The Economic Impact of Artificial Intelligence and Robotics: A Quantitative Analysis Using Australian Input-Output Tables

Derek Harris

Abstract

This paper presents a quantitative analysis of the macroeconomic impact of artificial general intelligence (AGI) and AI-enabled robotics using a multi-sector input-output model calibrated to the 2023–24 Australian National Accounts. The model traces the effects of AI-driven total factor productivity (TFP) improvements through three channels: price reductions, income redistribution, and demand stimulus. Three TFP scenarios are modelled — Low, Central, and High — representing economy-wide Domar-weighted TFP increases of 28.3%, 56.7%, and 121.3% respectively.

The central finding is that the model, when fully specified, broadly replicates orthodox economic outcomes: real GDP grows substantially in all scenarios, and net employment falls only modestly in the base case (-3.2% central) or rises in the high scenario (+8.2%). However, this benign outcome rests critically on a mechanism — the creation of new employment-intensive industries absorbing income that exceeds the saturation capacity of existing demand — that standard economic theory predicts but cannot guarantee. In the central base case, 12.3% of the projected workforce is employed in such new industries; in the high scenario, 54.1%. If this channel does not function — because demand for genuinely novel human activities does not materialise at wage rates that make employment worthwhile — outcomes deteriorate sharply, and a large structural decline in employment cannot be ruled out.

The analysis also finds that the marginal propensity to consume (MPC) capital income is a parameter of first-order importance. When capital owners consume a smaller fraction of their income, the Keynesian multiplier collapses, employment falls dramatically, and the benefits of productivity growth accrue almost entirely to capital holders. The ratio of labour to capital income shifts materially against labour in all scenarios, from an initial 52.5%/47.5% split to 41.3%/58.7% in the central case and 29.1%/70.9% in the high case — even under the most benign employment outcomes. A further sensitivity analysis of the AI investment parameters (payback period and depreciation rate) finds that higher AI investment flows are net stimulatory due to the asymmetric effect of concentrated GFCF injections versus proportional income drains, producing substantially stronger GDP and employment outcomes when AI deployment intensity is higher.

---

1. Introduction

The arrival of large language models and AI-enabled robotics has prompted extensive commentary on their economic consequences. Much of this commentary oscillates between utopian abundance narratives and dystopian displacement fears. Both poles tend to underspecify the mechanisms through which productivity gains translate into income, employment, and welfare outcomes. This paper attempts a more systematic treatment, using the rigour of a multi-sector input-output framework to trace the full general-equilibrium effects of AI-driven TFP improvements through the Australian economy, adopted as an analog of a typical advanced economy.

The analysis addresses several questions. Does AI-driven productivity growth necessarily reduce aggregate employment, or can demand responses offset displacement? Through what mechanisms do productivity gains reach households, and how does their distribution affect macroeconomic outcomes? What role might new industries play in absorbing labour released from sectors subject to AI substitution? And how does the marginal propensity to consume capital income affect the system's response?

The paper is structured as follows. Section 2 describes the data and model structure. Section 3 details the key assumptions and parameters. Section 4 sets out the model mechanics. Section 5 presents base-case results across the three TFP scenarios. Section 6 analyses sensitivity to three sets of parameters: the MPC of capital income, the allocation of excess demand to new industries, and the AI investment parameters (payback period and depreciation rate). Section 7 examines the new industry channel in detail, including its implications for the orthodox recovery narrative. Section 8 discusses income distribution dynamics. Section 9 considers model limitations and the question of historical stabilising mechanisms. Section 10 concludes.

---

2. Data and Model Structure

2.1 Input-Output Framework

The model is built on the 2023–24 ABS Input-Output Tables for Australia, covering 115 industries aggregated into 39 functional industry groups plus two innovation industry categories (discussed below). The core analytical framework is the Leontief quantity model:

$$\mathbf{x} = (\mathbf{I} - \mathbf{A})^{-1}\mathbf{f}$$

where $\mathbf{x}$ is the vector of gross outputs, $\mathbf{A}$ is the technical coefficients matrix derived from intermediate transactions, and $\mathbf{f}$ is the vector of final demands. The inverse $(\mathbf{I} - \mathbf{A})^{-1}$ is the Leontief inverse, capturing both direct and indirect production requirements across the supply chain.

The model incorporates a price system alongside the quantity system. TFP improvements reduce the unit cost of production in each industry, which in turn reduces output prices proportionally to the share of TFP gains competed away. The economy-wide price deflator is computed as a final-demand-weighted average of industry price changes, using Domar weights to ensure consistency between the TFP aggregate and individual industry contributions.

