e61 Appendix | 19 February 2026

APPENDIX: FISCAL SUSTAINABILITY IN AUSTRALIAL: INSIGHTS FROM CONSOLIDATED GOVERNMENT ACCOUNTS

Jack Buckley, Matthew Maltman, and Matt Nolan

Example summary box

This is the Online Appendix for the e6-McKinnon Fiscal Sustainability Reports. The content covered is:

- A discussion of revenue raising and why it is not the focus of the report.
- Additional contextual graphs on health and education to support the report.
- An explanation of the datasets used and why.
- And explanation of the methods applied for cross-country trend comparisons (Synthetic Control), clustering of spending functions, the simulation of a bond market shock, and decomposition of spending drivers (Shapley Decomposition).

Further information about these exercises is available on request, and future notes will provide additional context on these scenarios.

Additional Graphs and context

1. Revenue context

The revenue challenge for state and federal governments is complex, and will be discussed further in future work. The high level challenge ahead is highlighted below, providing a motivation for the main reports focus on spending pressures and patterns as a means to restore fiscal balance.

Federal Budget 2025-2026 projections suggest that – even with cuts to the second tax threshold – the deficit will be closed via bracket creep. This is the reduction in the real value of the tax-thresholds, thereby pushing up the average tax rate paid by the average worker.

Figure 1: ETR by Gross Taxable Income

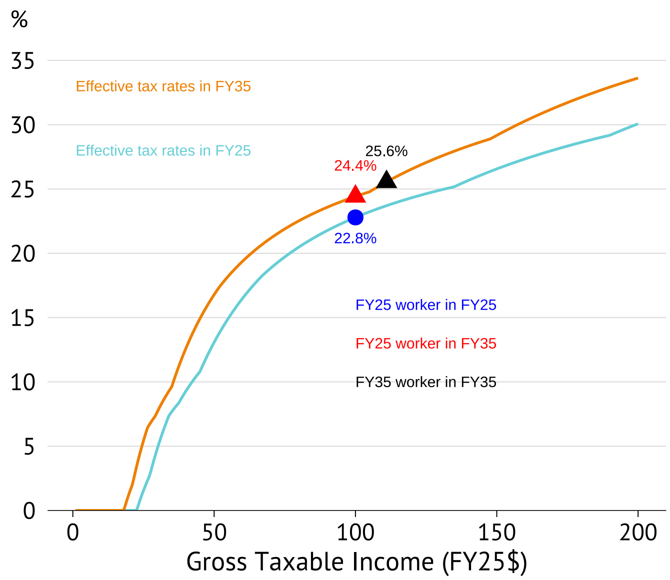

* Markers for Avg Earnings in FY25 (red), FY35 (black) and extra (blue)
** Deflated line assumes 2.5% inflation for 10 years
*** Effective Tax Rate includes Single Medicare Levy but excludes the Medicare Levy Surcharge.

However, this is both an inefficient and inequitable way to raise the tax revenue needed to meet current spending projections. The inefficiency is reflected in the marginal excess burden (MEB) of raising an additional dollar of tax revenue. This is the value of lost economic activity when a dollar is taken from a private individual and given to government, due to private actors switching to otherwise less valuable alternatives. In Australia KPMG (2010) estimates offer the most comprehensive evaluation of relative MEB's across the tax system – with transaction taxes such as Stamp Duty very inefficient, while rent taxes has as land tax and the resource rent tax are very efficient.

Figure 2: Marginal Excess Burdens

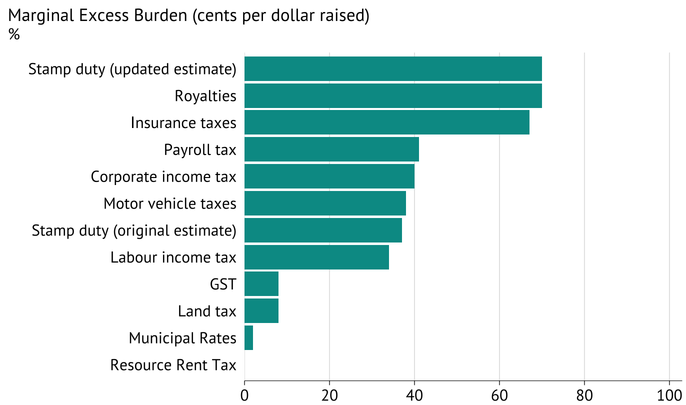

* Estimates from KPMG (2010) report for the Australia's Future Tax System Review.
** Updated Stamp Duty Estimates used by PBO come from Cao et al (2015)
Sources: KPMG Econtech; PBO

More recently Nassios and Giesecke (2025) estimated a MEB – at current tax rates – for income taxes and Goods and Services Tax (GST) that is roughly equivalent. The approximate 24% MEB means that for every dollar raised from these tax bases costs the private sector $1.24. In other words, an additional 24 cents of valuable economic activity is wasted for each additional dollar raised.

These bases are relatively more efficient than some of the tax bases that state governments rely on for raising additional revenue – such as stamp duty.

The marginal excess burden found for the GST is relatively high compared with international literature. The reason for this is the relatively narrow base of Australia's GST system – with approximately half of consumer spending covered by the GST, compared to 58% across the OECD, and 96% in countries like New Zealand. Instead of increasing tax rates, broadening these bases may be an effective mechanism to raise revenue at a low economic cost.

However, even with these higher excess burdens the work of Nassios and Giesecke highlight that there would still be a significant economic benefit in reducing taxes like stamp duty, and replacing them with a higher rate of tax on GST or income. Part of the challenge of reforming away from current inefficient taxes – like stamp duty – towards a broader more efficient tax is the coordination around revenues across government. Taxes like stamp duty are received by the states, while the income tax is collected and used by the Federal Government.

Reluctance at both the state and federal level to significantly change the tax system suggests that the ability to raise near-term revenue may be constrained. Adding to this constraint is the growing share of income tax as a proportion of national income. By 2024 income tax revenue was equivalent to 18.7% of GDP – the highest share of national income taken in income tax in history. PBO projections indicate that this income tax share will remain steady in the next five years, as the increase in tax revenue from fiscal drag counteracts the decline in revenue from weaker nominal growth. Following that, income tax burdens will continue rising – the main mechanism that will reduce the government budget deficit.

