Grasping at the invisible – Bayesian models for estimating latent economic variables

Prepared by Hugh Miller and Isabella Lyons

Presented to the Actuaries Institute
2021 All-Actuaries Virtual Summit
27 April – 21 May 2021

Abstract

Actuaries have long used official statistics and market data such as the unemployment rate or the traded price of bonds to set assumptions. However, correct interpretation of these quantities often depends on elements that cannot be directly observed. For example:

- Short-term interest rates are expansionary or contractionary depending on the 'neutral' interest rate, which is generally not observed
- The severity of unemployment is usually measured relative to the non-accelerating rate of unemployment (NAIRU).

Such quantities are increasingly of interest to actuaries – for instance IFRS 17 allows more judgement about the composition of the yield curve. The challenge in estimating these is heightened by the fact that quantities such as the neutral interest rate evolve over time.

Despite actuaries being pioneers in early work on credibility theory, Bayesian approaches have largely given way to frequentist approaches. This is to our loss, since such methods have advanced substantially over the past couple of decades and they are well-suited to problems with latent variables or noise around underlying 'true' values.

The paper gives an introduction to modern Bayesian model-fitting using the stan package and then builds and applies models to estimate the latent economic variables listed above. Findings in the paper are be supported by corresponding published code so results can be reproduced. By having public models that can be routinely updated, these models will fill a gap in the Australian market, where regular estimates are not routinely available.

Keywords: Bayesian; neutral interest rate; natural interest rate; NAIRU; Stan

1 Introduction

1.1 Background

One of the longest-running tensions in Statistics is the one between Frequentist and Bayesian approaches. This section gives an overview of how this has played out over the past 70 years, with a bias towards actuarial applications. The presentation of this tension below is highly subjective, and possibly overstates the depths that Bayesian have fallen to (there has remained a solid active research stream continuously), but we feel it has instructive value. More thorough reviews of the tension can be found in works such as Pratt (1965), Barnett (1999), Rubin (1984) and Bayarri and Berger (2004).

Actuarial science, with its heavy use of statistics and data analysis, has felt this tension too. Many early actuarial models (even going back to Whitney, 1918) used the notion of credibility to construct Bayesian-type estimates used across reserving, pricing models and other key assumptions. Key historical references are Bailey (1945, 1950), Bühlmann (1969) and Bühlmann and Straub (1970). Goulet (1998) offers an overview of the history of actuarial credibility.

But, in the authors' opinion, much of the development in statistical approaches since then has followed pragmatic lines rather than philosophical ones. Before widespread computation, Bayesian statistics offered an attractive setup – a person could take a collection of previous evidence (for example, historical claim experience) and then express this as a relatively simple set of numbers (priors); there was then a theoretically convenient way to integrate new data that allowed some regard for both the historical and the new information.

With the increased availability of data and computation, even as far back as the eighties and nineties, frequentist approaches soon demonstrated considerable value. Entire claim histories could be modelled at once, rather than applying simplified update rules. Measures could be estimated and tested that no longer relied as heavily on statistical assumptions. Generalised linear models allowed high-accuracy multivariate models to be fit on reams of data, enabling greater complexity to be tested and embedded in statistical work.

This trend only continued in the early decades of the 21st century. Datasets grew even larger. Modern machine learning methods rely heavily on frequentist ideas, with little emphasis on priors and more focus on 'letting the data speak'. Frequentist approaches have thus become the dominant paradigm for statistical learning.

Of course, Bayesian statistics has continued to be used over this time. But Bayesian models have been slower to adapt to a data-heavy world. The class of techniques for computing more complex designs with arbitrary priors remains many times slower than common frequentist approaches.

But lately there are signs of a resurgence of Bayesian models. Computational power and methods have continued to improve, so that such models are even more efficient, and more accessible to researchers. And with this is the recognition that there are many important problems where the Bayesian setup is natural.

Furthermore, many modern modelling approaches draw on Bayesian ideas. For example, penalised regression approaches, such as the lasso (Tibshirani, 1996) give partial weight to parameter estimates, in a way that can be reinterpreted as taking the model of a Bayesian posterior. Similarly other model types with shrinkage terms often have Bayesian interpretations.

One domain where Bayesian models have a clear design advantage is in situations where there are latent variables. These are quantities that are driving the observed values but are not directly visible. Economics has several important latent variables that are of constant interest to policymakers, such as the neutral interest rate and the non-accelerating inflation rate of unemployment. These latent variables are also important in actuarial work, as we discuss further in Section 3. However, there are no routinely updated public estimates of these quantities – an issue this paper seeks to address.

The structure of the remainder of the paper is as follows. Section 1.2 discusses developments in Bayesian software that have contributed to recent popularity. From there our paper is effectively in two halves:

- Section 2 explores the growing popularity of Bayesian models by covering three recent examples of Stan-based Bayesian models that demonstrate its flexibility
- Section 3 examines the issue of latent economic variables and presents our results for our implementation of various models.

Brief conclusions are given in 4. Our code and results related to this paper have been posted online on github[1].

1.2 Software for Bayesian models

Once the specification of a Bayesian model is sufficiently complex, analytic solutions become too complex and computational alternatives are needed. The key family of models that produce Bayesian solutions are Markov Chain Monte Carlo (MCMC) models, which are sampling techniques that will traverse the posterior distribution in a wide set of circumstances. While types of MCMC models have been suggested since the 1950s (for example, Metropolis et al., 1953), computational and theoretical breakthroughs were needed that prevented the ideas being used widely until the early 1990s (see for example Gelfand and Smith, 1990). This in turn affected the availability of software that made such techniques available to a broader range of researchers. The key software developments since that time have been:

- BUGS (Bayesian inference Using Gibbs Sampling) software was developed from 1989 onwards, with version 0.1 released in 1993. It implemented Gibbs sampling techniques for a flexible range of model structures. It came in various flavours, including WinBUGS for Windows. See Lunn et al. (2009) for a more detailed history.
- JAGS (Just Another Gibbs Sampler) was based on BUGS and released in 2007. It is implemented in C++ which made it more easily transportable across computing environments. It also corrected some of the limitations of BUGS, including better separation of model data and assumptions.
- Stan (not an acronym – named after researcher Stanislaw Ulam) was released in 2012 as a modern interface for Bayesian models, built by a development team including Andrew Gelman, Bob Carpenter, Matt Hoffman, and Daniel Lee. It implements some newer methods, such as the Hamiltonian Monte-Carlo which has more efficient estimation of complex models (particularly those with correlated parameters). It has good integrations with other software like R.

With the development of JAGS and Stan in particular, users now have accessible free software to specify sophisticated Bayesian models. While Stan is not more efficient than JAGS in all situations, it appears to be the most popular choice for many researchers.

