



Does the Budget get in the way of good policy?

Sam Reinhardt, Parliamentary Budget Officer,
Address to the Australian Conference of Economists, 8 July 2026

Good afternoon, and thank you to the ACT Branch of the Economic Society of Australia – particularly Bob Breunig and Maria Racionero – for all their work in bringing this ACE conference together. I also add my acknowledgement of the Traditional Custodians of the land on which we meet today, and pay my respects to Elders past and present.

The question I am discussing today is: does the budget get in the way of good policy?

It is a deliberately provocative question. But I want to be clear at the outset: I am not arguing that Australia's budget process is broken. Quite the opposite. Australia has a strong and credible budget framework. Our institutions, processes and budget documents are widely respected internationally — and so are Australia's budget outcomes.

We also have a high level of public engagement with the budget and fiscal sustainability. That matters. Public understanding and scrutiny are part of what helps produce good outcomes.

But even strong systems create incentives. And those incentives shape the way policies are developed, presented and debated.

Today I want to look at the budget not as the policy debate itself, but as one of the main ways that debate is framed. The budget is a bit like a map: incredibly useful, but it's not the actual territory. And sometimes the way the map is drawn — or folded — can make some parts of the landscape easier to see than others.

For most of my talk I'm going to explore eight examples of how the construction or presentation of the budget might create incentives for how policy is designed. Some will be familiar to those who follow budgets closely: the no policy change assumption, the focus on the underlying cash balance, the forward estimates period, and the treatment of new spending. Others are more technical, but still important: the contingency reserve, point estimates and estimates variations, the whole-of-government perspective, and what I will call the 'morning after effect' of winners and losers.

The common theme is not that any of these features are flaws. Each has a purpose. But if we are not careful, they can narrow the debate — shifting attention away from the merits of policy and towards the way those policies appear in the budget.

So, let's spend some time looking over this map we call the budget.

Section 1 – The forward estimates period

I'm going to start with the longer-term impacts of policy and speak about the forward estimates period.

Policy impacts in the budget are primarily shown over the forward estimates – the upcoming budget year plus three years.

As of the late 1980s, budgets were produced on a rolling one-year basis. The four years that our budgets now cover can tell us a great deal more about the path of policy and longer-term costs than a one-year budget.

However, the forward estimates do not always capture the full implications of a policy, particularly those that grow exponentially over time such as programs that are indexed.

Our standard approach at the PBO is to include the medium-term costs for all our costings, this includes those in our Election Commitments Report. We also focus on the medium-term budget baseline in our flagship report, Medium-Term Budget Outlook.

A historical example of the medium-term costs showing a more complete picture is the decision to freeze the indexation of fuel excise from March 2001.

The original budget estimate in 2001-02 was \$2.5 billion over 4 years. If the medium-term costs were included, they would have been around \$18.1 billion over 10 years.

The point here is not to suggest that a large medium-term cost, by itself, implies that a policy lacks merit.

But understanding the cost of a policy over the long run is crucial to an informed debate about the policy.

Section 2 – No policy change assumption

The second example I want to talk about is what sounds like a technical assumption, but has very real consequences for the way we understand the budget: the assumption of no policy change. This assumption sounds banal until you realise it is doing quite a lot of heavy lifting in the budget.

Budgets – Commonwealth, state and territory – are built on the basis that existing policy settings continue, and further policy changes beyond those already announced are not included.

It's a reasonable assumption. A budget can't include the impact of future decisions that governments haven't made yet. Budgets are a statement of current plans, not hypothetical future plans. In practice, that means that when some programs reach the end of their funding period, the budget assumes they end and are not replaced by a similar program

That can introduce a degree of systematic optimism into budget forecasts. The forward estimates can show expenses declining, not because governments have made a clear decision to reduce services or activity, but because the budget does not assume future governments will extend, renew or replace terminating programs.

So, while the no policy change assumption is a sensible and transparent way to prepare a budget, it can still lead the published outlook to understate future expenses relative to what is likely to occur in practice.

I want to be clear about what I am, and am not, saying here. This is not an argument that every program should automatically receive ongoing funding. In some cases, programs end because they were designed to be temporary and have a clear end of life. The funding for a major project ends on completion. In other cases, governments fund a program for a fixed period so that in the future the program can be assessed to determine whether it is working and whether it represents value for money. That discipline is a good thing.

