Speech

The Neutral Rate: The Pole-star Casts Faint Light

RESERVE BANK OF AUSTRALIA

Luci Ellis[*]
Assistant Governor (Economic)

Keynote Address to Citi Australia & New Zealand Investment Conference
Sydney – 12 October 2022

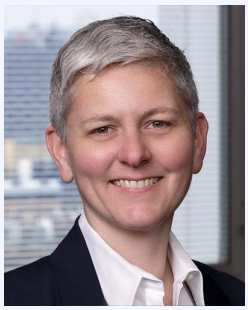

Thanks to Citi for the opportunity to speak today. My topic is the so-called 'neutral interest rate'. The concept has attracted a lot of attention in recent months. It's no surprise why. Until earlier this year, policy interest rates – in Australia, the cash rate target – had been held at very low levels to support economies through the COVID-19 pandemic and into recovery. Such low levels are understood to be unusual – they are an emergency setting for extraordinary times.

But now, with unemployment close to a 50-year low and inflation currently very high, policy stimulus is no longer needed. It's natural in these circumstances to think of policy getting back to a 'more normal' setting – although one could argue that times are still far from normal, just in a different way to the previous two years. Regardless, it's understandable that in recent months people have been wondering what 'normal' policy might be. To the extent it means 'no longer providing stimulus', this then becomes a question of what 'neutral' policy might mean. Economic theory contemplates a risk-free or policy rate that can be considered 'neutral'. This rate – whatever it is – is usually labelled 'r^*'.

Today I'm going to talk about how a neutral rate might be defined and measured, what might determine its level, and how we might respond, knowing that it could change. In doing so, I want to emphasise that the neutral interest rate is a long-run concept. There are other ways of thinking about this; some people make a distinction between 'long-run r^*' and 'short-run r^*'.[1]

What is it?

When the policy interest rate is low enough, it stimulates demand. And when it is high enough, it restrains demand. Somewhere in the middle, there must be a rate that is neither contractionary nor expansionary.

That's one way to frame the neutral rate. Certainly that's how we have defined it publicly in the past. Dig a little deeper, though, and there are some subtleties in this definition that are worth unpacking.

Consider what it means for policy to be stimulatory: growth would tend to increase, unemployment decline and inflation accelerate. But what if some negative shock hits the economy that offsets the effect of monetary policy? Then we wouldn't see those tendencies be borne out in the data. This means that, when we think about policy being expansionary or contractionary, underneath those concepts is always the question: 'compared with what?'. Well, compared with the situation if we had held interest rates at a different level.

Rather than seeing the economy expand or slow in response to the stance of policy, at times policy will only be cushioning the effects of forces working in the other direction. The neutral rate is therefore best thought of as a long-run concept. It is the interest rate that keeps growth at trend and inflation constant at target, when no other shocks are hitting the economy. Put another way, it is where the policy rate settles once all shocks have played out.

Still another definition stems from the first – for growth to be constant at its trend rate and inflation to be constant at target, the levels of desired saving and desired investment must balance each other. So the neutral rate can also be described as the level of the policy rate that balances desired saving and investment. This framing can be useful when thinking about the determinants of the level of the neutral rate, and therefore what might cause the neutral rate to change.

Before turning to the question of how to measure the neutral rate, we should note some important features that are implied by these definitions:

1. **The neutral rate is not directly observable**. Like all the other so-called 'star' variables, such as the NAIRU or the rate of growth in potential output, it has to be inferred from data.
2. **It is mostly a global concept**. Capital can flow across borders. This means that the desired saving and desired investment that must be balanced are *global* desired saving and *global* desired investment. Shifts in desired saving and investment globally therefore influence the neutral rate across economies.[2]
   That said, country-specific factors also matter. For example, economies with above-average trend growth may attract more investment, leading to a higher neutral rate. Neutral real interest rates can therefore differ across economies even if capital is fully mobile across borders. Neutral rates expressed in nominal terms can also differ if trend inflation differs across countries. So different inflation targets imply different nominal neutral rates, even though the neutral real rate won't necessarily be different.
3. **It is an emergent property of the economic system**. The neutral rate arises from the combined behaviour of many actors in the economy. It is not an 'exogenous' variable from outside the system and shouldn't be treated as such. Rather, causation runs both ways – economic outcomes can influence the level of the neutral rate, and vice versa. So we should expect that it can change over time, and that these changes have underlying economic explanations. Changes in the neutral rate are not random or external shocks. Like the other 'star variables', the neutral rate is an economic fundamental, but fundamental does not mean fixed.
4. **The neutral rate is a *real* interest rate.** Economic theory holds that it is real, inflation-adjusted interest rates that matter for people's saving and investment decisions. People look though the expected inflation component of nominal rates; lenders expect to be compensated for it, and borrowers expect to have to provide that compensation.