Figure 3: Income tax trends

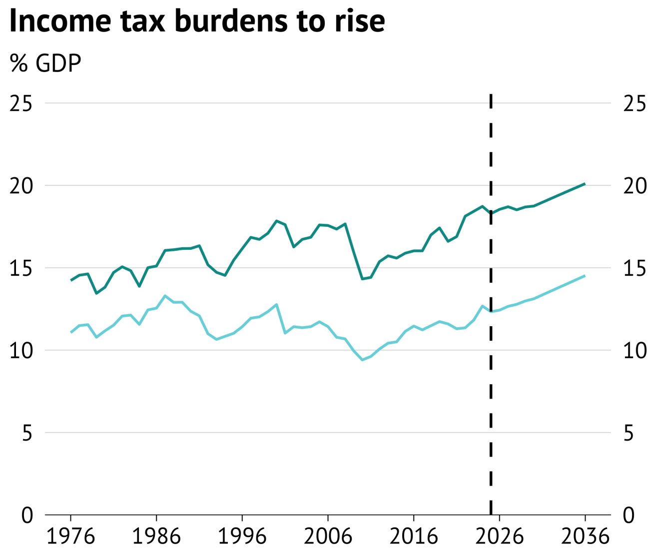

* Until 2000, cash accounting concepts are used. Post 2000 income tax is accrual based.
** Projections sourced from PBO Medium Term Budget Outlook, with an allowance for non-individual or corporate income tax.
*** Total income tax includes individual, corporate, fringe benefit, resource rent, and superannuation taxes.
Sources: e61; PBO

In an environment where tax reform is challenging and current tax burdens are historically high, it is important to ask whether spending restraint may be more appropriate than increased taxes. For this reason the main report focused on government spending in order to consider relative value for money and inform this question.

2. Health

In this section we outline additional graphical facts regarding consolidated health spending, which are used to support claims made in the main report.

Figure 4: Real and % GDP health trends

Consolidated health spending growth

Growth in real spending per capita and spending as a share of GDP.

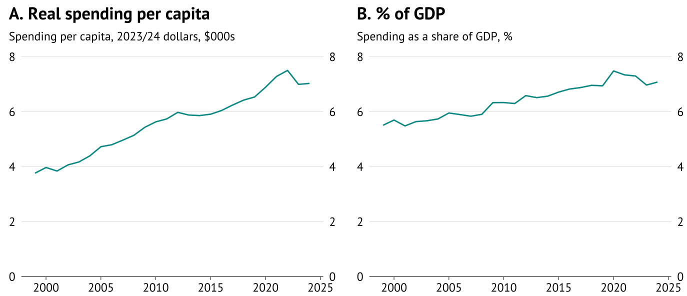

Sources: ABS; e61

Figure 5: Expense trends

Consolidated health expenditure by area and type

Spending as a share of GDP

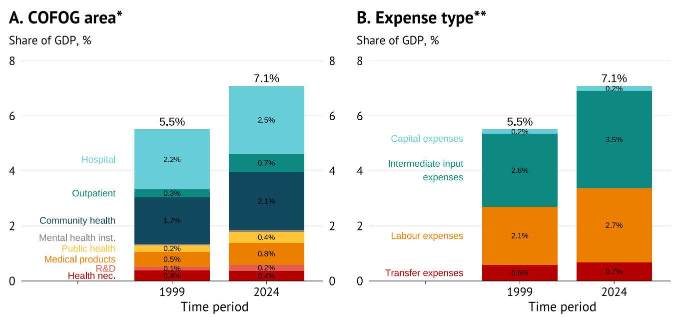

* This chart reports consolidated government expenditure broken down according to the Classification of the Functions of Government (COFOG) areas.
** Capital expenses include ETF classes depreciation (ETF 124); capital transfer expenses (ETF 126); interest expenses (ETF 127); and other property expenses (ETF 128). Intermediate input expenses are defined as non-employee expenses (ETF 123), which includes use of goods and services, social benefits to households in goods and services and production tax expenses. Labour expenses include superannuation and other employee expenses (ETF 121 and 122). Transfer expenses are defined as current transfer expenses (ETF 125), which includes grants, subsidies and monetary transfers to households.
Sources: ABS; e61

Figure 6: Expense shares

Consolidated health expenditure by area and type

Share of expenditure on health

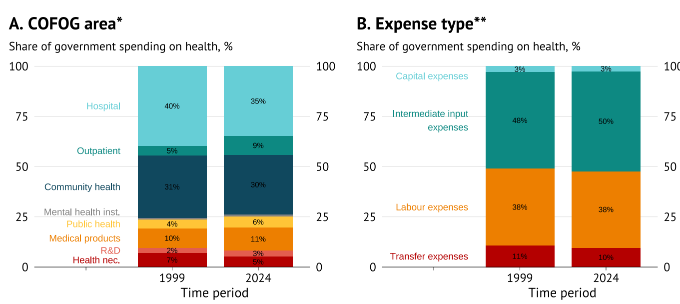

* This chart reports consolidated government expenditure broken down according to the Classification of the Functions of Government (COFOG) areas.
** Capital expenses include ETF classes depreciation (ETF 124); capital transfer expenses (ETF 126); interest expenses (ETF 127); and other property expenses (ETF 128). Intermediate input expenses are defined as non-employee expenses (ETF 123), which includes use of goods and services, social benefits to households in goods and services and production tax expenses. Labour expenses include superannuation and other employee expenses (ETF 121 and 122). Transfer expenses are defined as current transfer expenses (ETF 125), which includes grants, subsidies and monetary transfers to households.
Sources: ABS; e61

Figure 7: Health age adjustment

Age driven spending growth

Drivers of spending growth since 2000: composition vs intensity

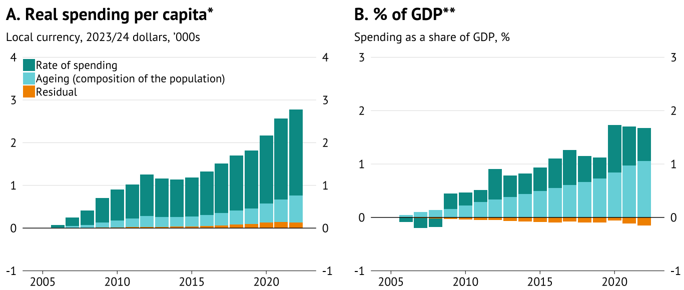

* This chart uses a shift-share decomposition to separate changes in age-adjusted outcomes into: (i) a composition effect due to changes in the population age structure (holding age-specific rates fixed at 2000 levels), (ii) a rate effect due to changes in age-specific rates (holding the 2000 population structure fixed), and (iii) a residual capturing interaction effects.
** This chart uses a shift-share decomposition to separate changes in age-adjusted outcomes into: (i) a composition effect due to changes in the population age structure (holding age-specific rates (% GDP per % pop) fixed at 2000 levels), (ii) a rate effect due to changes in age-specific rates (holding the 2000 population structure fixed), and (iii) a residual capturing interaction effects.
Sources: ABS; e61 Institute