[1] https://github.com/hug-mi/Economic_latent_models

2 Recent examples of Bayesian models

Part of the resurgence mentioned in 1.1 is the increased usage of Bayesian models in academia and industry. We present three contemporary examples of this. All use Stan-based models, with publicly available code for interested readers.

2.1 Estimating COVID-19 infections

A defining feature of the COVID-19 pandemic is the large amount of data regularly available to both policy-makers and the public. However, the data does not reveal everything. An important example of this is that reported cases are based on the date when symptomatic people are *confirmed* positive, rather than the date they actually contract the virus (infection date). This matters – the lag between infection and detection means that a day's infection rates cannot be accurately estimated until (roughly) a fortnight later, and that this estimate will be subject to uncertainty. Furthermore, there are incorrect ways to infer the pattern of infections from the observed data.

To illustrate, first consider the toy example shown in Figure 1. Here 100 infections follow an expected distribution of reporting over the subsequent two weeks, even though all the infections relate to a single day – here a capped Poisson distribution with mean 6.

Figure 1 – Assumed distribution of reported cases relative to infection date

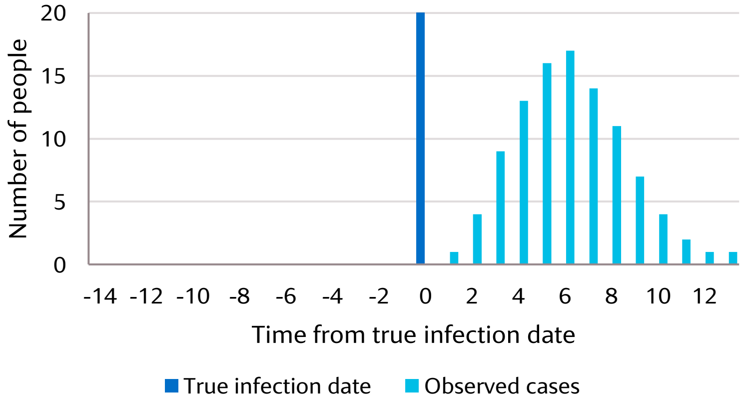

Suppose we observe the light blue distribution in Figure 1 and we know the shape of the distribution between infection and cases. Figure 2 shows two less useful ways that these infections could be projected backwards to estimate infection dates. The purple 'shift' curve subtracts six days from the observed cases. While this recovers the correct mean, it does not create the sharp spike of infections that underlies the true data. The yellow 'convolve' shape is even wider – this is the pattern obtained if the known probabilities are applied in a backwards direction (so a report on day 1 could have come from an infection during any of the previous 14 days, using the probabilities in Figure 1). This produces a much wider range than the true infection pattern, since it ignores the observed pattern of reported cases. Rather, we should seek to 'deconvolve' the probabilities to obtain a sharper infection pattern, as seen in Figure 1. See Gostic et al. (2020) for an extended discussion on this. Bayesian models are well suited to the task, particularly combined with the need to estimate an evolving reproduction number through time, and so can obtain a more realistic estimate of the historical rate of infections over time.

Figure 2 – Two potential methods for projecting reported cases backwards in time to infection date

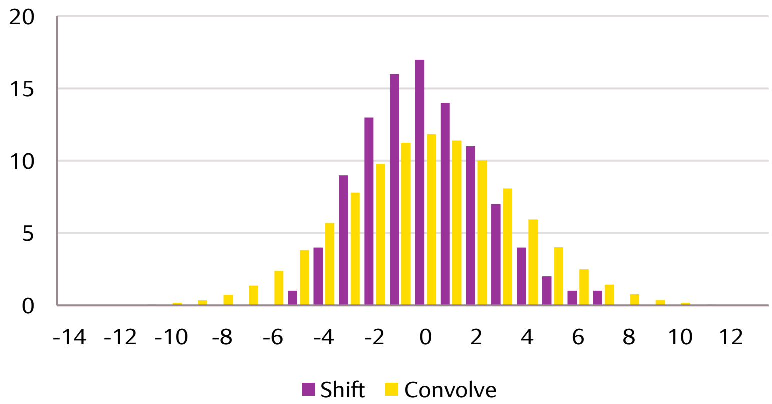

We apply such an approach using the *epinow2* package in R, which implements the thinking in Gostic et al. (2020). The model has a Stan backend to perform this deconvolution and (time-varying) infection rate evolution on time series data of COVID-19 reports. We have applied this to NSW's local infection numbers up to the 28th of July 2020[2]. This was during the second wave, where a consistent number of cases had been reported for about two weeks. Estimated numbers by infection date are shown in Figure 3.

Figure 3 – Estimated infection dates using NSW locally acquired case data. Grey bars show reporting dates and coloured lines are estimated infection numbers

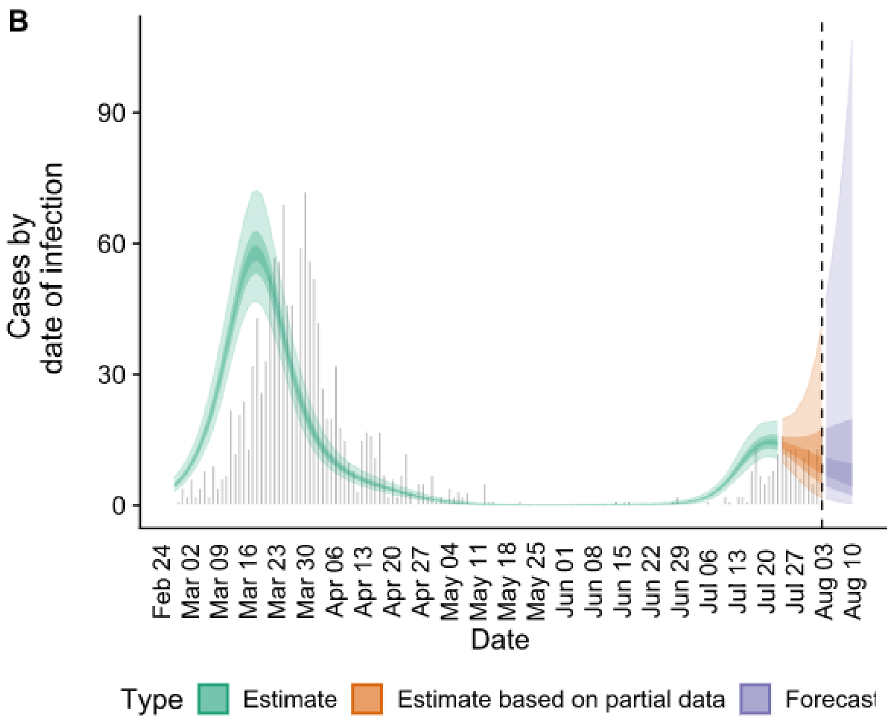

Several things are immediately apparent from the fit:

- The first wave of infections in March started to subside in mid-March, which was before some of the strictest lockdown methods were introduced (although the stricter rules no doubt sped the subsequent decline). This suggested the early measures, combined with behavioural change, contributed to flattening the curve.

