+1}\left(\widehat{\boldsymbol{\theta}}(\boldsymbol{Y}^t)\right)\right]\left[U^*_{t+1|t+1}\left(\widehat{\boldsymbol{\theta}}(\boldsymbol{Y}^t)\right) - U^*_{t|t}\left(\widehat{\boldsymbol{\theta}}(\boldsymbol{Y}^t)\right)\right]\right].\end{aligned}$$

This conditional variance has three terms. The first term can be thought of as the 'indirect' contribution of parameter re-estimation, the second is the 'direct' contribution of incoming data (holding parameter estimates fixed) and the third is a covariance term. (2024b) quantified the expected magnitude of revisions by computing the second term under the assumption that the parameters governing the DGP are known and equal to the MLE. That analysis therefore abstracted from the effect of parameter re-estimation and ignored uncertainty about the 'true' model parameters, and so could be interpreted as providing a lower bound on the true degree of revisability. In large samples, parameter uncertainty will vanish and the revisability measure in (2024b) would hence provide an accurate guide about the expected magnitude of actual revisions (abstracting from data revisions).

To approximate the conditional distribution of revisions ($\Delta_{t:t+1}$), I use Monte Carlo methods. The approach involves two key assumptions. First, I assume that the usual asymptotic normal approximation of the sampling distribution of the MLE is approximately equivalent to a Bayesian posterior distribution for the model parameters.[12] This means I can capture uncertainty about the parameters governing the true DGP by randomly drawing parameter vectors from a normal distribution centred at the MLE. Accounting for uncertainty about the parameters governing the DGP means that the Monte Carlo distribution of input data has more dispersion than if I were to condition on a particular value of the model parameters (e.g. the MLE). Second, I assume that model shocks are normally distributed.[13] This allows me to approximate the conditional distribution of $\Delta_{t:t+1}$ by drawing $\boldsymbol{y}_{t+1}$ from a normal distribution with mean and variance given by outputs from the Kalman filter, evaluated at each draw of the parameters.[14]

The following algorithm details how I approximate the conditional distribution of revisions:

11 See (2023b) for an overview of state-space models and methods, including the Kalman filter, or Hamilton (1994) for a textbook treatment.

12 This approach is motivated by the Bernstein von Mises theorem, which states that (under some conditions) the sampling distribution of the MLE centred around the true parameter value asymptotically coincides with the posterior distribution centred around the MLE. The same idea underlies the approach that we use to approximate uncertainty about the NAIRU (2023a).

13 Normality of shocks is also assumed in (2024) when computing the probability of spare capacity in the labour market.

14 Sequentially drawing parameters and then data given the draws of the parameters means that I am drawing from an approximation of a Bayesian posterior predictive distribution for the data.

1) Draw $\boldsymbol{\theta} \sim N(\widehat{\boldsymbol{\theta}}(\boldsymbol{Y}^t), \boldsymbol{\Omega}(\boldsymbol{Y}^t))$, where $\boldsymbol{\Omega}(\boldsymbol{Y}^t)$ is the variance-covariance matrix of the MLE obtained using data up to time t.[15]
2) Draw $\widetilde{\boldsymbol{y}}_{t+1} \sim N(\widehat{\boldsymbol{y}}_{t+1|t}(\boldsymbol{\theta}), \boldsymbol{F}_{t+1}(\boldsymbol{\theta}))$, where: $\widehat{\boldsymbol{y}}_{t+1|t}(\boldsymbol{\theta})$ is the expected value of $\boldsymbol{y}_{t+1}$ given data available at time t, conditional on the model parameters being equal to the draw of $\boldsymbol{\theta}$ from Step 1); and $F_{t+1}(\boldsymbol{\theta})$ is the associated prediction error variance. Both quantities can be obtained using the Kalman filter.
3) Augment the history of input data up to time t with the simulated data in time $t+1$ and let $\widetilde{\boldsymbol{Y}}^{t+1} = (\widetilde{\boldsymbol{y}}'_{t+1}, \boldsymbol{y}'_t, \ldots, \boldsymbol{y}'_1)'$. Evaluate:
    a. $U^*_{t+1|t+1}\left(\widehat{\boldsymbol{\theta}}(\boldsymbol{Y}^t)\right)$ by applying the Kalman filter to $\widetilde{\boldsymbol{Y}}^{t+1}$ with model parameters $\widehat{\boldsymbol{\theta}}(\boldsymbol{Y}^t)$.
    b. $U^*_{t+1|t+1}(\widehat{\boldsymbol{\theta}}(\widetilde{\boldsymbol{Y}}^{t+1}))$ by applying the Kalman filter to $\widetilde{\boldsymbol{Y}}^{t+1}$ with re-estimated model parameters $\widehat{\boldsymbol{\theta}}(\widetilde{\boldsymbol{Y}}^{t+1})$.
4) Repeat Steps 1-3 $1{,}000$ times and approximate conditional variances or standard deviations by Monte Carlo sample analogues.

This exercise generates a model-implied distribution of potential revisions. An alternative would be to run a real-time estimation exercise, where the model is recursively estimated each quarter using available data. The sample standard deviation of revisions in such an exercise tells us about the average magnitude of revisions over time, but it could give a misleading sense about how prone our *current* estimates are to revision. That said, real-time estimation exercises require weaker assumptions and can capture other sources of revisions, such as revisions in input data. The two different approaches can be viewed as complementary.

The NAIRU models include exogenous (pre-determined) variables on the right-hand side of the Phillips curve measurement equations, such as the unemployment rate and inflation expectations. Accounting for estimation of the coefficients on these variables requires making assumptions about these variables' values in time $t+1$. I hold these variables constant in time $t+1$ at their time-t values.

15 I use the Huber-White 'robust' sandwich estimator of the variance-covariance matrix.