Figure 8: International health spending

Government spending on health as a share of GDP

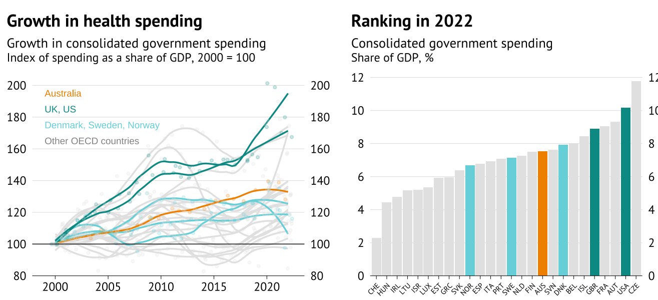

Sources: e61; OECD

Figure 9: International preventative healthcare spending

Spending on primary and preventative healthcare

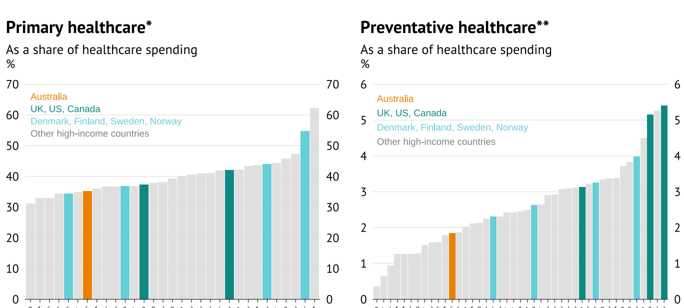

* According to the WHO, Primary Health Care is a whole-of-government and whole-of-society approach to health that aims to ensure the highest possible health and well-being for people and communities and the equitable distribution of health services. Furhter detailed can be found at: https://www.who.int/news-room/fact-sheets/detail/primary-health-care.
** Preventative healthcare includes spending on information, education and counseling programmes; immunization Programmes; early disease detection programmes; healthy condition monitoring programmes; epidemiological surveillance and risk and disease control programmes; preparing for disaster and emergency response programmes
Sources: e61 Institute; WHO

Figure 10: International preventative healthcare spending 2

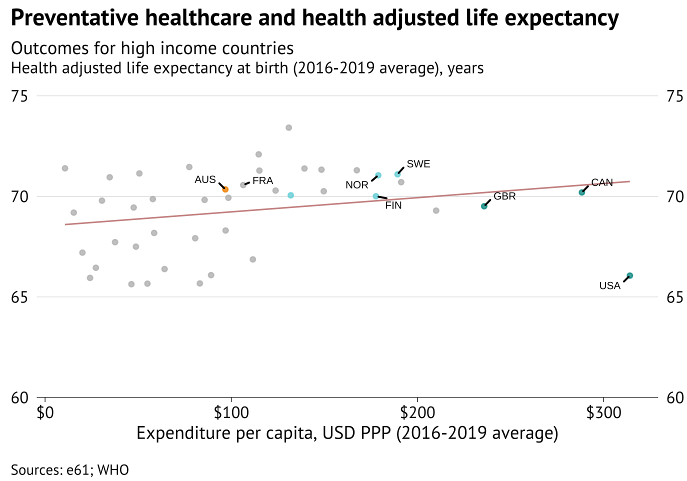

3. Education

Figure 11: Real and %GDP Education trends

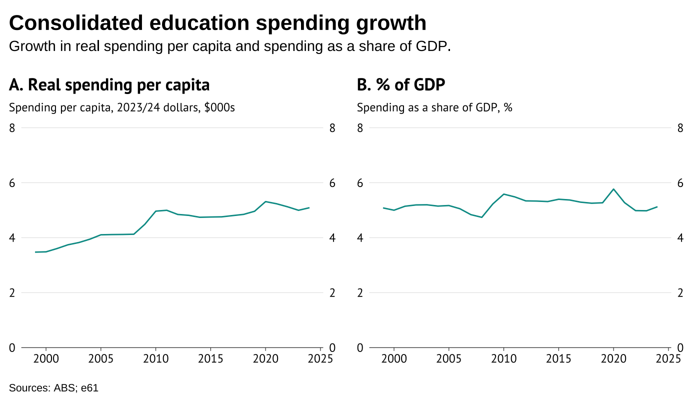

Figure 12: International preventative healthcare spending 2

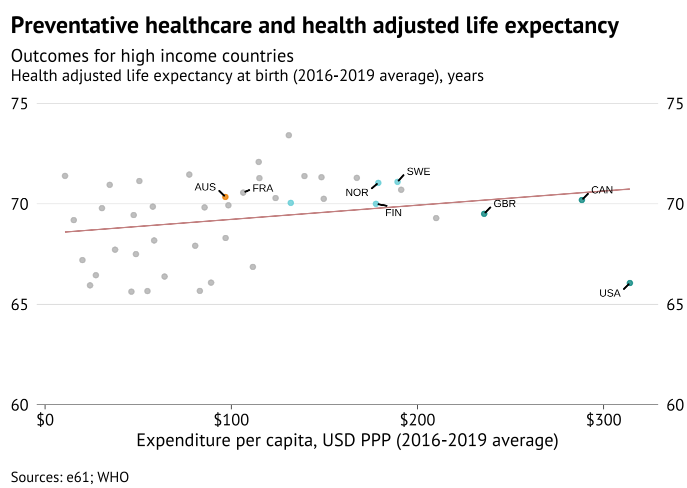

Data and methods

1. Data used

Four key data sources are used in this report to comment on government spending trajectories: The PBO's collection of information from historic budget documents, the ABS Government Finance Statistics, the OECD Fiscal Decentralisation Database, and ABS National Accounts measures of general government expenditures.

Each of these measures reflects a different lens on spending, as described below.

The ABS Government Finance Statistics (GFS) used are a disaggregated version of Table 939 on the ABS website, which were requested from the ABS. This reflects general government expenditures –- thereby not counting public non-financial and financial entities that operate at arms-length.

We use the ABS GFS as the gold standard measure of the functions where spending occurs, and to provide aggregates about spending at the general government level. This data is separated into current expenses and net acquisition of non-financial assets. When added together this gives total expenditure.

The OECD Fiscal Decentalisation Database provides consolidated statistics on general government expenditure based on two classifications –- one that focuses on the level of government spending the funds and another based on the entity that provides funding. We focus on the *spent by* measure in order to make federal vs non-federal comparisons. This is only produced on a total expenditure basis, but provides the only information that can be compared across countries.

The ABS National Accounts figures provide a longer time series than most figures which can be used to evaluate long-term trends in government consumption, investment, and transfers to households. However, they are measured on a different basis than the accounting concepts used for government financial accounts which means these are generally used as a secondary source of analysis where the longer time series is required.

For aggregated consolidated statistics and projections, we use Budget 2026 and PBO National Fiscal Outlook 2026 consolidated data – with this data based on roughly the same general government expenditure definitions as those used by the ABS. Projections from the ABS and Australian Treasury Intergenerational report are also used for population and economic variables, as well as some projections.

2. Synthetic Control Approach

The synthetic control (SC) method estimates what Australia's government expenditure-to-GDP ratio (E/GDP) would have been in the absence of Australia-specific shocks after 2007. The approach constructs a weighted combination of other OECD countries -— forming a "synthetic Australia" that closely replicates Australia's pre-2008 behaviour.

The dependent variable is nominal general-government total expenditure as a share of nominal GDP, sourced from the OECD consolidated general-government dataset. To remove level differences, each country's series is indexed to 1 in 1999. Because Australia's fiscal year ends in June, observations for Australia are adjusted using a half-year averaging approach (half of FY_t and half of FY_{t+1}).