From real back to nominal

Real interest rates are not readily observed, however, and monetary policy is conducted in nominal terms. This means that once you have inferred the level of the neutral real rate, you have to translate that back to the nominal rate by reflating using some measure of inflation. The choice of which inflation rate to use is not clear cut. If you are focusing on these longer run conceptions of the neutral rate, a longer run or trend rate of inflation makes most sense. But there are other choices.

The saving and investment decisions that inform the level of the neutral rate depend on expectations of the future. As such, it makes sense to use some measure of inflation expectations, ideally over a multi-year horizon. This adds another layer of uncertainty to the calculation. Inflation expectations are imperfectly measured at every horizon, and it is not obvious which horizon should be used. A one-year horizon is a common choice (and can be seen in some graphs in this speech) – but it is a choice driven largely by data availability, rather than being well grounded in theory or empirical evidence. Another approach is to use the inflation target – in our case, the

midpoint of the target band. This can be justified by the definition of the neutral rate as the rate that keeps inflation at target and growth at trend.

Given the difficulties and judgements involved in choosing a forward-looking inflation rate, people sometimes resort to using recent observed inflation instead. That's understandable. Most of the time, when inflation is stable, it won't put you too far wrong.

Conceptually it isn't quite right, though. People's decisions depend on their beliefs about the future, not solely on the recent past. And there are times when using past inflation produces very different estimates of current real rates than a trend or forward-looking measure. We are definitely in one of those times. Current inflation is very high, and expected to stay high in the short term. But beyond the next year, inflation expectations remain well anchored inside the target range, both here and overseas (Graph 1).

Using the current or near-term inflation rate would imply that real interest rates were very negative at the moment. That doesn't seem like the right conclusion to draw; if it were true, borrowing behaviour would be a lot more exuberant than it currently is. So while we can learn something from that calculation – and I show it later – it isn't a preferred measure.

Graph 1

Inflation Expectations

Australia, inflation swaps measure

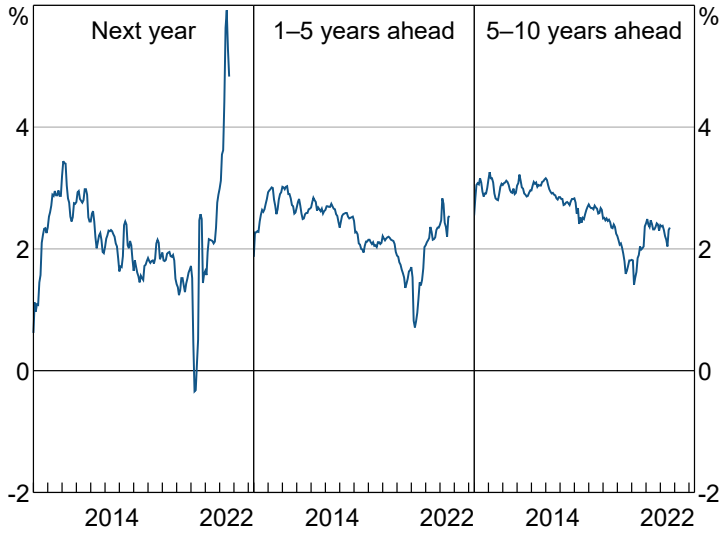

Sources: Bloomberg; RBA

Measuring the real rate could also be subject to perception biases. It's possible that saving and investment decisions depend more on people's and firms' expectations about growth in their *own* prices or wages, than on their expectations about inflation in the whole economy. If these own-price expectations are systematically biased, they will not aggregate to the expected rate of inflation in the whole economy. That is, when you ask people what they expect inflation across the whole economy to be, you will get a different answer than if you could aggregate all these own-price expectations that actually drive people's behaviour. If these biases are large enough, the real rate you infer by subtracting an economy-wide measure of inflation expectations won't be the real rate that people actually perceive and act on.