But when we look across the budget as a whole, the cumulative effect matters. If a large number of grant programs are assumed to finish, total grant spending can appear to fall significantly as a share of GDP over the forward estimates. Declines like this are standard in budget projections, but they rarely eventuate in full because governments later announce new programs, extend existing ones, or respond to emerging pressures.

At times in the past, there has also been a related issue with unlegislated measures – decisions that remain in the budget but have not been passed by the Parliament. Where those measures sit in the budget for long periods (years at a time), and there is limited likelihood they will be delivered, they too can make the fiscal outlook look more optimistic than it may ultimately prove to be.

A different, but related, example is personal income tax. The budget assumes no future tax cuts. As wages grow if we assume tax thresholds remain unchanged more income is taxed at higher rates, and personal income tax revenue rises as a share of the economy. On one level, that is exactly what the budget should show: the fiscal position under current policy. But history tells us that governments regularly choose to return some of that fiscal drag through income tax cuts. So even if the budget accurately presents the fiscal position based on the government's policies, it may not be the best guide to what will actually happen over time.

There are also potential upside risks – including things like commodity prices and temporary programs that raise revenue – but these are not systematic upsides. Taking the commodity price example, there is always the potential that these will be underestimates or overestimates and there are plenty of budgets where we have experienced the later – the revenue has been lower than forecast.

This matters for good policy, particularly for longer-term fiscal strategy. We may wrongly assume the budget position will simply improve over time, when this improvement is partly generated by the simple mechanics of the no-policy-change assumption that underpins every budget. If so, then there is a risk that pressure for more structural reform is reduced.

Section 3 – the Underlying Cash Balance (UCB)

My third example is the way the budget debate can be limited if it is focused on one fiscal aggregate to the exclusion of others. In Australia, the fiscal aggregate we usually focus on is the underlying cash balance, or UCB.

The UCB gives Australians a clear and simple indicator of the overall shape of the budget. It supports fiscal discipline, and it provides a measure that can be readily understood and compared over time.

But the simplicity that makes the UCB useful can also make it incomplete.

Some major forms of government support are not captured in the UCB. Government investment funds, equity injections – such as those used for Snowy Hydro, the NBN or Inland Rail – and concessional loans, such as student loans, can all have significant fiscal and economic effects without showing up as conventional spending in the underlying cash balance.

These are sometimes described as 'off-budget' measures. More precisely, they are captured in other fiscal aggregates notably the headline cash balance, which includes loans, equity investments and other financial transactions of government.

There can be very good policy reasons to use these mechanisms, and the concern is not that these instruments are inherently poor policy. These may be effective ways of delivering

government support, but should be assessed on their merits, risks and full fiscal implications. A large investment fund, or billions of dollars in equity or concessional loans, may have little immediate impact on the UCB – and in some cases may even improve it – despite involving real fiscal risks and real economic effects.

Those risks and costs matter. They can include the risk that future returns to government are lower than expected, the repayment of debt required to finance the investment, and the broader effects of government intervention in markets. In some cases, the UCB impact may not reflect the fiscal impulse or the economic substance of the policy.

The overall health of the government's balance sheet should also remain a focus. The amount of debt on issue matters both because it needs to be repaid in future budgets, but also because the balance sheet is a reflection of the ability to respond to future shocks.

We as economists, analysts and commentators have an important role to play. We need to look beyond the headline or underlying cash measure and ask what the policy is trying to achieve, what risks it creates, and what it costs in economic and fiscal terms.

The broader point is that no single metric can capture every aspect of the fiscal position. That is why the PBO will continue to report on, and provide information to help people understand, the budget beyond the cash balance.

This is not a new issue in Australia. In general, the HCB and UCB series move together. In recent years the HCB has come in lower than the UCB, reflecting increased investment in financial assets for policy purposes. The highlighted areas show periods where the HCB was higher than the UCB. This primarily reflects the sale of government assets during this time, specifically the Commonwealth Bank, Qantas and Telstra.

For the budget nerds in the audience, and those who were actually there at the time: the UCB has been reported only since the 1990s. The main budget balance had been the *headline* cash balance for a very long time up until then. But the sale of government assets during this time, added a lot of revenue to the budget balance (the HCB at the time). The critique was that the HCB presented the budget too favourably, given these assets could be sold only once.

The underlying cash balance did not count those asset sales, so it was considered to present a better ongoing measure of the state of the budget.

It's interesting that currently some commentators critique the UCB as being too optimistic about the budget, and the HCB is suggested to be the better measure of the 'true' picture.