2.1 Model setup

- Target unit: Australia
- Intervention year: 2007
- Pre-intervention period: 1999–2007
- Donor pool: All available OECD economies with continuous data from 1999 onward, excluding clear outliers (e.g., Ireland).

The SC algorithm selects non-negative donor weights (summing to one) that minimise the mean squared prediction error (MSPE) between Australia and the weighted average of donors in the pre-intervention window. The resulting synthetic path provides a counterfactual trajectory of E/GDP after 2007 if Australian had continued to move with similar countries.

The dataset used for this are harmonised OECD spending measures from Organisation for Economic Co-operation and Development (OECD) (2026).

Setup and notation

Let $j = 1, \ldots, J+1$ index countries, where $j = 1$ denotes Australia and $j = 2, \ldots, J+1$ is the donor pool of OECD economies with continuous expenditure data from 1999 onward (excluding Ireland). [1] Let $T_0 = 9$ denote the number of pre-intervention years (1999–2007) and T the total number of periods (1999–2022).

The outcome variable is the ratio of nominal general-government total expenditure to nominal GDP, indexed to unity in the base year:

$$Y_{jt} = \frac{E_{jt}}{G_{jt}} \Big/ \frac{E_{j,1999}}{G_{j,1999}} \tag{1}$$

where E_{jt} and G_{jt} denote nominal expenditure and nominal GDP for country j in year t. Indexing to 1999 removes cross-country level differences, so that the synthetic control exercise identifies changes in spending trajectories rather than differences in spending levels.

Fiscal-year adjustment for Australia

Australia's fiscal year ends on 30 June, so an observation labelled year t in the raw data spans the period July $t-1$ to June t. To align with the calendar-year observations of donor countries, the Australian series is re-expressed using a half-year averaging rule:

$$Y_{1t}^{\text{adj}} = \tfrac{1}{2}\,Y_{1t} + \tfrac{1}{2}\,Y_{1,t+1} \tag{2}$$

That is, the calendar-year t value assigned to Australia is the simple average of the raw observations for years t and $t+1$. All subsequent analysis uses Y_{1t}^{adj} in place of Y_{1t} for Australia.

The synthetic control estimator

For country $j = 1$ (Australia), define two potential paths for Australia: Y_{1t}^{I}, the outcome observed under the post-2007 Australia-specific shocks, and Y_{1t}^{N}, the counterfactual outcome that would have been observed if Australia had continued to evolve in a similar way to other countries on the basis of demographics. The change following the GFC is then given by:

$$\alpha_{1t} = Y_{1t}^{I} - Y_{1t}^{N} \tag{3}$$

However, Y_{1t}^{N} has to be estimated based on pre-GFC data. The synthetic control method estimates it as a convex combination of donor-country outcomes:

$$\hat{Y}_{1t}^{N} = \sum_{j=2}^{J+1} w_j^{*}\, Y_{jt} \tag{4}$$

where the weights $\mathbf{W}^{*} = (w_2^{*}, \ldots, w_{J+1}^{*})'$ satisfy the constraints:

$$w_j \geq 0 \quad \forall j, \qquad \sum_{j=2}^{J+1} w_j = 1 \tag{5}$$

The post-GFC gap is estimated by:

$$\hat{\alpha}_{1t} = Y_{1t} - \hat{Y}_{1t}^{N} = Y_{1t} - \sum_{j=2}^{J+1} w_j^{*}\, Y_{jt}, \qquad t > T_0 \tag{6}$$

1 Ireland is excluded from the donor pool as a clear outlier, owing to the distorting effect of multinational profit-shifting on its nominal GDP series.

A positive gap indicates that Australia's expenditure-to-GDP index exceeded what would have been predicted by the trajectory of comparable OECD economies after 2007.

This should not be viewed as a causal estimate of the effect of the Global Financial Crisis on debt ratios in Australia. Donor countries were also *treated* by the same global shock at that time. Instead this should be treated as a descriptive exercise to understand if Australian government expenditure trends exceeded those that would be expected given the trends experienced in other similar countries.

Weight selection

Let $\mathbf{X}_1$ be a $k \times 1$ vector of pre-intervention predictor values for Australia and $\mathbf{X}_0$ be the corresponding $k \times J$ matrix for the donor pool. The optimal weights solve:

$$\mathbf{W}^* = \arg\min_{\mathbf{W}} \|\mathbf{X}_1 - \mathbf{X}_0\,\mathbf{W}\|_{\mathbf{V}} \tag{7}$$

subject to (5), where the $\mathbf{V}$-norm is:

$$\|\mathbf{X}_1 - \mathbf{X}_0\,\mathbf{W}\|_{\mathbf{V}} = \left(\mathbf{X}_1 - \mathbf{X}_0\,\mathbf{W}\right)'\mathbf{V}\left(\mathbf{X}_1 - \mathbf{X}_0\,\mathbf{W}\right) \tag{8}$$

$\mathbf{V}$ is a $k \times k$ symmetric positive semi-definite matrix of predictor importance weights. Following Abadie et al. (2010), $\mathbf{V}$ is itself selected by a nested optimisation: for each candidate $\mathbf{V}$, the inner problem (7) yields $\mathbf{W}^*(\mathbf{V})$; the outer problem then chooses $\mathbf{V}$ to minimise the pre-intervention mean squared prediction error (MSPE) on the outcome path:

$$\mathbf{V}^* = \arg\min_{\mathbf{V}} \frac{1}{T_0}\sum_{t=1}^{T_0}\left(Y_{1t} - \hat{Y}^N_{1t}(\mathbf{W}^*(\mathbf{V}))\right)^2 \tag{9}$$

This nested structure is implemented directly by the `tidysynth` package in R.

2.2 Application

Two specific specifications were estimated that provided similar results.

1. Baseline model – Predictors: average pre-2007 E/GDP index.
2. Demographics-augmented model – Adds the average pre-2007 population shares aged 65–74 and 75+, capturing ageing profiles across countries.

Baseline model. A single predictor: the average pre-intervention value of the outcome index over the window $t \in [1999, 2007]$:

$$\mathbf{X}_1^{\text{base}} = \left(\bar{Y}_{1,\,\text{pre}}\right), \qquad \bar{Y}_{j,\,\text{pre}} = \frac{1}{T_0}\sum_{t=1999}^{2007} Y_{jt} \tag{10}$$

Demographics-augmented model. Three predictors: the average pre-intervention outcome index plus average pre-intervention population shares for the 65–74 and 75+ age groups:

$$\mathbf{X}_1^{\text{aug}} = \left(\bar{Y}_{1,\,\text{pre}}, \quad \bar{s}^{65\text{-}74}_{1,\,\text{pre}}, \quad \bar{s}^{75+}_{1,\,\text{pre}}\right)' \tag{11}$$

where $s^{65\text{-}74}_{jt}$ and s^{75+}_{jt} are the shares of country j's population aged 65–74 and 75 or above in year t, sourced from the OECD Population Statistics, and bars again denote pre-intervention averages over 1999–2007. The motivation for including these controls is that ageing population structure is a primary driver of health and social protection spending; controlling for pre-intervention demographic profiles reduces the risk that the synthetic control conflates Australia-specific fiscal policy choices with demographic trends common across OECD countries.