How is it measured?

The definitions of the neutral rate just described relate to a long run when all the shocks have worked their way through. To infer it, you have to identify the confounding shocks and strip out their effects on the economy. That is no small task. Fortunately, we have data and econometric models, so the cause isn't completely hopeless.

Before we get into those more sophisticated approaches to measurement, it is worth mentioning one helpful rule of thumb: in the long run, the neutral real rate is highly likely to be positive. This is because of time preference. People would rather have a dollar now than a dollar tomorrow or next year. They demand some compensation for having to wait, even if the payoff next year is risk-free. Because the policy rate represents the risk-free short-term interest rate, it embeds people's underlying time preference. So we should expect the neutral real rate also to be positive, though it could be small. And when you add in the condition that inflation should be at target, it implies a nominal neutral rate in Australia of at least 2½ per cent.

Because the definition of the neutral rate can be framed in different ways, and because of the inherent difficulties in measurement, at the Bank we use nine different models to estimate it. These nine models are derived from three broad approaches. For each approach, we construct three models, each with minor variations.

- *Approach 1: The neutral rate is inferred from financial market pricing.* It estimates the neutral rate as the average of expected future short-term real interest rates derived from a model of the term structure of nominal and real interest rates, adjusting for term premia.[3] The underlying assumption is that the real interest rate should revert to its neutral rate over time as the effects of shocks fade. The three models estimated using this approach assume different horizons over which the interest rate settles at neutral – three, five and 10 years.
- *Approach 2: A statistical model is used to infer the neutral rate at each point in time as the level of the real policy rate that would prevail if inflation is at target, output is at potential and the economy is at full employment.*[4] The model updates its estimate of the neutral rate each quarter based on how much some variables deviate from the model's predictions for the previous quarter. This method makes strong, but standard, assumptions about the structure of the economy, including the relationship between output and inflation and the gap between the cash rate and the neutral rate. The three specific models we use allow for varying emphasis on different drivers of inflation – for example, domestic versus global factors.
- *Approach 3: A flexible statistical model is used to forecast the future value of the policy rate once all cyclical influences have dissipated, which is then used as an estimate of the neutral rate.* The idea is that this forecast is the trend value to which the policy rate will converge based on structural influences.[5] In practice, this method makes fewer assumptions about economic structure than the second approach. However, its estimates will only be plausible if a forecast of the future actual policy rate is actually a good estimate of the neutral rate.

We then take a simple average across the nine models to derive a central estimate of the neutral real rate. This is a standard approach to dealing with model uncertainty. Nonetheless, it is important to remember that each model is estimated with considerable uncertainty, and we are combining these highly uncertain estimates. You can see the wide range of the different model estimates in this graph; they span a range from –0.5 per cent to 2 per cent, with the model average just under 1 per cent (Graph 2).

Graph 2

Real Neutral Interest Rate

Average of the model estimates

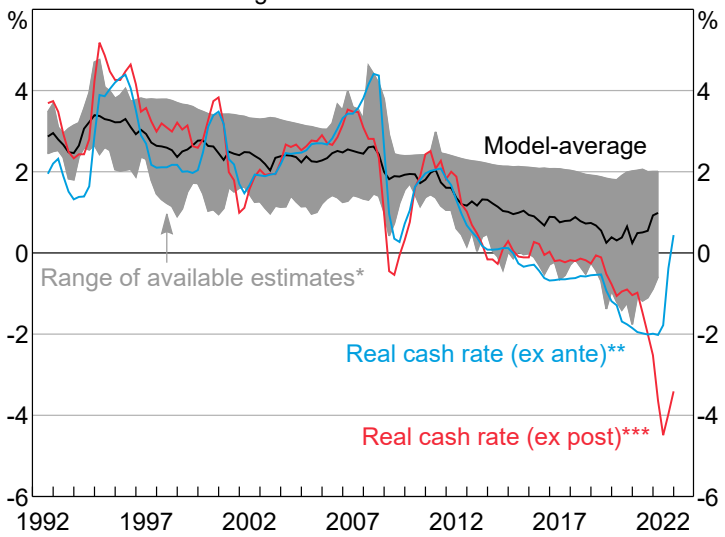

* Range of central estimates corresponding to available models.