2.2 Industry Coverage

The 39 industry groups cover the full spectrum of Australian economic activity, from primary production (agriculture, mining) through manufacturing, construction, infrastructure, and services. Three sectors receive special treatment in the model due to structural transitions associated with AI and decarbonisation:

Sectors with exogenous demand reductions: Fossil fuels (coal, oil, gas), motor vehicles, automotive repair, tertiary education, health care, aged care, and pharmaceuticals are assigned negative exogenous demand adjustments reflecting structural changes independent of the income effect — decarbonisation, Mobility as a Service (MaaS) substitution, improved health outcomes reducing system demand, and AI substitution for formal education and research.

Sectors with exogenous demand increases: Metals and minerals extraction (copper, rare earths, lithium for AI and energy transition infrastructure), electricity generation (renewables and AI data centre demand), and road transport (MaaS growth) receive positive exogenous demand adjustments.

New industry categories: Two innovation residual categories are added to capture economic activity not present in the current ABS tables: *Human-centric innovation* (labour ratio 75%, average wage A$58,300 per FTE) and *AI-augmented innovation* (labour ratio 25%, average wage A$135,000 per FTE). These categories receive the income that exceeds the saturation capacity of existing industry demand, modelled using the logistic income demand function described in Section 4.2.

2.3 Calibration to National Accounts

All monetary values are calibrated to the 2023–24 National Accounts. Base-year GDP is approximately A$2.7 trillion. The labour share of factor income (P1) is 52.5% and the capital share (P2) is 47.5%, consistent with ABS data. The current account trade surplus is A$40.7 billion.

---

3. Key Assumptions and Parameters

3.1 TFP Scenarios

Three scenarios are modelled. The defining parameter in each is the assumed labour productivity improvement for each industry group — the reduction in full-time equivalent labour required to produce the same output — alongside an equivalent capital efficiency improvement. TFP is then computed endogenously from these input reductions, weighted by gross output (the Domar method).

The scenarios represent materially different trajectories for AI deployment:

- **Low scenario**: Labour and capital input reductions of roughly half the central values. Represents a world where AI augments workers and capital modestly, improving efficiency in high-automation-potential sectors but leaving most of the economy broadly unchanged. Domar-weighted TFP increase of **28.3%**.
- **Central scenario**: Moderate-to-substantial input reductions across most sectors. In high-displacement industries (finance, professional services, construction, logistics, health care) labour inputs fall 95%; in lower-displacement industries they fall 15–45%. Represents a world of significant but uneven AI deployment over a medium-term horizon. Domar-weighted TFP increase of **56.7%**.
- **High scenario**: Maximum-plausible input reductions across all sectors — 95% in most capital-intensive and cognitive-labour-intensive industries. Represents a trajectory approaching AGI-enabled near-full substitution of cognitive labour across the economy. Domar-weighted TFP increase of **121.3%**.

Industries are assigned TFP parameters reflecting their susceptibility to AI and robotics substitution. Extraction support services, pharmaceutical manufacturing, construction trades, road transport, logistics, finance, professional services, health, education, and personal services carry the highest displacement assumptions. Food services, personal services, and entertainment carry lower assumptions, reflecting the resilience of human-proximity and experience-based activities.

3.2 Income Distribution

When an industry achieves a TFP improvement, the resulting gains are distributed across three channels using industry-specific parameters:

- **Share competed away to prices** (reducing output prices and benefiting consumers)
- **Share captured by labour** (increasing real wages or employment income)
- **Share captured by capital** (increasing profit margins or return on investment)

The labour share parameter requires clarification. The model uses a multiplier (set at 0.50 in the base case) applied to each industry's actual labour share of value-added. This means the effective proportion of TFP gains flowing to labour varies by industry, typically in the range of 20–35% rather than a uniform 50%. Industries with high labour intensity — public administration, education, health care — see relatively more TFP gains flow to labour; capital-intensive industries — finance, real estate, utilities — see relatively less.

The remaining proportion of TFP gains not captured by labour are split equally between price reductions and capital income enhancement.