[2] Data from https://data.nsw.gov.au/nsw-covid-19-data/cases

- There is reasonable evidence that the stable numbers in the second wave observed to the model date corresponds to a falling rate of infections; if infections were constant or rising we would have expected an increasing number of reports instead. This leads to a ‘stable numbers is good news’ situation, with all else equal a decrease in reported cases should follow.
- The orange ‘nowcasting’ and the purple projection both carry considerable uncertainty, with rapidly expanding confidence intervals. This reflects the inherent uncertainty introduced by the lag between infections and reported cases, as well as the changing reproduction number over time.

While these insights around COVID-19 infections are not new, we seek to point to the value in a Bayesian approach to understanding how reported cases translate to the timing of people contracting the virus.

2.2 Election forecasting

Bayesian setups are also useful for situations with evolving data, where predictions can vary over time as new information is collected. The Economist magazine ran a Bayesian prediction model[3] for the 2020 US presidential election. The model could:

- Incorporate new polls (local and national) as they were collected
- Factor in allowances such as non-response bias, state-level correlations and consistent pollster biases
- Give plausible confidence margins over time.

The main output for the vote share is shown in Figure 4.

Figure 4 – Predicted voting share from the Economist presidential election model

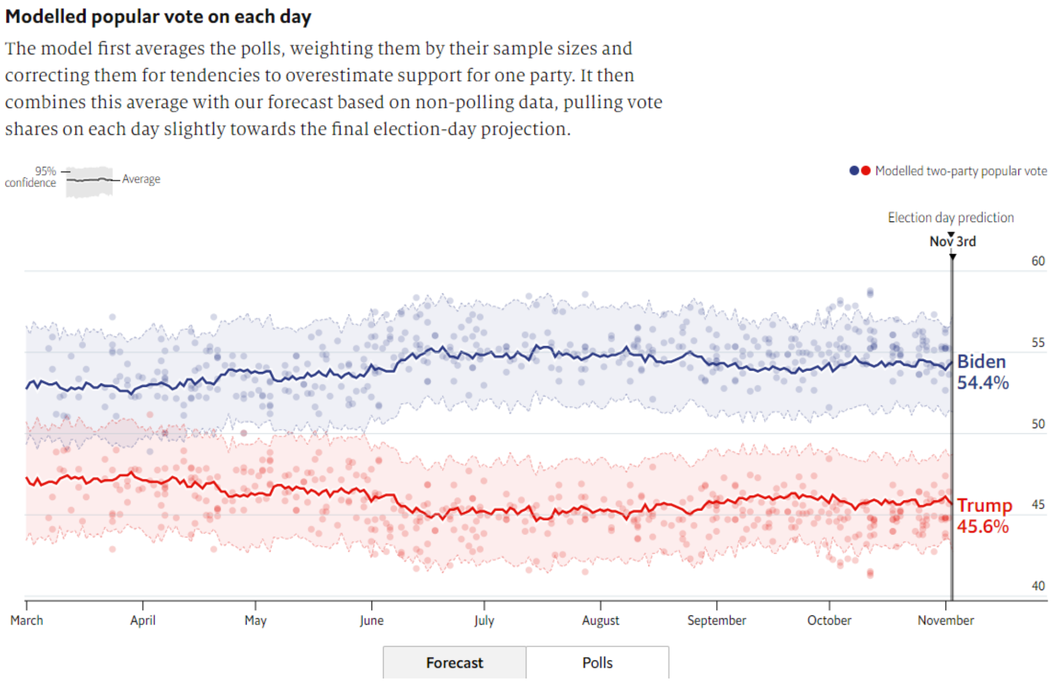

Source: Economist.com

Readers may be interested in comparing the predictions to the true two-party popular vote split, which was 52.3% to 47.7%. While the true result sits within the expected margin of error, it represented the second time that polls had underestimated Trump’s election performance.

---

[3] See https://projects.economist.com/us-2020-forecast/president

2.3 Actuarial reserving

Traditional triangle reserving models are straightforward to implement as Bayesian regression models. Glenn Meyers' monograph[4] (Meyer, 2019) gives an instructive look at:

- How a range of common reserving models can be implemented as Bayesian models in Stan
- Useful diagnostics to assess the goodness of fit for these models
- How many common frequentist approaches can understate the true variability in a reserving context
- The value of distributional information that arises from the Bayesian model.

The value of such models therefore extends to very traditional actuarial contexts.

[4] See https://www.casact.org/sites/default/files/2021-02/08-Meyers.pdf

3 Bayesian models for latent economic variables

3.1 Latent economic variables in economics

Economics is characterised by the intersection of theoretical models and time series data. Many core concepts are latent, in that we cannot directly observe them, but we can infer information about them from what we can observe.

Further challenges relate to the movement of latent quantities over time, and relationships between them and other economic quantities. One particularly stark example of change over time is the collapse in the Phillips Curve – the relationship between unemployment and inflation. Figure 5 shows how inflation has moved with the unemployment rate in Australia. The sustained growth seen since the early 1990s combined with an effective management of inflation has meant that there has been few instances of low inflation tied with high unemployment, nor high inflation tied with low unemployment. This comes through as a weakening of the Phillips curve slope, something that is generally regarded as a key foundation of economic management. The trends echo those seen internationally too.

Figure 5 – The Phillips curve in Australia over time; relationship between inflation and unemployment, 1984-2019

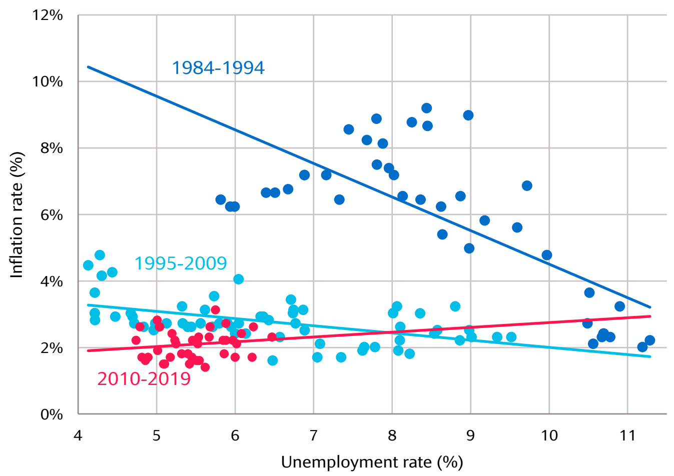

For fuller discussion of the Phillips curve see Hooper et al. (2020) and references therein. It argues that even if the national-level Phillips curve has flattened (in their case in the USA), it is not dead; there is still strong evidence for its applicability at a subnational level, as more local labour markets can experience more extreme swings, with corresponding inflationary pressures.