** The ex-ante real cash rate is deflated using expected inflation; December quarter 2022 value is a preliminary estimate.

*** The ex-post real cash rate is deflated using year-ended trimmed mean inflation; December quarter 2022 value is a preliminary estimate.

Source: RBA

How has it changed?

The other point to note about the estimates in the graph are that they have declined over the course of several decades, and noticeably so in the wake of the global financial crisis (GFC). This decline is common globally (Graph 3), which is to be expected as the neutral rate largely rests on global factors. The causes of this common downward trend are less clear-cut, however. Many possible explanations have been proposed in the literature. Some are more plausible than others, and of course there could be more than one thing going on.

Graph 3

Real Neutral Interest Rates*

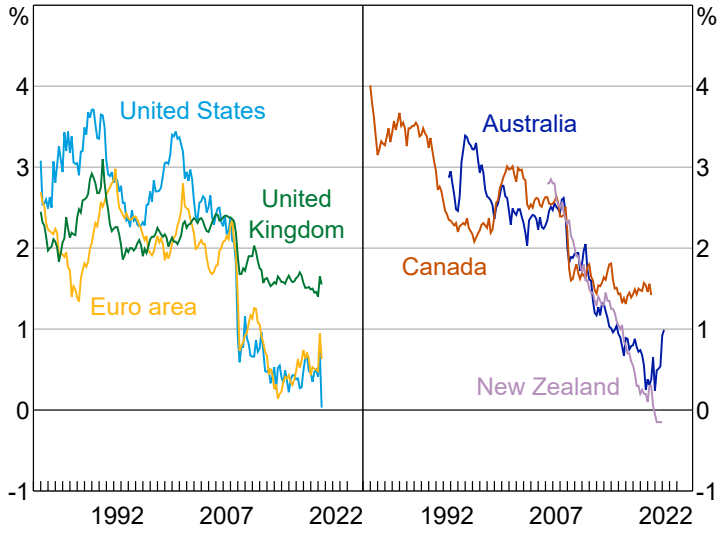

* Last data point for United States, United Kingdom, euro area and Canada is June 2020.

Sources: Holston, Laubach and Williams (2017); RBA; RBNZ

One possible cause is the apparent decline in potential output growth, which has been a multi-decade trend common across advanced economies (Graph 4). Lower potential growth reduces expected future demand. In turn, it reduces firms' incentive to invest and so decreases the interest rate that balances investment with desired saving.

Graph 4

Neutral Interest Rates and Potential GDP Growth*

Select advanced economies

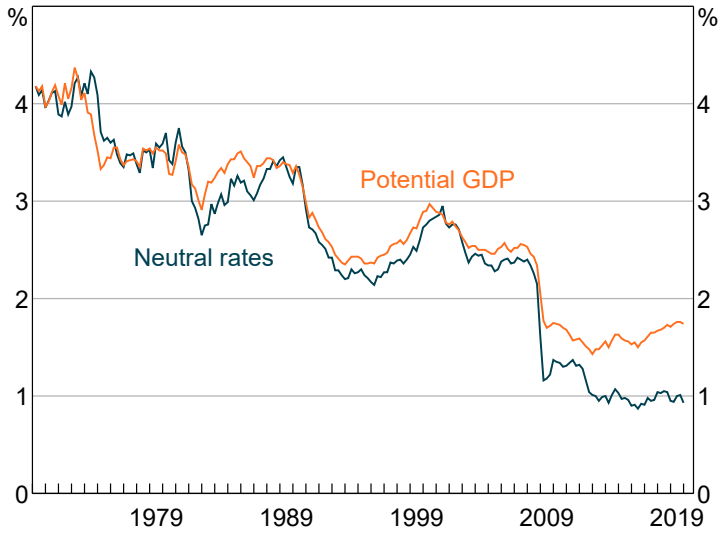

* United States, euro area, United Kingdom and Canada; PPP-weighted; last available data point December 2019.

Sources: Holston, Laubach and Williams (2017); OECD; RBA

Slower growth in potential output can stem from lower productivity growth or population ageing, or a combination of the two. Productivity growth has slowed over the past few decades, in Australia and elsewhere (Graph 5).[6] Candidate explanations for this include reduced innovation, increased regulation, slower adoption of innovation, reduced business dynamism and reduced competition.[7] The trend has been obscured lately by the effects of the pandemic. As best as we can tell, though, this slow productivity performance has continued below the surface.