Both specifications use the same donor pool and time window. The augmented version improves pre-intervention fit and tests whether demographic structure accounts for part of the post-2007 divergence. For this reason, we have presented the augmented model in the text.

To interpret the synthetic control exercise we can focus on the gap between observed and synthetic Australia after 2007 measures the extent to which spending diverged from the counterfactual trend implied by other advanced economies.

The synthetic Australia is built based on the weighted average value of a set of donor countries. Donor weights show which countries most closely match Australia's pre-intervention path (e.g., United States, Norway, New Zealand, Israel, Iceland). High weights reflect similarity, not causality.

This approach is used to consider overall general government expenditure, as well as consolidated expenditure on health and education. In each case the demographic adjusted model is shown in the main text.

2.2.1 Donor weights of exercises

Below we illustrate the donor weights for each of the three synthetic control exercises in the report.

Figure 13: Donor countries and weights: Aggregate

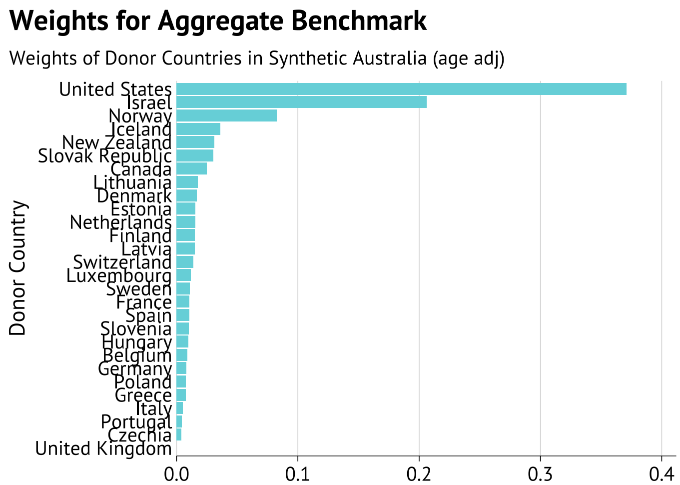

Figure 14: Donor countries and weights: Health

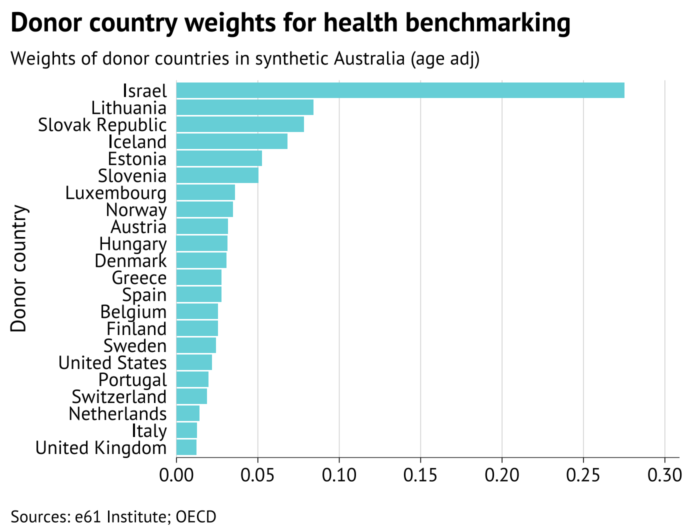

Figure 15: Donor countries and weights: Education

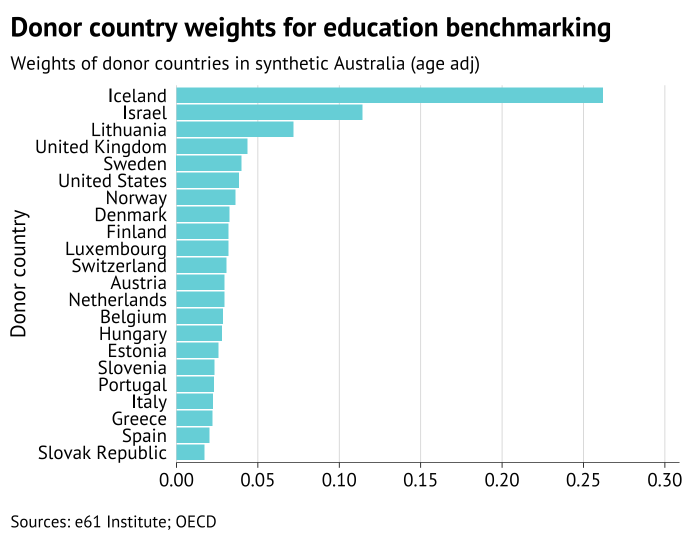

3. Cluster Analysis

Cluster analysis is used to understand groups of spending functions that tend to comove in the data. Given this clustering we are able to establish a broader set of functions that appear to have similar spending drivers. We use ABS Government Finance Statistics (GFS) data on consolidated general government spending by COFOG function from 1999 onward. The spending variable considered reflects current spending and so excludes volatile net capital expenditures. The data are aggregated by subfunction and converted into a time-series panel where each row represents a function and each column a financial year. We apply a hierarchical cluster analysis using features extracted from each function's spending trajectory:

- Mean and volatility of spending (average and standard deviation)
- Skewness and kurtosis (asymmetry and tail behaviour)
- Autocorrelation (persistence from year to year)
- Trend slope (average rate of increase)

Each function is therefore summarised by a six-dimensional feature vector describing its dynamics. We then:

1. Standardise these features to remove magnitude effects;
2. Compute a Euclidean distance matrix between functions;
3. Apply Ward's hierarchical clustering to determine the clusters, which groups categories to minimise within-cluster variance.

The resulting scatter plot then shows which spending categories share similar temporal patterns (although for this representation we focus on only the first two principle-components for illustrative purposes). Finally, each cluster's average trajectory is normalised to equal 1 in 1999 to show relative growth through time.

4. External Bond Market Shock

This section documents a stylised scenario used to illustrate debt dynamics under an externally-driven increase in borrowing costs. This scenario builds off of the PBO Medium Term Outlook, assuming that the economy progresses in line with these projections until FY30. Follow that there is a temporary 150bp shock to borrowing costs occurs for three years with a permanent 50bp increase in the exogenous component of borrowing costs.

This scenario focuses on the build up of debt and endogenous GDP and fiscal responses following this interest rate shock. However, a broader recession or economic consequences are not included. As real life bond market shock would likely be driven by factors that drive down confidence or impose other costs on the Australian economy (i.e. weaker export prices) which would exacerbate the increase in overall debt.

The scenario is intended to be *transparent and decomposable*: each layer adds a single mechanism to a common baseline so that differences in outcomes can be attributed to a specific channel (interest rates, the GDP channel, and automatic fiscal stabilisers). The framework is deliberately simplified. It does not represent a full macro-fiscal forecasting model, and therefore should be interpreted as an accounting exercise rather than a projection.