3.3 Government Fiscal Stance

The model incorporates a government revenue and expenditure feedback loop. Government final consumption spending adjusts in line with projected changes in tax revenue from production (indirect taxes P3 and P4), holding the government's fiscal balance constant by construction in nominal terms. No assumption is made, or needs to be made, in respect of income taxes. This represents a broadly neutral fiscal stance — the government neither injects nor withdraws net demand beyond what is generated by the automatic stabilisers of the tax system. It implies that if productivity growth expands the tax base, government expenditure expands proportionally; it does not assume additional stimulus spending or austerity.

This formulation is conservative in that it does not model potential fiscal interventions — such as direct income transfers funded by capital taxes — that might be deployed in response to distributional pressures. It also does not assume that current deficits or surpluses are corrected. The constraint modelled is simply that government debt should not grow faster than GDP, which is consistent with a neutral to mildly contractionary real stance given that nominal GDP rises in all scenarios.

It is worth observing here though that in some of the more extreme scenarios, where employment declines dramatically, Governments may struggle to maintain this posture if welfare payments overwhelm Government revenue raising, and deficit funding, capacity.

3.4 Demand Saturation and Income Elasticity

Demand for most goods and services does not increase proportionally with income at high income levels. The model uses a logistic demand function to capture this saturation:

$$\mathbf{D(IC) = M / (1 + \exp(-J \cdot M/(M-1) \cdot IC + \ln(M-1)))}$$

where M is the saturation multiple (set at 2.0), IC is the aggregate income change, and J is the income elasticity parameter calibrated to each industry. The saturation multiple of 2.0 implies that, at the limit of income growth, households would consume at most double their current quantities of existing goods and services, noting though that the shape of the demand curve is asymptotic to M. Income above the logistic demand curve is diverted to the new industry categories rather than increasing demand for existing products.

It should also be noted that a number of industries, particulaly those low on people's hierarchy of needs, like food, education, and health care, or which are primarily Government supplied, like defence, are assumed to have a zero income elasticity. That is, they are assumed already be at the point of demand saturation. Mechanically, the main consequnce is that these industries don't factor into the calculation of average labour intensity, and hence employment. However, any effect is almost certainly swamped by the more consequential assumptions around the labour productivity potential of AI and robots, which in turn is the major driver of projected labour intensity.

3.5 External Sector and Trade Balance

Exports are projected at the industry level as the product of (i) the base-year export value, (ii) any exogenous demand shock applied to the industry, (iii) the projected price level for that industry, and (iv) a linear price-elasticity response. The exogenous demand shock captures structural changes independent of the income and price channels — most consequentially, a 90% reduction in fossil-fuel-related exports under decarbonisation, with smaller offsetting increases in metals and minerals extraction and electricity. The price-elasticity response captures the fact that lower domestic prices increase foreign demand for Australian output.

Competing imports are scaled at the using-column level: their nominal value tracks the nominal change in the corresponding industry intermediate use or final demand category, then deflated at projected domestic prices. This implicitly assumes that imports deflate at the same rate as domestic production — consistent with treating Australia as a representative advanced economy facing global AI-driven productivity gains rather than as a small open economy importing from an unchanged foreign price level. Currency and terms-of-trade effects are excluded, consistent with the closed-financial-system limitation noted in Section 9.

To close the external sector, a single uniform imports adjustment factor is applied across all imports such that the projected nominal export-to-import ratio matches the base-year ratio of 1.067. Under this construction, the X/M ratio is preserved in nominal terms (and approximately in real terms, given the equal-global-deflation assumption). The absolute trade balance therefore scales with the volume of trade. As nominal trade volumes fall — driven primarily by the structural reduction in fossil-fuel exports net of price-elasticity stimulus — the absolute trade balance falls proportionally. In the central scenario, the nominal trade balance falls from A$40.7bn to A$29.4bn; in the high scenario, to A$28.3bn.

It is worth being explicit about what this construction does and does not assume. It does not assume that AI-driven productivity gains improve Australia's terms of trade or generate an export-led growth boom; nor does it assume the opposite, that an AI-rich economy loses competitiveness to lower-cost producers. The constraint is structurally neutral: the external sector neither absorbs nor injects net demand relative to the base proportions. The decline in the absolute trade balance across scenarios is a mechanical consequence of the structural reduction in fossil-fuel-related exports modelled as an exogenous demand shock. A reader who believes AI deployment would meaningfully shift Australia's external position in either direction would need to adjust the exogenous demand parameters in column Z of the Scenario Parameters sheet accordingly.

Terms of trade treatment: The model's GDP projections measure real output at domestic producer prices and do not adjust for changes in the terms of trade (the ratio of export prices to import prices). This is deliberate, for three reasons.