Graph 5

Labour Productivity*

March 1993 = 100, log scale, quarterly

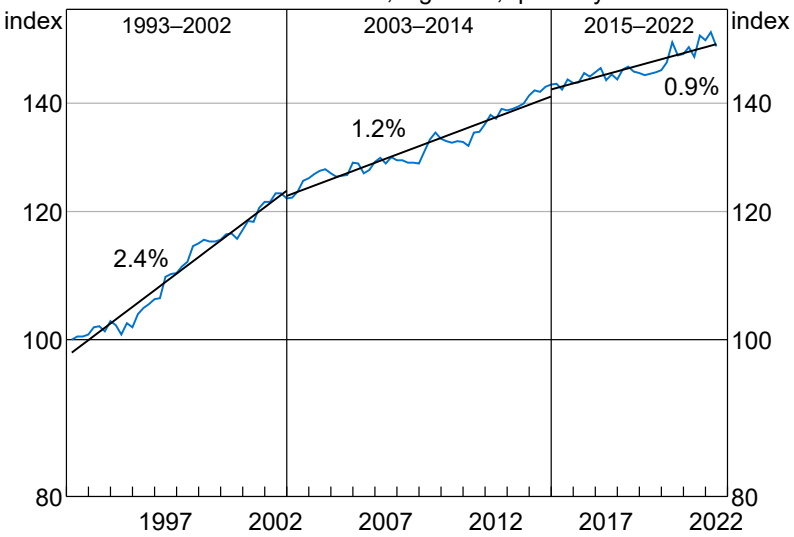

* GDP per hour worked; black lines denote linear trend; labels show average annual growth.

Sources: ABS; RBA

The effect of slower population growth and population ageing is a bit more nuanced. In principle, population ageing should lower saving because more people are retired and drawing down on their savings. But this effect seems to be overshadowed by the need to increase saving as longevity increases – that is, people need to save more to fund a longer period in retirement.

Other factors beyond potential growth could also help to explain the decline in neutral rates. One of these is reduced risk appetite.[8] The uncertainty that prevailed following the GFC dampened firms' appetite to invest; instead, the focus turned to cost-cutting. It also encouraged households to increase precautionary saving and reduce debt. In Australia, this was evidenced in the broadly stable ratio of debt to income, even though lower interest rates would ordinarily encourage leverage. Risk aversion also showed up as increased demand for safe assets – such as liquid and highly rated bonds – by global investors, especially central banks in emerging economies and pension funds in advanced economies.

A related explanation of the decline in the neutral rate is the 'global savings glut' hypothesis. Recall the earlier discussion of how the neutral rate balances global desired saving and global desired investment. If some factor drives up desired saving without also expanding investment opportunities, the neutral rate necessarily falls. According to this explanation, following the Asian financial crisis, governments in emerging Asia built up foreign exchange reserves to prevent something similar happening to them again. The counterpart to that reaction was that these countries ran large current account surpluses – they switched from being net importers of capital to net exporters. The resulting increased global saving flowed into advanced economies and put downward pressure on interest rates there.

An alternative explanation of a higher desired saving rate focuses more on the advanced economies than emerging Asia. According to this view, deregulation, lower inflation and other factors resulted in debt expanding relative to incomes. This lifted servicing burdens relative to incomes at any given (nominal) interest rate, thereby raising the saving represented by repayments of principal.

Another explanation for lower neutral rates is that they have been driven by increased inequality.[9] This relies on the observation that saving rates tend to be higher for high-income households, and high-income households' share of overall income has increased since the 1980s. This is the same period over which neutral interest rates have fallen. Because global saving and investment influence the domestic neutral rate, increased inequality

globally can reduce the neutral rate in Australia even though income inequality has not increased by as much here.[10]

These explanations for lower neutral rates are largely based on structural factors. Another possibility is that cyclical influences might be persistent enough that standard models read them as affecting the neutral rate. One such influence might be the global stance of fiscal policy. Its effects are hard to disentangle: expansionary fiscal policy boosts demand in the short run, but might reduce expected future growth in the long run if people focus on the need to fund current fiscal spending. That said, if fiscal policy is persistently tight, it is likely that monetary policy will have to be persistently accommodative to offset the contractionary influence on the economy. And if your model of the neutral rate does not explicitly include the global stance of fiscal policy, it will interpret this as the neutral rate having fallen.