The scenario uses historical cash balance data (PBO) to anchor initial fiscal stocks and flows, and then projects debt under simple behavioural assumptions for nominal GDP growth, the interest rate path, and the primary balance. Stock-flow adjustments are included explicitly, recognising that cash balances and observed debt changes differ for a variety of valuation and financing reasons.

4.1 Data inputs and constructed series

The historical dataset contains (annual) series for receipts, payments, gross debt and nominal GDP. Where available, interest is either observed directly or inferred from an effective interest rate. The final historical year (in this case FY30), denoted $t = 0$, provides the initial conditions: gross debt-to-GDP b_0 and nominal GDP Y_0. Nominal GDP growth is

(12)
$$g_t \equiv \frac{Y_t}{Y_{t-1}} - 1.$$

The effective interest rate is measured as

(13)
$$r_t \equiv \frac{I_t}{D_{t-1}},$$

where I_t denotes interest payments and D_{t-1} is the previous year's gross debt stock. When I_t is not directly observed, it is inferred using r_t (or vice versa) given the available series.

The primary balance (share of GDP), with surplus positive, is:

(14)
$$pb_t \equiv \frac{R_t - (P_t - I_t)}{Y_t},$$

where R_t is receipts and P_t is total payments.

Debt changes need not equal the cash deficit. Define the stock-flow adjustment (share of GDP) as:

(15)
$$sfa_t \equiv \frac{\Delta D_t - (P_t - R_t)}{Y_t},$$

where $\Delta D_t \equiv D_t - D_{t-1}$.

In the projections, sfa_t is set using a simple rule (the trailing historical average over the most recent decade) to preserve the empirical tendency for persistent differences between cash balances and debt accumulation.

4.2 Debt dynamics and interest payments

Let $b_t \equiv D_t/Y_t$ denote gross debt as a share of nominal GDP. The law of motion for debt-to-GDP is:

$$b_t = \frac{1 + r_t}{1 + g_t}\, b_{t-1} - pb_t + sfa_t. \tag{16}$$

Equation (16) makes explicit the three proximate drivers of the debt ratio: the interest-growth differential acting on the inherited debt stock, the primary balance, and stock-flow adjustments.

Interest payments as a share of GDP are:

$$i_t \equiv \frac{I_t}{Y_t} = \frac{r_t}{1 + g_t}\, b_{t-1}. \tag{17}$$

The overall balance (share of GDP) is then:

$$ob_t \equiv \frac{R_t - P_t}{Y_t} = pb_t - i_t. \tag{18}$$

Borrowing costs are decomposed into: (i) a baseline anchor rate path, (ii) an exogenous shock, and (iii) a debt-dependent risk premium:

$$r_t = r_t^\star + r_t^{add} + \underbrace{\phi\left(b_{t-1} - b^\star\right)}_{\text{debt-sensitive premium}}. \tag{19}$$

The parameter ϕ is a semi-elasticity capturing the idea that higher debt may raise risk premia (or term premia) through market perceptions of fiscal risk, rollover risk, or inflation compensation.[2]

The calibration used in this scenario is $\phi = 0.0003$, equivalent to 3 basis points for each 1 percentage point increase in b (in ratio units) based on estimates in Yan and Brittle (2010). The reference level $b^\star$ is set to the last observed historical debt ratio, ensuring the premium is zero at the initial condition.

4.3 Baseline projection

The baseline projection fixes nominal GDP growth and the anchor interest rate at constant values, and assumes a gradual improvement in the primary balance from the latest observed level. Specifically:

- Nominal GDP growth: $g_t = \bar{g}$ for all projection years.
- Anchor interest rate: $r_t^\star = \bar{r}$ for all projection years.
- Primary balance: pb_t improves gradually via a small increase in receipts (share of GDP) over the first two years, holding primary spending (share of GDP) constant.
- Stock-flow adjustments: $sfa_t = \overline{sfa}$, where $\overline{sfa}$ is the average of the last 10 historical observations of sfa_t.

Debt sensitivity is switched off in the baseline ($\phi = 0$) so that differences across scenarios can be interpreted relative to a common reference.

The scenario is designed to mimic an externally-originating repricing of long-term rates that spills over to domestic borrowing costs. The shock path is imposed as an additive wedge:

$$r_t^{add} = \begin{cases} 150 \text{ bps} & t = 1, 2, 3 \\ 50 \text{ bps} & t \geq 4, \end{cases} \tag{20}$$

with debt sensitivity activated ($\phi > 0$). This creates a mechanical increase in interest payments, which then feeds back into debt via Equation (16), and (through (19)) further increases borrowing costs when debt rises above the reference level.

To aid interpretation, the scenario is constructed in layers:

1. **Baseline** (no shock; $\phi = 0$).
2. **Rates-only** (shock r_t^{add}; $\phi > 0$; g_t and pb_t held at baseline).
3. **Rates + GDP** (adds a transient GDP growth response to the rate shock).
4. **Rates + GDP + fiscal** (adds automatic stabilisers via a pass-through from the GDP gap to the primary balance).

2 This reduced-form specification abstracts from maturity structure, issuance strategy and monetary policy responses. It is intended to capture a broad sensitivity of average borrowing costs to the debt ratio.

4.4 Transient GDP channel

The GDP channel is implemented as a finite impulse response mapping from the interest rate shock to deviations in nominal GDP growth. Let Δg_t denote the deviation of nominal growth from its baseline value $\bar{g}$. We specify:

$$\Delta g_t = \gamma \sum_{k=0}^{K} \beta_k \, r_{t-k}^{add}, \tag{21}$$

with coefficients $\{\beta_k\}$ chosen such that:

$$\sum_{k=0}^{K} \beta_k = 0. \tag{22}$$

Condition (22) implies that the level of nominal GDP experiences a temporary shortfall but eventually returns to its baseline path (i.e. the cumulative growth deviation sums to approximately zero). The GDP *gap* is defined as the cumulative deviation in growth:

$$gap_t \equiv \sum_{j=1}^{t} \Delta g_j. \tag{23}$$

The scale parameter γ is calibrated so that, under a shock of 150 bps sustained for three years, the cumulative growth shortfall over the first three years equals 3.5 per cent.[3]

Nominal growth in the GDP-channel scenarios is then:

$$g_t = \bar{g} + \Delta g_t. \tag{24}$$

To capture automatic stabilisation, the GDP gap is mapped into changes in the primary balance using a reduced-form semi-elasticity. Specifically:

$$\Delta pb_t = \eta \cdot (100 \cdot gap_t), \tag{25}$$

where gap_t is expressed in log-level units (approximately percentage deviations) and η is the primary-balance semi-elasticity per 1 per cent change in the GDP level.[4]

Changes in the primary balance are allocated predominantly to receipts (reflecting the typical greater cyclicality of revenue than primary spending in the short run). Denote baseline receipts and primary spending shares as rec_t^0 and ps_t^0. Then:

$$rec_t = rec_t^0 + \lambda \Delta pb_t, \tag{26}$$

$$ps_t = ps_t^0 - (1 - \lambda)\Delta pb_t, \tag{27}$$

with $\lambda = 0.85$ in the implementation, so that 85 per cent of the cyclical movement in the primary balance occurs through receipts.