Neutral interest rate

The neutral interest rate is defined as *'the interest rate equates saving and investment at the level of income that is consistent with full employment and stable inflation.'* (McCririck, & Rees, 2017). In an economic environment when unemployment is sustainably low, it is the point at which a higher level of rates would be contractionary and lower rates expansionary or stimulatory. The implied balance between saving and investment also means that the rate may change over time; indeed, there have been many structural factors that have led to experts identifying a fall in the neutral interest rate.

The challenge is identifying the neutral rate along with any movements over time. Estimation is important; if the rate has fallen significantly over the past decade, then a policy rate that may have once been expansionary could actually be contractionary.

From an econometric view, we can see evidence of the level of the neutral interest rate through the interaction between rates, economic growth (as measured by GDP), inflation and unemployment. If unemployment rates are low, growth at or above trend, and inflation remains under control then rate settings may be close to neutral. However, if inflation is rising strongly at the same time, then rates could be below neutral. Conversely, if inflation and growth are both tracking lower than expectations, then that could be evidence that rates are above neutral.

There is also an obvious distinction between the *nominal* neutral interest rate (after inflation) and the *real* neutral interest rate (excludes inflation). All discussion below refers to the real rate unless otherwise specified. Lower inflation rates will obviously reduce the nominal rate without necessarily affecting the neutral interest rate.

A common model applied in many countries is the 'HLW model' of Holston et al. (2017)[5]. It suggests the natural interest rate has fallen up to 2 percentage points over the past decade for advanced economies, using a sophisticated Kalman filter to estimate the latent value of the neutral interest rate.

NAIRU

The non-accelerating inflation rate of unemployment (NAIRU) is the lowest sustainable level of unemployment an economy can maintain without causing a material increase in inflation. Estimation of the NAIRU is typically done through examination of some version of the Phillips curve; the relationship between inflation, unemployment and other key sources of cost pressure (such as wages or even oil prices) over time.

A key reference for our research is the 2017 Reserve Bank research bulletin on the NAIRU (Cusbert, 2017). There is also a recent working paper released by the Commonwealth Treasury (Ruberl et al., 2021).

Term premia

Term premia are the additional yield on bonds at various durations on top of expected interest rates. As with the other economic quantities, these cannot be directly observed but can be estimated. Estimation is useful for understanding longer-term interest rate expectations. Term premium are beyond the scope of this paper, but we note that the AOFM has started publishing regular estimates of term premia[6]. See also the related working paper (Jennison, 2017).

3.2 Neutral interest rate model

We have built a neutral interest rate model in Stan that follows the structure of the HLW model (Holston et al., 2017). Notation is show in Table 1. Table 2 describes the equations used to define the model and Table 3 gives the parameter results from our fit.

Table 1 – Notation summary

Item	Description
y_t	The log of national output (Chain-value GDP, ABS series A2304402X) at time t. y_t^* is potential log-output and $\tilde{y}_t = (y_t - y_t^*)$ the corresponding output gap
r_t^*	The neutral rate of interest at time t. The actual case rate is r_t
g_t	Trend growth (quarterly) at time t
z_t	Other components apart from growth affecting r_t^*
π_t	Inflation at time t. Based on CPI series from ABS 6401.06. π_t^e denotes inflation expectations at time t, developed according to Appendix A.1.
ϵ_t	Error term at time t, for various models, with corresponding superscripts

[5] https://www.newyorkfed.org/research/policy/rstar

[6] https://www.aofm.gov.au/publications/research/estimation-term-premium-within-australian-treasury-bonds

Table 2 – Equations used in our Neutral interest rate model

Description	Equation
Output gap	$\tilde{y}_t = (y_t - y_t^*)$
Neutral rate	$r_t^* = (4g_t + z_t)$
IS Curve	$\tilde{y}_t = y_t^* + a_1\,\tilde{y}_{t-1} + a_2\tilde{y}_{t-2} - \frac{a_r}{2}[(r_{t-1} - r_{t-1}^*) + (r_{t-2} - r_{t-2}^*)] + \epsilon_t^{\tilde{y}}$
Aggregate supply	$\pi_t = (1-\beta_1)\pi_t^e + \frac{\beta_1}{3}(\pi_{t-1} + \pi_{t-2} + \pi_{t-3}) + \beta_2\tilde{y}_{t-1} + \epsilon_t^{\pi}$
Potential output evolution	$y_t^* = y_{t-1}^* + g_t + \epsilon_t^{y^*}$
Growth rate evolution	$g_t = g_{t-1} + \epsilon_t^{g}$
Other components of neutral rate evolution	$z_t = z_{t-1} + \epsilon_t^{z}$

Table 3 – Parameter priors and posteriors (March 1965 to March 2020 model)

Param.	Term	Prior dist.	Prior (μ, σ^2)	Posterior mean	80% CI
a_1	$\tilde{y}_{t-1}$	Normal	(1.1,0.25)	0.87	(0.75,0.99)
a_2	$\tilde{y}_{t-2}$	Normal	(-0.2,0.25)	0.03	(-0.08,0.14)
a_r	$(r_{t-1} - r_{t-1}^*) + (r_{t-2} - r_{t-2}^*)$	Inv gamma	(0.15,1)	0.04	(0.03,0.05)
β_1	$\pi_{t-1} + \pi_{t-2} + \pi_{t-3}$	Beta	(0.625, 0.13)	0.01	(0.00,0.04)
β_2	$\tilde{y}_{t-1}$	Normal	(0.5,0.09)	0.23	(0.15,0.33)
$\epsilon_t^{\tilde{y}}$		Inv gamma	(1,1.98)	0.01	(0.01,0.01)
ϵ_t^{π}		Inv gamma	(1,1)	1.27	(1.19,1.36)
$\epsilon_t^{y^*}$		Inv gamma	(1,1)	0.30	(0.21,0.40)
ϵ_t^{g}		Inv gamma	(0.25,1)	0.06	(0.04,0.09)
ϵ_t^{z}		Inv gamma	(0.8,1)	0.63	(0.31,1.05)