This omission might go some way to explaining a puzzle surrounding standard estimates of the neutral rate. It has fallen, but policy rates fell even more and were consistently below estimates of the neutral rate for a decade (Graph 6). Despite this, inflation stayed low, both in Australia and in other advanced economies. Something must have been offsetting this apparently expansionary policy; fiscal consolidation in the wake of the GFC, especially in Europe, could have been part of that.

Graph 6

Nominal Neutral Interest Rate

Range using different inflation measures

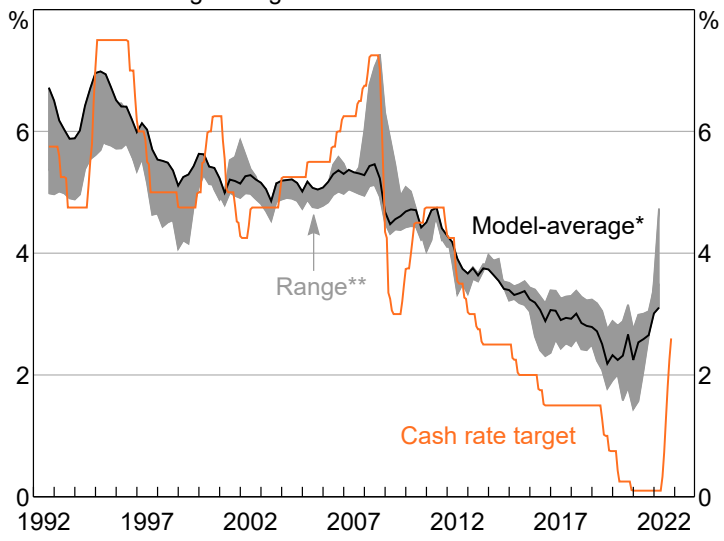

* Inflated using expected inflation.

** Inflated using the inflation target or trimmed mean inflation.

Source: RBA

What are the consequences of a low neutral rate?

It is not necessarily good or bad for the neutral rate to be at a particular level – it just is what it is. Given the possible role of demography and inequality in bringing the neutral rate to low levels, it is plausible that low neutral rates persist for the foreseeable future. If so, there are some implications to consider.

A low neutral rate in a low-inflation world means that the economy could be more frequently constrained by the effective lower bound on nominal interest rates. It becomes harder to stimulate the economy when needed. The central bank must resort to other measures beyond reducing the cash rate to support the economy. This is one reason why central banks have inflation targets that are above zero.

Another implication is that it can make it harder to achieve a desired rate of return. So investors with a mandate to achieve an absolute return could end up taking on risk that they would otherwise not take.

Alternatively, the neutral rate could increase at some point. This is also plausible: the factors holding back productivity growth could dissipate; Asian governments will not necessarily continue to accumulate foreign reserves at the rate they did in the early 2000s; and fiscal consolidation might not be as much of a priority as it was in Europe in the early 2010s.

An increase in the neutral rate has important implications, which will depend crucially on the cause of the increase. For example, those who look forward to a future of faster productivity growth must recognise that this will likely imply higher real rates as well as higher trend growth. This matters for investors, borrowers and fiscal authorities alike. These implications are mostly benign, but not if global productivity growth picks up and Australia does not keep pace.

A future increase in the neutral rate would also have implications for balance sheets. Lower average nominal interest rates mean that you can service more debt for the same repayment. If the reason average nominal rates are lower is a permanent disinflation – as occurred after Australia adopted inflation targeting in the 1990s – the resulting expansion in balance sheets is sustainable. But if average (real) rates increase again, some deleveraging would be needed. That process is unlikely to be painless.

If that were to occur, interest rate buffers become particularly valuable. In Australia and some other countries, banks determine allowable loan sizes by assessing a borrower's capacity to pay based on an interest rate that is higher than the currently prevailing rate. These measures are designed to ensure that borrowers don't overstretch and end up in distress from normal cyclical moves in interest rates. But it is easy to see that these buffers also provide some margin to cushion the adjustment if average real rates were to increase in the future. In this way, short-term prudence also supports long-term prudence.