4.5 Interpretation and decomposition

Because the scenarios are nested, differences in the projected debt path can be interpreted as marginal contributions of each mechanism:

- The *rates-only* scenario isolates the mechanical impact of higher interest costs and the debt-sensitive risk premium.
- The *GDP* layer adds the effect of weaker nominal GDP growth on the debt ratio through the denominator (Y_t) and via the interest-growth differential in Equation (16).
- The *fiscal* layer adds endogenous deterioration in the primary balance due to automatic stabilisers, which further raises debt and (through (19)) can amplify borrowing costs.

3 This calibration is intended to deliver a plausible order-of-magnitude macroeconomic response for a severe, externally-driven repricing. It is not a structural estimate. The magnitude was chosen based on estimates from Ballantyne et al. (2019) indicating an approximate 0.9% decline in NGDP from a one-year 100bp increase in the cash rate – with the shock considered here occuring for three years, and 50% larger.

4 The mapping is implemented in ratio-to-GDP terms, consistent with standard fiscal semi-elasticity concepts. In practice, semi-elasticities depend on tax bases, labour market conditions and the composition of the shock; the parameter used here is illustrative.

A useful accounting identity for interpreting results is the year-t change in the debt ratio:

$$\Delta b_t \approx \left(\frac{r_t - g_t}{1 + g_t}\right) b_{t-1} - pb_t + sfa_t, \tag{28}$$

which highlights that, conditional on inherited debt, increases in r_t and decreases in g_t both raise the snowball component, while cyclical fiscal effects operate through pb_t.

This framework is intentionally parsimonious and omits several features that would matter for a full assessment of fiscal sustainability:

- **Maturity structure and refinancing**: average interest costs respond gradually to market rates depending on the term structure of debt.
- **Endogenous policy responses**: discretionary consolidation (or stimulus) is not modelled.
- **Monetary policy and inflation**: the GDP response is imposed in reduced form and does not separate nominal growth into real activity and inflation, which may influence the nature of any policy response.
- **Risk and uncertainty**: outcomes are deterministic and do not incorporate stochastic shocks or confidence intervals.

Accordingly, the scenario should be read as an illustrative stress test of accounting mechanics rather than a probabilistic forecast.

5. Shapley Decomposition of Drivers

Interpreting a change in government spending requires evaluating which drivers were important in determining what has happened. By estimating the contribution of varying drivers (e.g. population ageing, changes in relative prices) we are able to understand more about why government spending in Australia has increased, and also project what may happen with government spending in the future.

In order to quantify these contributions we apply a Shapley-value decomposition to a regression model of government spending as a share of GDP. The Shapley approach is attractive as it provides a principled and transparent allocation of the explained change in spending to a set of correlated drivers, avoiding sensitivity to arbitrary ordering assumptions.[5]

The decomposition is best interpreted as an *accounting attribution* within a reduced-form statistical model, not a structural causal identification exercise. Specifically, the Shapley contributions describe how much of the observed change in the spending share can be attributed to changes in each driver *given the estimated historical relationship* between spending and the full set of drivers. This implies that two things:

- The estimated change in spending due to a given factor embeds both **mechanical drivers** of such changes (i.e. greater spending on the aged pension as more individuals age over 67) and **policy related drivers** that may be correlated with these factors (i.e. pressure to increase spending on the aged pension as a greater share of the population is aged over 67).
- These estimates are subject to concerns around omitted variables, and the fact that some orderings of marginal contributions may be closer to the *true* contribution than others.

5.1 Shapley decomposition

Let y_t denote government spending as a share of GDP in year t. In the preferred specification, spending is measured using National Accounts components and expressed relative to GDP. All plots in the main text focus on total government expenditure, however results are qualitatively similar when focusing on only government consumption or government demand (consumption plus investment):

$$y_t \equiv \frac{G_t}{Y_t}, \tag{29}$$

5 In many macro–fiscal applications the candidate drivers are mechanically and economically correlated (e.g. age shares co-move; unemployment and terms of trade co-move with activity; relative prices trend with structural change and economic growth). In such settings, sequential decompositions can be highly order-dependent. The Shapley value averages marginal contributions across all possible orderings and is commonly used to address attribution under correlation.

where G_t is government expenditure (in the application, a measure combining government final consumption expenditure, gross fixed capital formation, and cash transfers where relevant) and Y_t is GDP.

Let $x_t \in R^K$ collect K observable drivers. In the application these include demographic composition (age-group population shares), relative-price pressures (a standardised ratio of the government spending deflator to the GDP deflator), and macro conditions (e.g. unemployment and terms of trade, standardised). We estimate a no-intercept linear model:

$$y_t = x_t'\beta + \varepsilon_t. \tag{30}$$

The no-intercept specification is convenient when key regressors are shares that sum to one (e.g. age shares). In such cases, including an intercept alongside a full set of shares induces near-collinearity; removing the intercept yields an interpretable mapping from composition to the level of the spending share. More generally, the model is used as a parsimonious empirical summary of co-movement between y_t and x_t.

We estimate (30) over an annual sample (financial-year or calendar-year aggregation), excluding extreme outlier years where the relationship may be dominated by one-off shocks (e.g. the pandemic period). The model is estimated using OLS and all variables are aligned to the same annual frequency.

5.2 The decomposition problem

Consider two years t_0 and t_1. The observed change in the spending share is:

$$\Delta y \equiv y_{t_1} - y_{t_0}. \tag{31}$$

A natural fitted counterpart is:

$$\widehat{\Delta y} \equiv \widehat{y}_{t_1} - \widehat{y}_{t_0} = x_{t_1}'\widehat{\beta} - x_{t_0}'\widehat{\beta}, \tag{32}$$

with residual (unexplained) component $\Delta y - \widehat{\Delta y}$.

A challenge is to allocate $\widehat{\Delta y}$ across individual drivers when x_t is multivariate and correlated. In general, $\widehat{\Delta y}$ cannot be uniquely written as a sum of driver-by-driver contributions without further conventions. Shapley-value methods supply an axiomatic allocation that (i) exhausts the explained change, (ii) treats drivers symmetrically absent evidence to the contrary, and (iii) is robust to ordering choices.[6]

5.3 Shapley-value decomposition

Let $\mathcal{K} \equiv \{1, \ldots, K\}$ index drivers, and let $S \subseteq \mathcal{K}$ denote a subset of drivers. For each subset S, estimate the restricted model:

$$y_t = x_{t,S}'\beta_S + \varepsilon_{t,S}, \qquad \widehat{\beta}_S = \arg\min_{\beta_S} \sum_t (y_t - x_{t,S}'\beta_S)^2, \tag{33}$$

where $x_{t,S}$ denotes the vector of regressors in subset S. Define the fitted change using only subset S:

$$\widehat{\Delta y}(S) \equiv x_{t_1,S}'\widehat{\beta}_S - x_{t_0,S}'\widehat{\beta}_S. \tag{34}$$

The Shapley contribution of driver $i \in \mathcal{K}$ to the fitted change is:

$$\phi_i \equiv \sum_{S \subseteq \mathcal{K} \setminus \{i\}} \frac{|S|!\,(K - |S| - 1)!}{K!} \left(\widehat{\Delta y}(S \cup \{i\}) - \widehat{\Delta y}(S) \right). \tag{35}$$

Equation (35) averages the *marginal contribution* of adding driver i to a model across all possible subsets S not containing i, with weights equal to the probability that S precedes i under a uniformly random ordering of drivers.