Figure 6 – Estimated neutral interest rate, Stan model

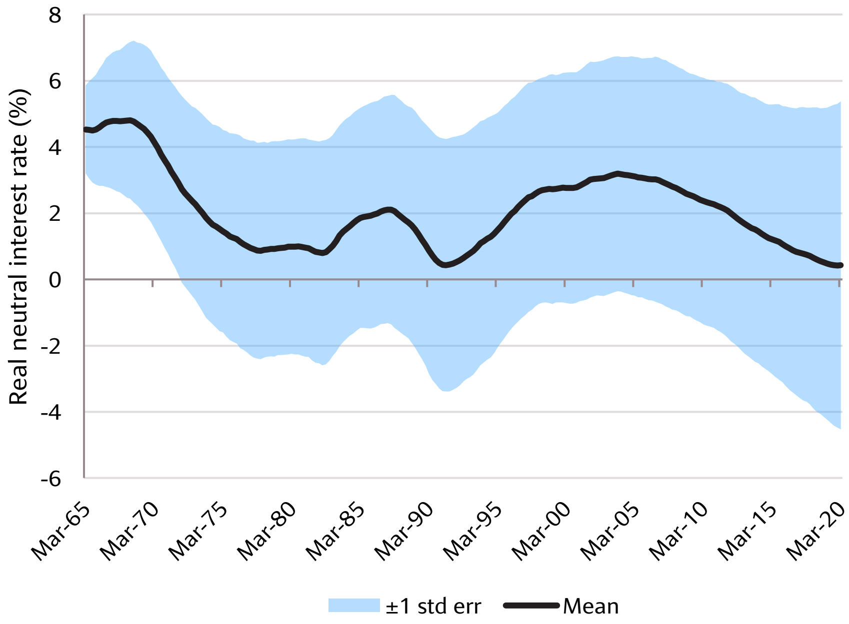

Our estimated neutral interest rate over time is shown in Figure 6. We observe:

- The estimated neutral rate has fallen significantly since the mid-2000s decade. The estimated rate at December 2003 is 3.2% (which equates to 5.7% nominal assuming 2.5% inflation), whereas it falls to just 0.4% at March 2020, a 2.8% difference.
- The uncertainty is large, as evident by the wide standard error shaded region. This appears to be a feature of our particular model setup, with wide confidence intervals persisting across a range of different model variations. We believe one source of the uncertainty is expressing the neutral interest rate as a sum of two other latent variables, z_t and g_t. This has an effect of compounding uncertainty. Additionally, fitting the model across a longer time period tends to increase the error, since the uncertainty appears to 'build up' over time for the model setup. However, we note that overall uncertainty in the point estimate does not necessarily translate to uncertainty in the decrease seen over the past decade. Wide confidence intervals are a feature of other neutral interest rate models too.
- The model is (disappointingly) sensitive to setup:
  - Inclusion of COVID-19-affected quarters significantly affects the fit, due to the abrupt shock to some time series. The fit above excludes quarters after March 2020, but we discuss extensions to the model required to adjust for the pandemic in Appendix A.2.
  - The tightness of the prior on ϵ_t^g and ϵ_t^z materially affect how 'wiggly' the estimated neutral rate in Figure 6 appears. So there has been some judgement applied in the selection of these priors.
  - The choice of time periods can also significantly affect results. We believe much of this relates to the inclusion of outliers, which can have a disproportionate impact on model estimates and posteriors.

  Despite this, we believe our proposed setup produces reasonable outputs – the estimates are intuitive and compare well to other potential estimates (see Figure 7).

The uncertainty and sensitivity does mean that a certain amount of humility is required when assessing the model, consistent with the comments in (McCririck, & Rees, 2017). The neutral interest rate is hard to estimate, and reviewing a range of models can be useful for understanding some of this uncertainty.

One way we can test the reasonableness of the Stan model is to compare to a Kalman filter based model, using modified code from the HLW model. We have fit this model on data from December 1983 and compare the outputs below. While the equation setup of the models is the same, the optimisation approach is different; the HLW approach constrains the variability of error terms more tightly over time.

The figures align reasonably well; the HLW model actually projects a larger fall in the neutral interest rate, as well as carrying a higher estimate through the early decades of the fit. In either case, a large structural shift in interest rates appears to have taken place.

Figure 7 – Comparison of the Stan model to HLW model, applied by authors to Australian data.

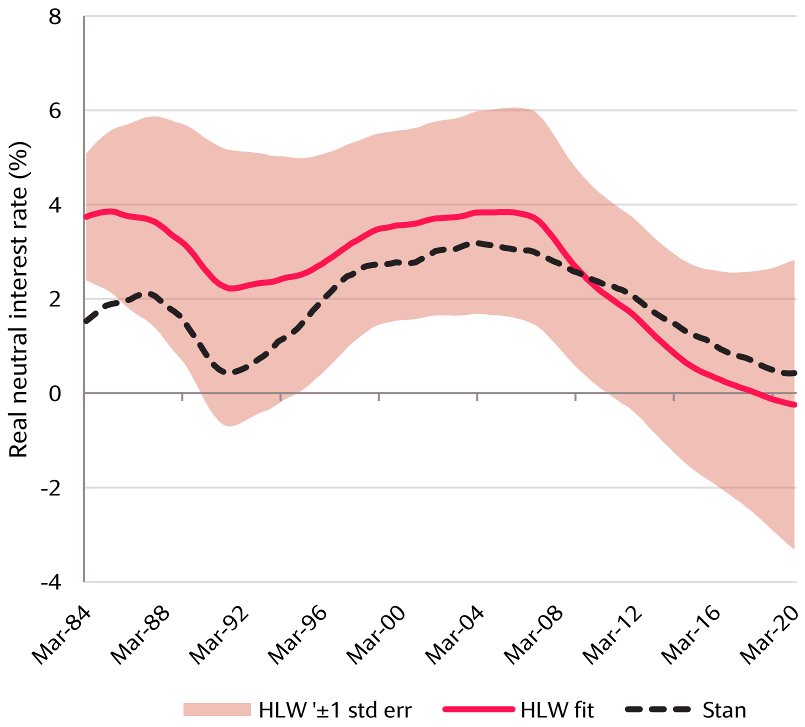

Code to produce HLW fit based on the code available at https://www.newyorkfed.org/research/policy/rstar

It is also worth recognising that even though we a primarily interested in estimating the neutral interest rate, other quantities are also estimated. These are shown in Figure 8.

Figure 8 – Other economic quantities estimated by the Stan neutral interest rate model