First, the structural background of the model includes a 90% reduction in fossil-fuel-related exports as an exogenous decarbonisation shock. This reduction almost certainly causes some deterioration in the terms of trade relative to the base year, given Australia's historical comparative advantage in fossil fuel exports. However, this deterioration occurs identically across *all* AI and robotics scenarios — it is a feature of the decarbonisation trajectory, not the AI trajectory being analysed. Consequently, it cancels in any comparison between scenarios and does not affect the relative analysis that is the primary purpose of this paper.

Second, the incremental terms-of-trade effect attributable to AI deployment itself is genuinely ambiguous in sign. AI capital investment — hardware, data centre infrastructure, advanced robotics — has high import intensity, which would tend to worsen the terms of trade by increasing demand for foreign-produced capital goods. But economy-wide productivity gains simultaneously reduce the domestic cost of production, which strengthens export competitiveness and could improve the terms of trade. The relative magnitudes of these two effects cannot be determined without assumptions that are essentially arbitrary given current data.

Third, the trade balance constraint described above means that any increase in import demand for AI capital must be offset elsewhere in the economy through expansion of domestically-produced exports or import substitution. The model does not specify which sectors provide this offset, but its existence means the terms-of-trade question is largely about the composition of trade adjustment rather than its aggregate scale.

For these reasons, the model makes no adjustment to real GDP to reflect terms-of-trade effects. The GDP projections should be interpreted as measures of domestic productive output at constant prices. Translating these into real purchasing power measures that account for the terms of trade would require additional assumptions outside the scope of the present analysis.

---

4. Model Mechanics

4.1 Price Channel

AI-driven TFP improvements reduce unit production costs in each industry. The share of cost reduction passed on to output prices is governed by the "competed away" parameter, reflecting competitive market pressures. The economy-wide price deflator is constructed as a Domar-weighted average of industry-level price changes, adjusted for the input-output structure so that cost reductions in upstream industries propagate downstream. In the central scenario, the aggregate price deflator falls to 90.1% of the baseline — implying a roughly 10% reduction in the general price level.

This price deflation increases real purchasing power for all final demand components. For households, it is equivalent to an income increase without a nominal income change. The model captures this through the price elasticity channel: lower prices stimulate additional quantity demanded according to each industry's price elasticity of demand.

A critical design feature is that the income effect embedded in the price elasticity (the Slutsky income effect) is explicitly excluded from the income demand function. This prevents double-counting: the demand response to lower prices is captured by the price elasticity formula; the demand response to actual income gains (from labour and capital income) is captured separately through the income channel. Without this separation, the model would overstate household demand by counting the purchasing power effect of price reductions twice.

4.2 Income Channel

Nominal income from labour (P1) and capital (P2) changes in line with the TFP distribution parameters. The income channel then operates through the marginal propensity to consume: a fraction of additional income flows back into expenditure on goods and services, amplifying the initial demand stimulus.

Aggregate income change (IC) is computed net of three deductions:

1. **Price-induced demand**: income that is economically "consumed" by the additional real spending triggered by lower prices (this income cannot simultaneously fuel income-elastic demand expansion).
2. **AI investment deduction**: GFCF is held constant in nominal terms other than investment in AI and robotics. The AI / robotics investment represents a change from the steady state and needs to be funded. To ensure closure of the savings / investment loop, it is assumed that income is redirected from households via savings to the AI / robotics investment.
3. **Residential GFCF**: household investment in residential construction (which is technically GDCF rather than consumption) is assumed to be a function of real household income change (computed using the same logistic income-demand function as the consumption channels). This represents income redirected from current consumption into housing capital, completing the savings/investment closure alongside the AI/robotics investment.

The remaining IC represents the genuine income available to stimulate demand through the logistic income elasticity function.

4.3 Investment Channel

The model assumes that AI/robotics deployment requires capital investment at a scale calibrated by a one-year payback assumption: the productivity gain generated by AI in one year is sufficient to fund the construction of the AI capital stock. This determines the implied stock — approximately A$733bn in the central scenario and A$1,536bn in the high scenario. The annualised AI investment flow appearing in GFCF is the steady-state replacement of this stock at a 20% annual depreciation rate: approximately A$147bn in the central scenario and A$307bn in the high scenario. These figures apply to the base case parameters; the sensitivity of results to these parameters is analysed in Section 6.3.

The GFCF injection from AI investment creates additional demand directly in construction, equipment manufacturing, and ICT sectors, providing a demand offset to labour displacement in those industries.