As an aside, while the UCB wasn't reported before the 1990s, the historical tables in the back of budget paper 1 show what it would have been back to 1970, and the PBO has calculated it back to 1953. You can also find it in our *Historical fiscal data* spreadsheets.

Going back even earlier than that is more difficult, but we haven't given up yet and if anyone here is keen to trawl through old budgets to try to work it out then please talk to the PBO team later!

Section 4 – Taking a Commonwealth, state and territory focus

Stepping back from the budget details, the fourth example I want to discuss is the need to look at budgets from a whole-of-government perspective.

The Commonwealth budget attracts the most attention in national debates, but state and territory budgets are also central to understanding Australia's overall fiscal position, as well as in designing good policy.

As PBO's soon-to-be-released *Medium-Term Budget Outlook* publication will show, recent growth in total state gross debt presents a fiscal risk that warrants monitoring, with state debt increasing faster than Commonwealth debt.

The PBO has made a small contribution to this understanding through our annual publication called the *National Fiscal Outlook*, which adds up all the budget forecasts across the different governments to present a national picture.

Since the early 2000s total state debt has been a small fraction of total debt, but in recent years the proportion of state debt is increasing. This is expected to continue over the forecast period.

The International Monetary Fund noted in its latest assessment of Australia that, 'as the Commonwealth is viewed as a de facto guarantor of state debt by some credit rating agencies, higher sub-national debt could eventually impact Commonwealth borrowing costs.'

Considering both Commonwealth and state budgets also matters for policy design because many of the most important areas of public policy — including health and education — are shared responsibilities between the Commonwealth and the states and territories. The fiscal impacts are also shared.

That can shift the debate away from the merits of the policy itself and towards questions about which level of government pays, how much it pays, and whether the funding split is fair.

It can also make information and data sharing across jurisdictions more difficult. If governments are concerned that information may be used in a way that affects funding negotiations, they may be less willing to share the data needed for high-quality policy development.

So, if we want to understand and address the major challenges facing the country, we need to look beyond any single budget and take a genuinely whole-of-government view.

Section 5 – The 'morning after effect' (winners and losers)

A fifth example is something that has become a familiar feature of Commonwealth, state and territory budgets: the budget-night analysis of winners and losers.

By the morning after the budget, Australians are sorted into groups and told whether they are better off, worse off or unaffected by the measures announced the night before.

Given current cost-of-living pressures, that is understandable. Individuals and families want to know what support they will receive, and how quickly it will arrive.

But a focus on the immediate impact can obscure the longer-term and intergenerational effects of policy.

Part of the issue is the limits of modelling.

It is much easier to model what we call the 'morning-after effect': the immediate distributional impact, before the longer-term consequences become clear.

Microsimulation models allow us to estimate detailed distributional impacts with some confidence. They might show that an average 45-year-old would be \$1000 better off from 1 July.

But estimating how that same person will be affected in 10 years' time depends on many factors: their income, employment, educational and training opportunities and the broader economic environment.

There are also limits in the data available for this kind of analysis.

Income data is often more readily available than wealth data, which can be much harder to collect and, importantly, harder to match with income data.

As a result, policy analysis often focuses on a person's current income — a snapshot at a point in time — rather than their overall wealth, or their lifetime income and wealth. Measures of wealth can capture accumulated advantage or disadvantage over time in a way that annual income often cannot.

That difference matters. If the debate focuses only on what happens immediately, broader reforms with long-term or intergenerational benefits become much harder to explain, and much harder to sustain.

Large policy changes also often require governments to manage the position of people who are, or are perceived to be, immediately affected.

The introduction of the GST is a useful example. It was accompanied by personal income tax cuts and increases in social security payments, designed to offset the impact on households.

Those compensation arrangements were put in place before further negotiations on GST exemptions, which were later needed to secure passage of the legislation.

More than two decades later, questions remain about whether those exemptions have made the system more complex.

There are two points here. First, income measures do not always tell the full story of who is better or worse off particularly over time. Second, compensation to offset losses can be costly. In the case of the GST, that compensation was made easier, in part, by the stronger revenue conditions of the period.

That raises difficult questions. Is substantial compensation a necessary condition for major reform? And in a world of increasing fiscal pressure, can governments afford to offset every reform with large windfalls for those who may be worse off, even temporarily?