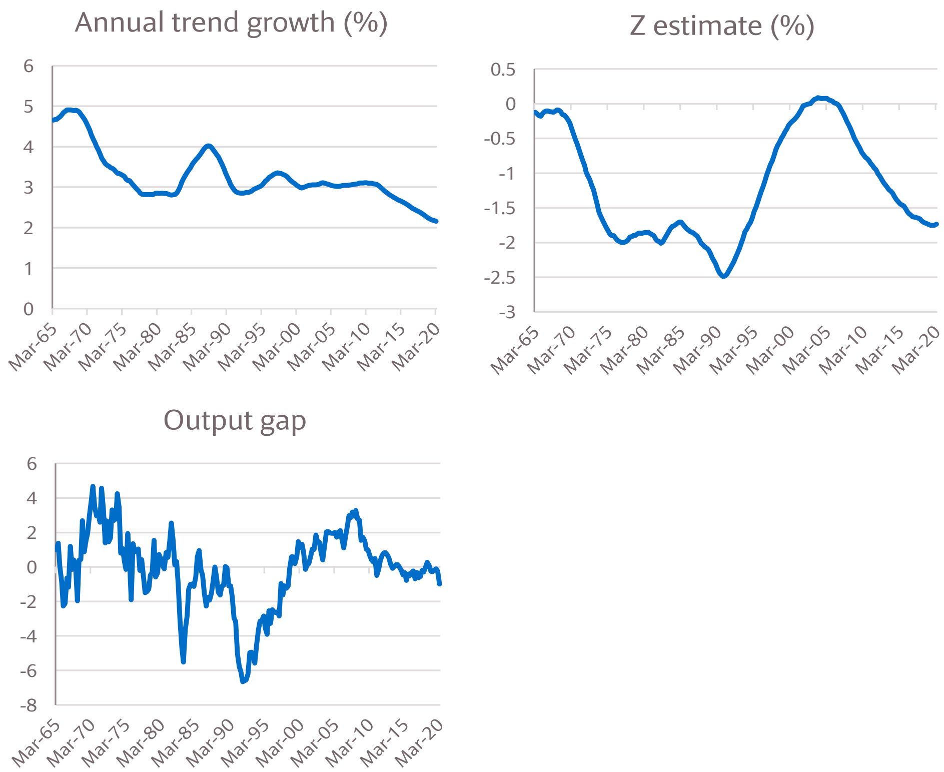

3.3 NAIRU model

The description and results of the NAIRU model follow in Table 4, Table 5 and Table 6. The structure follows that of Cusbert (2017). While the Phillips curve equation may appear intimidating at first glance, there is only one latent variable (the NAIRU), with other factors known, so this is a significantly simpler Bayesian model to estimate than the neutral interest rate model.

Table 4 – Notation summary

Item	Description
π_t	Inflation at time t. Based on CPI series from ABS 6401.06. π_t^e denotes inflation expectations at time t, developed according to Appendix A.1.
$\pi_t^{m,4}$	Annual growth in consumer imports price deflator
D_{76}	Indicator variable for timepoints prior to 1977
oil_t	Change in Brent crude oil price, time t
U_t	Unemployment rate, time t
$NAIRU_t$	NAIRU, time t
ulc_t	Unit labour cost growth, time t

Table 5 – Equations used in our NAIRU model

Description	Equation
Inflation (Phillips) curve	$\pi_t = \delta_p \pi_t^e + \beta_1 \pi_{t-1} + \beta_2 \pi_{t-2} + \beta_3 \pi_{t-3} + \phi\, ulc_{t-1} + \gamma_p \frac{(U_t - NAIRU_t)}{U_t} + \lambda_p \frac{\Delta U_{t-1}}{U_t} + \alpha_1 (\pi_{t-1}^{m,4} - \pi_{t-2}^{m,4}) + D_{76} \psi_p\, oil_{t-2} + \epsilon_t^p$
Labour cost growth	$ulc_t = \delta_{ulc} \pi_t^e + \omega_1 \pi_{t-1} + \omega_2 \pi_{t-2} + \gamma_{ulc} \frac{(U_t - NAIRU_t)}{U_t} + \lambda_{ulc} \frac{\Delta U_{t-1}}{U_t} + D_{76} \psi_{ulc}\, oil_{t-2} + \epsilon_t^{ulc}$
NAIRU evolution	$NAIRU_t = NAIRU_{t-1} + \epsilon_t^N$

Table 6 – Parameter priors and posteriors (March 1983 to December 2020 model)

Param.	Term	Prior dist.	Prior (μ, σ^2)	Posterior mean	95% CI
δ_p	π_t^e	Normal	(0.4,1)	0.37	(0.13,0.62)
β_1	π_{t-1}	Normal	(0.3,1)	0.38	(0.20,0.55)
β_2	π_{t-2}	Normal	(0.2,1)	0.27	(0.09,0.45)
β_3	π_{t-3}	Normal	(0.05,1)	0.00	(-0.18,0.18)
ϕ	ulc$_{t-1}$	Normal	(0.05,1)	0.00	(0.00,0.01)
γ_p	$\frac{(U_t - NAIRU_t)}{U_t}$	Normal	(-0.2,0.25)	-0.30	(-0.63,-0.02)
λ_p	$\frac{U_{t-1}}{U_t}$	Normal	(-0.4,1)	0.38	(-0.47,1.22)
α_1	$(\pi_{t-1}^{m,4} - \pi_{t-2}^{m,4})$	Normal	(0.01,1)	0.00	(0.00,0.01)
ψ_p	D_{76}oil$_{t-2}$	Normal	(0.1,1)	0.37	(0.13,0.62)
δ_{ulc}	π_t^e	Normal	(0.7,1)	1.11	(-0.21,2.45)
ω_1	π_{t-1}	Normal	(0.5,1)	0.49	(-0.76,1.78)
ω_2	π_{t-2}	Normal	(0.1,1)	-0.34	(-1.63,0.94)
γ_{ulc}	$\frac{(U_t - NAIRU_t)}{U_t}$	Normal	(-5,1)	-3.56	(-5.32,-1.89)
λ_{ulc}	$\frac{U_{t-1}}{U_t}$	Normal	(-2,1)	-2.07	(-4.00,-0.17)
ψ_{ulc}	D_{76}oil$_{t-2}$	Normal	(0.05,1)	0.03	(-1.87,1.97)
ϵ_t^p		Inv gamma	(0.3,0.5)	0.20	(0.18,0.22)
ϵ_t^{ulc}		Inv gamma	(1,0.5)	2.58	(2.31,2.90)
ϵ_t^N		Inv gamma	(0.5,0.25)	0.24	(0.11,0.43)

Figure 9 – NAIRU estimate from our Stan model, compared to actual unemployment

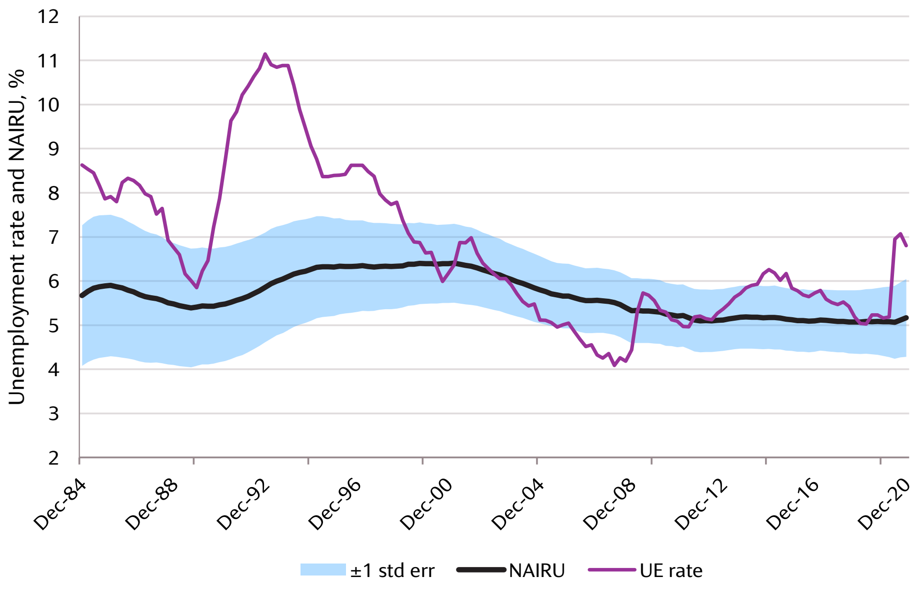

Figure 9 shows our model estimates, with the NAIRU being relatively stable over the past decade and averaging 5.1% over calendar 2020. In contrast to the Neutral interest rate model, confidence intervals are tighter and there is less (but still some) sensitivity to the choice of priors, start dates and including pandemic-affected periods.