4.4 New Industry Channel

Income that exceeds the saturation capacity of existing industries — where the logistic function approaches its upper bound — is modelled as flowing to two new industries not currently represented in the Input-Output tables. This mechanism replicates historical experience, where waves of automation have given rise to entirely new occupations that have absorbed the labour displaced by the automation. This residual demand is allocated between the two innovation categories according to the "share to human-centric industries" parameter. In the base case, 50% goes to human-centric innovation; a sensitivity scenario directs 90% to this channel.

Employment in the two new industries is then calculated from the income value of the residual demand and the assumed wage rates: A$58,300 per FTE for human-centric innovation and A$135,000 per FTE for AI-augmented innovation. The labour ratio of each category (75% and 25% respectively) determines the split between labour and capital income.

---

5. Base Case Results

Table 1 presents the headline results for the base case (MPC capital income = 100%, 50% of innovation demand to human-centric industries).

Table 1: Base Case Results

Metric	Low	Central	High
Domar-weighted TFP change	+28.3%	+56.7%	+121.3%
Aggregate price deflation	−5.0%	−9.9%	−20.2%
Real GDP change	+20.6%	+47.1%	+91.1%
Employment change	−3.9%	−3.2%	+8.2%
Real wage and salary income	+10.9%	+23.5%	+39.4%
Real capital income	+35.4%	+93.4%	+275.7%
Real aggregate household income	+22.6%	+56.7%	+151.7%
Labour share of income (P1): current → projected	52.5% → 47.5%	52.5% → 41.3%	52.5% → 29.1%
Capital share of income (P2): current → projected	47.5% → 52.5%	47.5% → 58.7%	47.5% → 70.9%
New industries as % of employment	3.1%	12.3%	54.1%

5.1 GDP and Employment

Real GDP rises strongly in all scenarios — +20.6% in the low, +47.1% in the central, and +91.1% in the high. The employment outcomes are strikingly benign. Net employment falls by only 3.9% in the low and 3.2% in the central scenario; in the high scenario, net employment rises by 8.2% above the baseline. This pattern closely resembles the historical experience of disruptive general-purpose technologies: initial displacement offset by demand-led expansion, with net employment broadly stable or rising at the aggregate level.

This result is not an artefact of optimistic assumptions. It reflects the interaction of three offsetting forces: (1) lower prices stimulating additional demand through price elasticity, (2) income redistribution stimulating additional demand through income elasticity, and (3) AI investment creating direct demand in capital goods sectors. The combined effect is sufficient to absorb the labour released by productivity improvements in the base case.

5.2 Income Distribution

While employment outcomes are benign, the income distribution picture is consistently and substantially unfavourable to labour. In all three scenarios, the labour share of national income declines. In the central case, P1 falls from 52.5% to 41.3%, a reduction of 11.2 percentage points. In the high case, P1 falls to 29.1%, reflecting a fundamentally different factor income structure in which capital claims more than twice the income share of labour.

In absolute terms, real wages and salaries still grow — by +23.5% in the central and +39.4% in the high case. Workers are not worse off in real terms. But the gains from productivity are distributed heavily in favour of capital: capital income rises by +93.4% in the central and +275.7% in the high, against labour income gains roughly one-third as large. The distribution of productivity gains increasingly resembles the early industrial revolution pattern of rising wages alongside a dramatic shift in factor shares.

This finding has important long-run implications. The Kaldor stylised fact — stable factor shares over long periods — has already been eroding since approximately 1980, alongside the first wave of information technology deployment. The present analysis suggests AI could accelerate this trend by an order of magnitude, producing a factor income structure with no historical precedent in an advanced market economy.

---

6. Sensitivity Analysis

6.1 Sensitivity to MPC of Capital Income

The marginal propensity to consume capital income is perhaps the most consequential parameter in the model. Capital income is concentrated among a relatively small number of households; these households in aggregate have higher savings rates and lower MPCs than the general population. The base case assumes MPC = 100% for capital income, consistent with full recycling of capital gains into consumption — an upper bound. Table 2 shows outcomes when this parameter is reduced to 50%.