Two properties are immediate:

6 The Shapley value is the unique attribution method satisfying efficiency (contributions sum to total), symmetry, dummy (irrelevant drivers receive zero), and additivity; see Shapley (1953).

$$\sum_{i=1}^{K} \phi_i = \widehat{\Delta y}(\mathcal{K}), \qquad \text{(efficiency / exact add-up)} \tag{36}$$

$$\Delta y = \sum_{i=1}^{K} \phi_i + \underbrace{\left(\Delta y - \widehat{\Delta y}(\mathcal{K})\right)}_{\text{residual / unexplained}} . \tag{37}$$

A key assumption of the Shapley approach is that there is no prior reason why we would favour a certain ordering above another. For example, if we had a theoretical reason why we should always consider the marginal contribution of unemployment conditional on the age distribution only ordering where this occurs should be averaged over. Our view was that not such ordering choice was theoretically justified, and have averaged over all potential orderings.

5.4 Estimation of time specific contributions for decomposition

The relationship between government expenditure and its observable drivers is unlikely to be structurally stable over long horizons. Changes in fiscal institutions and policy regimes may alter the sensitivity of expenditure to underlying determinants, leading us to break the estimates into four groups: 1980-89, 1990-2000, 2001-2011, and 2012-2024.

To accommodate potential parameter instability, the analysis does not impose a single coefficient vector over the full historical sample. Instead, Equation (5.3) is estimated separately over grouped time windows. Let $\mathcal{T}_j$ denote a partition of the historical sample into contiguous year blocks. For each block j, coefficients are estimated as:

$$\hat{\beta}_j = \arg\min_{\beta} \sum_{t \in \mathcal{T}_j} \left(y_t - x_t'\beta\right)^2 . \tag{38}$$

Shapley decompositions are then computed using the block-specific coefficients $\hat{\beta}_j$.

Allowing coefficients to vary across time blocks serves two purposes.

First, it reduces the risk that estimated relationships are driven by averaging across distinct policy and macroeconomic regimes. The responsiveness of expenditure to demographic structure, for example, may differ between periods characterised by rapid population ageing and periods of relative demographic stability.

Second, the approach aligns the attribution exercise with a more realistic interpretation of fiscal dynamics. Rather than assuming fixed behavioural or institutional elasticities, the decomposition reflects the empirical relationships prevailing within each period.

The resulting Shapley contributions should therefore be interpreted as *period-specific attributions*. They quantify how movements in drivers contributed to changes in the expenditure ratio under the reduced-form relationships observed within each time block.

This design choice recognises that fiscal responses evolve with policy frameworks and economic structure. It also mitigates concerns associated with long-sample parameter instability, which is particularly relevant for public finance aggregates observed over multiple decades.

While this strategy accommodates gradual shifts in relationships, it does not model the sources of parameter variation explicitly. Estimated coefficients may reflect changes in policy rules, compositional effects, or unobserved factors. Consequently, the decomposition remains descriptive rather than causal.

5.5 Projections and counterfactuals

This exercise is used to define *counterfactual* spending paths when ageing effects are fixed, and to produce projections based on the ageing anticipated in the ABS 2023 population projections.

For the counterfactual exercise the coefficients are estimated over the full period, rather than for one of the subperiods used in the main analysis. The reason for this is that we wanted the counterfactual spending path with fixed ageing to include institutional changes that will have occurred over the period of interest.

For the projection exercise marginal contributions from the final period are used (2012-2024), thereby presuming that policy habits and behaviours that developed over the past twelve years will persist going forward.

References

Abadie, A. (2021). Using synthetic controls: Feasibility, data requirements, and methodological aspects. *Journal of Economic Literature*, *59*(2), 391–425.

Abadie, A., Diamond, A., & Hainmueller, J. (2010). Synthetic control methods for comparative case studies: Estimating the effect of california's tobacco control program. *Journal of the American Statistical Association*, *105*(490), 493–505.

Abadie, A., Diamond, A., & Hainmueller, J. (2015). Comparative politics and the synthetic control method. *American Journal of Political Science*, *59*(2), 495–510.

Australian Bureau of Statistics. (2023). Population Projections, Australia, Latest Release.

Australian Bureau of Statistics. (2025a). Australian System of National Accounts, 2024–25 financial year.

Australian Bureau of Statistics. (2025b). Government Finance Statistics, Annual, 2023–24 financial year [Released 22 April 2025, Australian Bureau of Statistics].

Australian Treasury. (2023). *Intergenerational Report 2023* (Intergenerational Report). Australian Government Treasury. https://treasury.gov.au/sites/default/files/2023-08/p2023-435150.pdf

Azen, R., & Budescu, D. V. (2003). The dominance analysis approach for comparing predictors in multiple regression. *Psychological Methods*, *8*(2), 129–148.

Ballantyne, A., Cusbert, T., Evans, C., & Guttmann, R. (2019). Martin has its place: A macroeconometric model of the australian economy. *RBA Research Discussion Paper*, (2019-07).

Dunford, S. (2023). tidysynth: A tidy implementation of the synthetic control method [R package]. https://CRAN.R-project.org/package=tidysynth

KPMG. (2010, March). *CGE Analysis of the Current Australian Tax System* (Final Report) (26 March 2010). KPMG Econtech.

Lipovetsky, S., & Conklin, M. (2001). Analysis of regression in game theory approach. *Applied Stochastic Models in Business and Industry*, *17*(4), 319–330.

Nassios, J., & Giesecke, J. (2025). Inefficient at any level: A comparative efficiency argument for elimination of property transfer duties and insurance taxes. *Applied Economics*, 1–18. https://doi.org/10.1080/00036846.2025.2507977

Organisation for Economic Co-operation and Development (OECD). (2026). OECD Fiscal Decentralisation Database.

Parliamentary Budget Office. (2024). Tax Mix.

Parliamentary Budget Office. (2025). 2025–26 Medium-term Budget Outlook.

Shapley, L. S. (1953). A value for n-person games (H. W. Kuhn & A. W. Tucker, Eds.). *Contributions to the Theory of Games*, *2*, 307–317.

Shorrocks, A. F. (2013). Decomposition procedures for distributional analysis: A unified framework based on the shapley value. *Journal of Economic Inequality*, *11*(1), 99–126.

Yan, Y. H., & Brittle, S. (2010, September). *Reconsidering the link between fiscal policy and interest rates in australia* (tech. rep. No. 2010-04). Australian Treasury.