There has been recent commentary in the media and policy circles about a 'sub-5%' target for unemployment[7]. Recent Treasury modelling (Ruberl et al., 2021) produced estimates similar to ours, but found other estimates in the 4.5% to 5% range. This makes sense as this lower bound has only been tested once in recent decades (the mining boom prior to the global financial crisis in 2007 and 2008); unemployment fell near 4% and there was clear inflationary pression. However, inflation has consistently been weaker than expected since then, so a sub-5% unemployment rate has a manageable inflation risk. The lower bounds of the NAIRU have not been tested for some time.

Other commentators have also pointed to other structural factors that have been changing, particularly in recent years; for instance demographic shifts and increased rates of part time work. Some of these effects may have a bearing on the NAIRU. However they are not explicitly built into typical models but will likely be an active area of future research.

Internationally, there is evidence of slowly decreasing NAIRU estimates. We reproduce some of these estimates in Figure 10. Interesting, while the level of the NAIRU varies a lot by region, the underlying movement is similar, reflecting global correlations in both unemployment and inflation.

[7] See, for example, Peter Martin at https://theconversation.com/exclusive-top-economists-back-budget-push-for-an-unemployment-rate-beginning-with-4-159989

Figure 10 – International NAIRU estimates

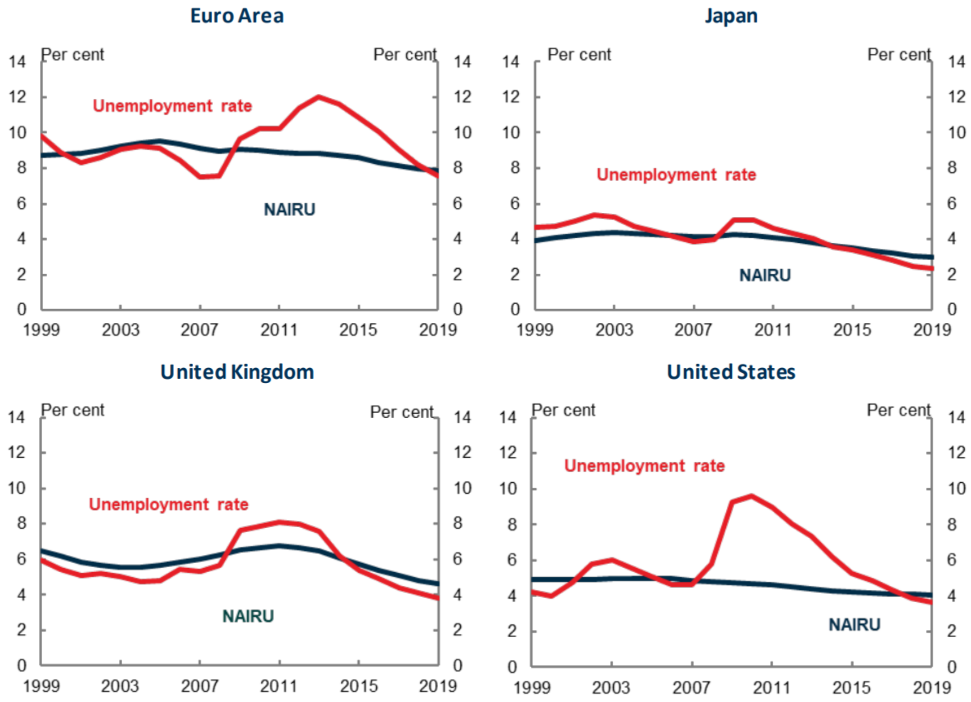

Source: OECD via Ruberl et al. (2021)

4 Conclusions and Implications to actuarial work

While the purpose of the paper is primarily to note the rising use of Bayesian models and apply them to some important economic questions, we note here some implications for actuaries.

4.1 Bayesian models

Most actuarial education still has Bayesian ideas and credibility models in the syllabus. The UK (and by extension Australia) teaches Bayesian statistics as part of the Foundation (Part 1) courses, but without much emphasis on modern MCMC approaches. We believe that expanding this teaching will be a worthwhile expansion to the actuarial toolkit.

Stan models will not reflect all actuarial work. In situations where data is plentiful and the targets of interest directly observable, the benefits of a Bayesian model may be limited and the costs material (e.g. longer fit times and more complex setup). However, a working knowledge of Bayesian ideas can still be useful for specific applications. For example, hierarchical setups are common in actuarial pricing, as are time series updates, both of which can benefit from a Bayesian thinking. Stan coding would also be a good complement to existing programming embedded in actuarial courses.

4.2 Latent economic variables

We conclude that:

- Latent economic variables have explanatory power and can be estimated regularly.
- There is strong evidence of a structural fall in the neutral interest rate, in the range of 2-3 percentage points over the past 10-20 years. Uncertainty is high.
- The NAIRU is in the vicinity of 5% and can be estimated reasonably efficiently by the Bayesian model. The lower bound has not been substantially tested over the past decade, so other literature suggesting a NAIRU in the 4.5% to 5% range, say, may be justified.

4.3 Actuarial assumption setting

Economic assumptions are core to many actuarial estimates. Some direct implications are:

- The lower neutral interest rate means that mean-reverting models that assume yields (or other related values) will return to pre-GFC levels are likely optimistic. Lower interest rates are a semi-permanent feature, even if some aspects (such as emergency COVID-19 related rates) are shorter-term. IFRS17 gives actuaries some flexibility around discount rate selection, so insights into neutral interest rate and expected yields are useful.
- An estimate of the NAIRU gives guidance on inflation risks. As unemployment approaches 4%, say, the models suggest there is evidence of potential higher inflation which will affect cashflows such as insurance liabilities. Actuarial estimates often do not directly reference unemployment rates, but having some regard could improve liability estimates.