Table 2: Results with MPC Capital Income = 50%

Metric	Low	Central	High
Real GDP change	+12.9%	+28.5%	+56.8%
Employment change	−10.2%	−20.7%	−52.4%
Real wage and salary income	+4.8%	+5.7%	−24.8%
Real capital income	+25.2%	+61.0%	+160.3%
New industries as % of employment	2.5%	5.6%	14.5%

The effect is dramatic. With MPC capital = 50%, central employment falls by 20.7% rather than 3.2%, and in the high scenario employment collapses by 52.4%. Real wage incomes, while positive in the low and central cases, turn sharply negative (−24.8%) in the high scenario as the collapse in employment overwhelms the per-worker productivity gain. Real capital income still rises strongly in all cases — it is the *labour* income that takes the damage.

The mechanism is the Kalecki multiplier effect. Capital income is distributed to a smaller, higher-saving population. If these recipients do not recycle their gains into consumption — whether by saving, investing offshore, or accumulating financial assets — the demand stimulus that sustains employment in the base case does not materialise. The economy's output capacity grows but its purchasing power does not, producing a demand-deficiency outcome despite rising potential GDP.

This sensitivity result has a direct real-world interpretation. Whether or not the benign base-case employment outcome is realised depends critically on whether the owners of AI capital — technology companies, their shareholders, and the asset owners who finance AI infrastructure — deploy their earnings in ways that sustain final demand. Historical evidence on this point is mixed.

The standard defence of MPC = 100% rests on the Keynesian savings-equals-investment identity (S = I). Savings must, by accounting definition, equal investment, so if capital owners save rather than consume, the funds must flow through the financial system back into productive investment — which then generates income for someone else. Budget conserved, demand maintained.

However, this breaks down in at least two ways that are directly relevant:

First, investment may not absorb all savings at full-employment output. If AI collapses the capital requirements per unit of output (fewer workers, less physical capital needed when AI does the work), then desired investment may simply be insufficient to absorb the savings of high-income capital owners. The equilibrating mechanism is falling interest rates — but if real rates hit zero, there is a liquidity trap. Money sits in the financial system and doesn't flow into real demand. This is exactly Summers's secular stagnation argument, and it applies with particular force to an AI-abundant economy.

Second, when capital owners save, they typically buy existing financial assets — shares, real estate, bonds. These are secondary market transactions. They drive asset prices up, but they don't create new productive capacity or new income for anyone outside the asset-owning class. The money circulates within the financial system. Quantitative Easing provided a large-scale natural experiment: the US Federal Reserve created ~$8 trillion; asset prices inflated dramatically; real demand growth was modest. The transmission from financial accumulation to real demand is genuinely leaky.

The post-1980 secular increase in corporate savings rates and the concurrent decline in labour's income share suggest that capital has, in practice, exhibited an effective MPC well below 100% in recent decades. If this pattern continues — or intensifies with AI — outcomes closer to the low-MPC scenario are plausible: a 50% figure is reasonable as a lower-bound sensitivity. The unknown is whether the longer-run equilibration (eventually, investment does absorb savings, rates adjust, etc.) will rebalance to achieve an effective MPC of 100%. To the extent it doesn't, there is a strong macro-economic argument for

Government to intervene, either directly or through wealth redistribution, to stimulate demand.

6.2 Sensitivity to New Industry Allocation

The second sensitivity varies the proportion of innovation demand directed to human-centric (labour-intensive) versus AI-augmented (capital-intensive) industries. In the base case this split is 50/50. Table 3 shows outcomes when 90% of innovation demand is directed to human-centric industries.

Table 3: Results with 90% to Human-Centric Innovation Industries

Metric	Low	Central	High
Real GDP change	+20.6%	+47.1%	+91.1%
Employment change	−2.2%	+3.9%	+43.3%
Real wage and salary income	+11.8%	+27.6%	+63.4%
Real capital income	+34.4%	+88.9%	+249.3%
New industries as % of employment	4.8%	18.4%	65.3%

GDP outcomes are identical to the base case (the total income is the same regardless of how it is allocated between the two innovation categories), but employment and wage income outcomes improve substantially. Central employment is now +3.9% and the high scenario produces a remarkable +43.3% employment increase — an economy that has grown to employ far more people than before the AI transition.

The 90% human-centric allocation represents a world in which most of the new activity generated by rising incomes is directed toward goods and services that require substantial human involvement: care, experience, craft, community, creative production, personal services at higher quality levels. This is not an unreasonable vision of how wealthy societies might direct their elevated purchasing power, but it requires both that the demand for such activities materialises and that humans can supply them at wage rates consistent with voluntary participation.