Or do we need to get better at analysing winners and losers in a more holistic way — one that captures intergenerational impacts, lifetime and wealth effects? If we can, we may be able to shift the narrative away from immediate gains and losses, and towards the broader question of whether a policy leaves the community better off over time.

Section 6 – The contingency reserve and decisions taken but not yet announced

Let me turn now to a more technical, but important, part of the budget: the contingency reserve.

The contingency reserve is an allowance within the budget for policy decisions or other expected costs that cannot yet be allocated to a specific program, or cannot be published separately.

Traditionally, the largest component is the conservative bias allowance. This recognises that the costs of government programs often turn out to be higher than initially estimated.

For example, if a major infrastructure project is expected to cost \$1 billion, the conservative bias allowance will include a buffer for possible cost increases. That allowance is reduced as the project comes closer in time, reflecting the fact that cost uncertainty falls over time. It is based on experience with specific projects and programs, rather than a broad judgement about the budget as a whole.

Other parts of the contingency reserve cover items that are too sensitive to disclose separately, including commercial-in-confidence matters or national security issues.

It also includes decisions taken by government but not yet announced — often referred to as DTBNYA. These are identified in the budget as a group, but not item by item.

There are good reasons for this. Governments may judge it prudent to set aside funding for a likely future cost before the details are finalised, or before negotiations are complete. For example, this might include future state programs where the broad funding envelope is expected, but the final amount has not yet been agreed.

For Commonwealth governments, this can be a difficult balance. If they do not make provision, they may be criticised for failing to put money aside. If they do make provision, they are asked how much has been set aside, what it is for, and how it compares with any final amount or agreement made. I do not want to understate how difficult those dynamics can be to manage.

Taken together, the conservative bias allowance, amounts not for publication, and decisions taken but not yet announced are all reported in aggregate.

A transparency challenge arises because, once a measure has first been funded through DTBNYA and is later announced, the original cost profile cannot generally be traced through future budget documents. That can make it harder for parliamentarians, analysts and the public to follow the fiscal impact of a decision that comes into the budget as a DTBNYA.

Section 7 – Point estimates and estimates variations

The next area I want to talk about is policy estimates, estimates variations and adjustments to program costs — in other words, what happens when the expected cost of a program changes, not because of a new government decision, but because the world has turned out differently from what was forecast.

These are a standard feature of the budget. When economic conditions change, when demand is higher or lower than expected, or when other program cost drivers move, departments update their forecasts. Those changes are recorded as estimates variations.

And as an agency that does a lot of costings, we know costings can be hard. They are especially challenging for programs in new, innovative or fast-changing areas.

Some policies are relatively straightforward to estimate. Changes to personal income tax rates or social security payments, for example, are built on large administrative datasets and well-established behavioural assumptions. But programs in emerging areas, or policies that create new tax bases, can be much harder to estimate reliably.

Governments report significant estimates variations, but smaller variations can still tell us a lot. Over time, they can show whether a program is tracking broadly as expected, or whether its costs are moving in a different direction.

That kind of information can be very valuable. A policy that looks like an efficient use of public money at \$100 million a year may look quite different if it grows to \$300 million a year.

Short of reporting every program the government runs in very fine detail, there is no simple solution to this. But it does mean that analysts, commentators and parliamentarians need to pay attention not only to new measures, but also to how existing programs are evolving over time.

NDIS costings are a case in point: early on, forecasts suggested the scheme would grow faster than it did – it cost less than expected. But as the scheme matured, each forecast expected growth to slow. As we know, that did not eventuate, and the scheme has grown significantly. More recent forecasts now reflect the government’s planned policy changes.

This example illustrates a broader point: costings are, by their nature, point-in-time estimates. When we prepare costings in the PBO, the two main sources of uncertainty are future economic conditions and the way people, businesses or institutions respond to the policy.

Every costing involves assumptions about both. And we know from experience that those assumptions can make a very large difference. Costings are estimates, not prophecies. If they were prophecies, the PBO would need a very different recruitment strategy.

The budget documents themselves are not meant to be predictions of the future. They are planning documents. They describe the government’s fiscal plan based on the policy settings and economic conditions known at the time. But to build any plan, you need a starting point — and that is where forecasts and costings come in.

It is a time-honoured budget game to look back and point out forecast errors. But in many cases, those changes are not the result of poor work. They reflect economic events, shifts in demand, or the inherent difficulty of estimating how people will respond to a new policy.