5 References

- Bailey, A.L. (1945) "A Generalized Theory of Credibility." Proceedings of the Casualty Actuarial Society 32: 13-20.
- Bailey, A.L. (1950) "Credibility Procedures, Laplace's Generalization of Bayes' Rule and the Combination of Collateral Knowledge with Observed Data." Proceedings of the Casualty Actuarial Society 37, 7- 23.
- Barnett, V. (1999). Comparative statistical inference (Vol. 522). John Wiley & Sons.
- Bayarri, M. J., & Berger, J. O. (2004). The interplay of Bayesian and frequentist analysis. Statistical Science, 58-80.
- Biihlmann, H. (1969) "Experience Rating and Credibility." ASTIN Bulletin 5, 157-165.
- Bühlmann, H., & Straub, E. (1970). Glaubwürdigkeit für schadensätze. *Bulletin of the Swiss Association of Actuaries*, *70*(1), 111-133.
- Cusbert, T. (2017). Estimating the NAIRU and the Unemployment Gap. *RBA Bulletin, June*, 13-22.
- Gelfand, A. E., & Smith, A. F. (1990). Sampling-based approaches to calculating marginal densities. Journal of the American statistical association, 85(410), 398-409.
- Goulet, V. (1998). Principles and application of credibility theory.
- Gostic, K. M., McGough, L., Baskerville, E. B., Abbott, S., Joshi, K., Tedijanto, C., ... & Cobey, S. (2020). Practical considerations for measuring the effective reproductive number, R t. *PLoS computational biology*, *16*(12), e1008409.
- Holston, K., Laubach, T., & Williams, J. C. (2017). Measuring the natural rate of interest: International trends and determinants. *Journal of International Economics*, *108*, S59-S75.
- Hooper, P., Mishkin, F. S., & Sufi, A. (2020). Prospects for inflation in a high pressure economy: Is the Phillips curve dead or is it just hibernating?. Research in Economics, 74(1), 26-62.
- Jennison, F. (2017). Estimation of the term premium within Australian Treasury Bonds. *Australian Office of Financial Management Working Paper*, *1*.
- McCririck, R., & Rees, D. (2017). The neutral interest rate. *RBA Bulletin*, *September*, 9-18.
- Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., & Teller, E. (1953). Equation of state calculations by fast computing machines. The journal of chemical physics, 21(6), 1087-1092.
- Meyers, G. (2019). Stochastic loss reserving using Bayesian MCMC models, 2nd edition. Arlington County, VA: Casualty Actuarial Society.
- Pratt, J. W. (1965). Bayesian interpretation of standard inference statements. Journal of the Royal Statistical Society. Series B (Methodological), 27(2), 169-203.
- Ruberl, H., Ball, M., Lucas, L., and Williamson, T. (2021). Estimating the NAIRU in Australia. Treasury Working Paper, 2021-01, accessed in May 2021 at https://treasury.gov.au/sites/default/files/2021-04/p2021-164397_estimatingthenairuinaustralia.pdf
- Rubin, D. B. (1984). Bayesianly justifiable and relevant frequency calculations for the applies statistician. The Annals of Statistics, 1151-1172.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. Journal of the Royal Statistical Society: Series B (Methodological), 58(1), 267-288.
- Whitney, A.W. (1918) The Theory of Experience Rating, Proceedings of the Casualty Actuarial Society, 4, 274-292.

Appendix A Technical details

A.1 Approach to inflation expectations

We have used the RBA series of inflation expectations (RBA statistics G3) to produce an estimate of inflation expectations comparable to the figures reported in Jennison (2017). The adopted estimate was

$$\alpha + \beta_1 x_1 + \beta_2 x_2 + \beta_3 x_3 + \beta_4 x_4 + \beta_5 x_5$$

With terms as per the table below.

Parameter	Value	Corresponding Term	Description
α	-0.29512		Intercept
β_1	0.00312	x_2	Union officials' inflation expectations – 1-year ahead
β_1	0.00000	x_2	Union officials' inflation expectations – 2-year ahead
β_1	0.05665	x_2	Market economists' inflation expectations – 1-year ahead
β_1	0.10977	x_2	Market economists' inflation expectations – 2-year ahead
β_1	0.53562	x_2	Break-even 10-year inflation rate

Values were then (heavily) smoothed using a rolling 19 quarter window. Selections were made to minimise squared error. A comparison is provided in the figure below. Our fit tracks Cusbert (2017) relatively closely but is more heavily smoothed.

Figure 11 – Adopted inflation expectations compared to those of Cusbert (2017)

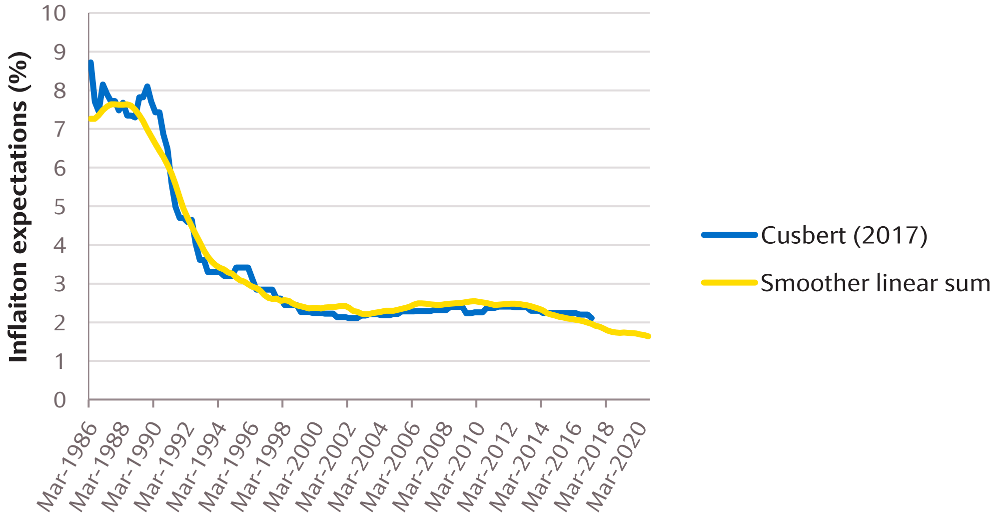

A.2 COVID-19 Adjustment to the neutral interest rate model

The neutral interest rate model overreacts significantly to the abrupt shocks of some time series, particularly output. To adjust for the economic downturn caused by lockdown measures, the output gap equation needs adjustment from that in Table 2.

This is done by assuming potential output has an additional constraint, with the severity of constraint based on the level of COVID-19 restrictions in place. The result is an adjusted output gap equation:

$$\tilde{y}_t = (y_t - y_t^*) - \phi c_t$$

Where ϕ is estimated to translate the effect of c_t, the stringency of restrictions in place at time t.

The adjustment is based on the method used by Holston et al. (2017) to adapt their model for the pandemic. Their choice of measure for stringency of restrictions is the Oxford Covid-19 Government Response Tracker[8], from which we have used the Australian index to adjust our model.

[8] https://github.com/OxCGRT/covid-policy-tracker