What does it all mean for policy?

Having considered all the possible drivers of the neutral real rate and the measurement difficulties, I'd like to briefly conclude by considering the question of how that translates into our policy practice.

The Reserve Bank is responsible for meeting a flexible inflation target. As set out in our various communications over recent months, the Board's objective is to get inflation back to target. What that implies for the policy rate depends on what else is going on. Some factors could be dragging on growth and inflation; others could be boosting them.

We make a holistic assessment of the forces that are affecting the economy and whether they might be reinforcing or offsetting the impact of monetary policy; it isn't a simple mapping from inflation forecast to desired policy stance. The framework that brings that assessment together is our forecasting process, including the risks around our central projections. While far from perfect, a full-blown whole-economy forecast reduces the chance that we would miss something important.

As I mentioned at the beginning of this talk, and has been clear for some time, the strong recovery of the economy and high rate of inflation mean that the previous setting of a cash rate near zero per cent is definitely no longer appropriate. But don't think of this as a mechanistic approach of 'we have to get back to neutral', or above neutral. The neutral rate is an important guide rail for thinking about the effect policy might be having. It is not necessarily a prescription for what policy should do.

'Neutral', then, is not a destination we necessarily reach, but more a pole-star to guide us. And even then, its location is sufficiently uncertain that we are perhaps better served by paying more attention to the ground as it shifts beneath our feet than to that faraway pole-star. We need to be mindful of the limitations of our instruments, and of the prospect that the stars themselves can realign. But as we navigate the narrow path to our intended goal, we welcome any faint light that those stars may cast.

Thank you for your time.

Endnotes

[*] Much of the material in this speech is drawn from work done by the Structural Analysis and Macro Modelling section in Economic Group for a briefing to the Board in July. Particular thanks are due to Ivan Roberts and Ivailo Arsov for leading this work, and to Matt Read for thought-provoking comments on this talk.

[1] See, for example, Bailey A (2022), 'The Economic Landscape: Structural Change, Global R * and the Missing-investment Puzzle', Speech given at the Official Monetary and Financial Institutions Forum, London, 12 July.

[2] See Rachel L and LH Summers (2019), 'On Secular Stagnation in the Industrialised World', *Brookings Papers on Economic Activity*, Spring.

[3] See Hambur J and R Finlay (2018), 'Affine Endeavour: Estimating a Joint Model of Nominal and Real Term Structures of Interest Rates in Australia', RBA Research Discussion Paper No 2018-02.

[4] The statistical method used to infer the neutral rate is the Kalman filter. See McCririck R and D Rees (2017), 'The Neutral Interest Rate', *RBA Bulletin*, September, pp 9–18. See also Holston K, T Laubach and JC Williams (2016), 'Measuring the Natural Rate of Interest: International Trends and Determinants', *FEDS*, No 2016-073.

[5] This method uses a time-varying parameter vector autoregression with stochastic volatility. See Lubik TA and C Matthes (2015), 'Calculating the Natural Rate of Interest: A Comparison of Two Alternative Approaches', Richmond Federal Reserve Economic Brief No EB15-10.

[6] Goldin I, P Koutroupmpis, F Lafond and J Winkler (2022), 'Why Is Productivity Slowing Down?', Oxford Martin Working Paper Series on Technological and Economic Change No 2022-08.

[7] Andrews D and D Hansell (2019), 'Productivity Enhancing Labour Reallocation in Australia', Australian Treasury Working Paper No 2019-06; Hambur J (2021), 'Product Market Power and Its Implications for the Australian Economy', Australian Treasury Working Paper No 2021-03.

[8] See Jones, B (2021), 'Uncertainty and Risk Aversion – Before and After the Pandemic', Keynote Address at the Minerals Week Australia-Asia Investment Outlook, Canberra, 2 June 2021.

[9] See Mian A, L Straub and A Sufi (2021), 'What Explains the Decline in r*? Rising Income Inequality Versus Demographic Shifts', Becker Friedman Institute for Economics Working Paper No. 2021-104; Rachel L and D Smith (2017), 'Are Low Real Interest Rates Here to Stay?', *International Journal of Central Banking*, 13(3), pp 1–42.

[10] See Productivity Commission (2018), 'Rising Inequality? A Stocktake of the Evidence', Commission Research Paper, August.