One way to deal with uncertainty is to present a range of estimates. But that can sit uneasily with the budget process, which often requires a single number. It is hard for a government to build a fiscal plan if a costing says the cost could be anywhere between \$100 to \$300 million a year.

There is a line on this theme attributed to the Texan rancher and former US President Lyndon B. Johnson. When presented with estimates in a likely range of values, he is supposed to have replied: ‘Ranges are for cattle. Give me a number.’

That instinct is understandable. Budgets need numbers. Decisions need numbers. But good policy also needs an appreciation of uncertainty. The challenge is not to pretend that estimates are precise, but to understand what drives them, how they change over time, and when those changes should prompt a fresh look at the policy itself.

Section 8 – Marginal spend and the overall budget

The last point I want to speak to is related to the previous one. The budget can narrow the policy debate by focusing attention on the marginal change — the new measures announced on budget night — rather than on the full stock of spending already built into the budget.

Since the 2018-19 Budget, excluding the COVID budgets, new budget measures to increase expenditure amounted to around 1 per cent of total payments across the forward estimates.

To put that in more familiar terms: if you were paying \$2000 a fortnight on your mortgage, 1 per cent would be less than the cost of a couple of extra coffees each week.

That is not to say new measures do not matter. They are the choices that have just been made, and they are the most recent decisions affecting the budget and therefore Australians. But if we

focus only on those choices, we can miss the much bigger question of where public money is already going.

The larger part of the budget is the spending that is already baked in: existing programs, existing entitlements, existing agreements and existing policy settings. It is not feasible, or even desirable, for every budget to reopen the entire spending base from first principles. But we should not lose sight of it. If we only focus on new measures, we risk giving far less scrutiny to the much larger set of commitments that continue from year to year.

Programs, once established, can also be very hard to unwind. They create expectations, delivery systems, stakeholder groups and, often, legitimate reliance by the people and organisations they support. So even where circumstances change, or where a program no longer represents the best use of public money, the practical and political barriers to change can be significant.

A related issue is tax expenditure, or revenue forgone. In substance, revenue forgone can operate very much like direct spending: it provides a benefit, it has a fiscal cost, and it reflects a policy choice. New measures generally show the revenue impact in the budget, while longstanding arrangements are reported separately through the Tax Expenditures and Insights Statement.

But if we are trying to understand the full scale of government support in a policy area, these measures need to be considered alongside direct spending. Economically, giving someone a dollar through a grant and giving up a dollar of tax revenue can have much the same budgetary effect.

The point, again, is not that new spending should receive less scrutiny. It should. But good budget analysis also requires us to look at the base: the programs, concessions and commitments that carry forward each year, often with far less attention.

Section 9 – Conclusion

So, does the budget get in the way of good policy? Not inherently—but the way we frame, measure and communicate the budget can shape the policy choices we make.

Across these examples, a common theme emerges: incentives matter and good-quality budget information and analysis matters. Whether it is the no policy change assumption creating optimism bias, the dominance of the underlying cash balance narrowing the lens, looking across levels of government, the focus on winners and losers, the short horizon of the forward estimates, or the role of the contingency reserve—each can subtly steer attention away from policy merit and toward budget presentation.

None of these features are flaws in isolation. In many cases, they reflect practical trade-offs: simplicity versus completeness, transparency versus confidentiality, certainty versus realism. Australia's budget framework is, by international standards, robust and credible. But a strong system still requires informed users.

The features and incentives of the budget rules and presentation can impact on policy design.

In a similar way to how we understand Goodhart's law: 'when a measure becomes a target, it ceases to be a good measure,' if we are not careful with how we assess the budget, we will create problematic targets and damage our measures.

For the budget to support good policy, we need to look beyond the headline numbers. That means:

- interrogating the assumptions that drive the forecasts
- considering impacts across the full fiscal balance sheet, not just the UCB
- taking a whole-of-government perspective
- looking beyond immediate distributional effects to long-term and intergenerational outcomes
- and placing greater weight on medium-term costs and risks.

Ultimately, the budget should be a tool to inform better decisions, not a constraint that narrows them. That requires ongoing transparency, careful interpretation, and a willingness from policymakers, analysts and commentators to engage with the complexity behind the numbers.

If we do that well, the budget won't get in the way of good policy—it will help us achieve it, ensuring that, 'the budget is not just a collection of numbers, but an expression of our values and aspirations.'

That is why the way we measure, present and debate the budget matters.