6.3 Sensitivity to AI Investment Parameters

Two parameters govern the scale of AI and robotics investment in the model: the payback period and the annual depreciation rate. Their effects on macroeconomic outcomes are counterintuitive on first inspection but internally consistent once the two-channel structure of the investment mechanism is understood.

Payback period determines the implied AI capital stock. The model assumes that AI deployment is economically rational: the productivity gain generated in one payback period equals the cost of the capital stock that produces it. A one-year payback implies a capital stock equal to one year's worth of cost savings; a two-year payback implies a stock twice as large. Longer payback periods may reflect higher upfront hardware and deployment costs, extended timelines for productivity benefits to fully materialise, or firstmover investment by firms accepting slower returns in exchange for competitive positioning.

Depreciation rate determines the annual investment flow required to maintain the stock at steady state. A 20% annual rate implies an average asset life of five years; a 30% rate implies approximately three and a half years. Higher depreciation — reflecting faster obsolescence of AI hardware and software — requires proportionally higher annual gross fixed capital formation to sustain the same stock.

The model operates two channels simultaneously for AI investment:

1. **GFCF injection**: Annual AI investment adds directly to gross fixed capital formation in equipment manufacturing and ICT sectors. This injection propagates through the Leontief inverse, amplifying demand throughout the supply chain upstream of those sectors.

2. **Income drain**: The same annual investment flow is treated as a deduction from household disposable income — representing the savings required to fund it — reducing the income available to stimulate consumption.

These two channels operate in opposite directions, but they are not symmetric in their macroeconomic effect. The GFCF injection is concentrated in sectors with substantial upstream linkages — equipment manufacturing, electronics, construction — where the Leontief supply-chain multiplier is high. The income drain, by contrast, is applied proportionally across the total income base of approximately A$4 trillion. A given dollar of annual AI investment therefore creates more demand through supply-chain amplification than it suppresses through its dilution of the aggregate income stimulus. The net effect is stimulatory: larger AI investment flows increase GDP, not reduce it.

This resolves what initially appears paradoxical. Doubling the payback period doubles the capital stock and, for a given depreciation rate, doubles annual investment. Both channels scale proportionally — but the net stimulus grows because the Leontief multiplier on the concentrated GFCF injection exceeds the proportional dilution on the aggregate income drain. Similarly, reducing the depreciation rate (say from 20% to 10%) halves annual investment for a given capital stock: this shrinks the GFCF injection more than it shrinks the income drain, reducing net GDP relative to the base case. The sign of the effect is the same for all combinations: more annual AI investment → higher GDP; less annual AI investment → lower GDP.

Table 4 presents results for a combined sensitivity: a two-year payback period and a 30% depreciation rate. This combination triples the base-case annual investment rate — the capital stock doubles (from the payback extension) and the depreciation rate on that larger stock rises from 20% to 30%, implying annual AI investment approximately three times the base-case level in the central scenario (A$440bn versus A$147bn).

Table 4: Results with 2-Year Payback Period and 30% Depreciation Rate

Metric	Low	Central	High
Real GDP change	+32.4%	+72.5%	+138.0%
Employment change	+5.4%	+22.6%	+87.1%
Real wage and salary income	+23.0%	+56.1%	+136.9%
Real capital income	+47.8%	+132.1%	+411.3%
Real aggregate household income	+34.8%	+92.2%	+267.4%
Labour share of income (P1): current → projected	52.5% → 47.9%	52.5% → 42.6%	52.5% → 33.8%
New industries as % of employment	3.7%	19.3%	66.2%
Implied AI capital stock (A$bn)	220	440	921
Implied annual AI investment (A$bn)	733	1,466	3,071

The contrast with the base case is substantial. Central GDP rises by 72.5% rather than 47.1%; employment increases by 22.6% rather than falling by 3.2%; and real household income rises by 92.2% rather than 56.7%. The fraction of employment in new industries increases from 12.3% to 19.3% in the central case, reflecting the greater income stimulus flowing to residual demand as the larger supply-chain multiplier effect pushes income higher.

The income distribution picture remains qualitatively similar across both scenarios — capital income grows proportionally faster than labour income — but both groups benefit more in absolute terms under the higher-investment scenario. Real wages rise by +56.1% in the central case, compared with +23.5% in the base case. The labour share of income (P1) is slightly better preserved (42.6% versus 41.3%), consistent with the higher employment base absorbing more of the productivity gains through labour demand.

It is important not to interpret these results as implying that investors should seek longer payback periods to maximise GDP. The payback period is an exogenous parameter representing market conditions — the rate at which AI deployment occurs and the capital intensity at which it is undertaken. The model results indicate that, across the range examined, higher AI investment intensity has net stimulatory macroeconomic effects; they do not imply that inefficient or overpriced investment is desirable. What the results do suggest is that the macroeconomic consequences of AI depend not only on the productivity gains it generates but also on the scale of investment required to achieve those gains: a more capital-intensive AI deployment path, counterintuitively, produces a stronger demand stimulus alongside the supply-side gains.

---

7. The New Industry Channel: A Critical Dependency

The most striking finding in Table 1 is the employment share of new industries. In the high base-case scenario, 54.1% of projected employment is in the two innovation categories —

industries that do not currently exist in their projected form in the national accounts. The base case with 90% human-centric allocation produces 65.3%. Even in the central base case, 12.3% of employment is in these new sectors.

To be explicit about what this means: the orthodox outcome — employment broadly sustained despite a doubling of economy-wide TFP — is only achievable if the economy creates a large volume of genuinely new jobs in occupations that do not currently exist. The existing industries, even after full demand response through the price and income channels, do not generate sufficient employment to replace what AI substitution removes at the aggregate level. The new industry channel provides the remainder.

This finding does not make the orthodox outcome impossible. Economic history provides ample precedent for the creation of entirely new industries in response to rising incomes and technological capability. Entire sectors — digital content creation, e-commerce, app development, social media, streaming entertainment, medical aesthetics, mental health services at scale — were negligible or non-existent thirty years ago. The mechanisms by which such industries emerge are well understood: rising incomes, falling costs, and entrepreneurial discovery of unsatisfied latent demand.

The question the model cannot answer — and that no quantitative model can answer — is whether the magnitude and pace of new industry creation will be sufficient to absorb the labour released by this particular wave of technological change. Three considerations suggest this is not guaranteed:

Comparative advantage may not hold at viable wages. Classical trade theory guarantees that human labour retains comparative advantage even when AI is absolutely superior at every task. But comparative advantage does not guarantee adequate *returns*. If AI can perform most cognitive and physical tasks at near-zero marginal cost, the terms of trade move heavily against human labour. Workers may retain a comparative advantage in some activities — care, authenticity, physical presence, embodied skill — while earning wages insufficient to support the living standards that rising aggregate income would imply for everyone else. The historical analogy is horses post-mechanisation: they retained comparative advantage in some domains for decades, but at economically marginal returns.

Demand for new human activities may saturate. The new industry channel requires that rising incomes translate into demand for human-provided services at meaningful scale. But the saturation dynamic in the model applies to new industries too, eventually. A society that has already satisfied demand for health, education, transport, food, shelter, entertainment, and, ultimately, self-actualisation, at high quality may reach a point where additional income predominantly flows into savings or capital assets rather than labour-intensive consumption. The path-dependence of consumption habits and the inequality of the income distribution (where the marginal propensity to consume additional income is lowest among those who receive the most) both constrain the channel.

Speed may overwhelm adjustment capacity. Historical transitions that economists point to as successful — the agricultural to industrial transition, the manufacturing to services transition — played out over decades to generations, allowing labour market reallocation through natural attrition, retraining, and generational replacement. If AI deployment is rapid — and the rate of capability improvement in recent years suggests it could be — the required reallocation may outpace the institutional mechanisms (education systems, labour market flexibility, geographic mobility) that historically enabled it.

The model represents the orthodox outcome as achievable in principle, but does not represent it as guaranteed in practice. The honest reading of the results is that the model has clarified the question, not answered it: what matters is whether the new industry channel functions. That is a social and political question as much as an economic one.

---

8. Income Distribution and Factor Shares

The factor share results deserve separate attention because they represent a structural change whose significance extends beyond the employment question. In every scenario modelled, with every parameter combination, the labour share of national income (P1) declines. The only variable is by how much.

Under the base case:

- Low scenario: P1 falls from 52.5% to 47.5%, P2 rises to 52.5%
- Central scenario: P1 falls to 41.3%, P2 rises to 58.7%
- High scenario: P1 falls to 29.1%, P2 rises to 70.9%

This shift is structural, not cyclical. It reflects the fact that AI substitutes for labour more extensively than it substitutes for capital, and that capital (broadly defined to include the AI assets themselves) captures a larger share of the productivity gains. It is important to emphasise though that this does not necessarily imply super-profits: the share of capital used in production increases, and this needs to be funded, so even with normal returns, the share of income flowing to capital will increase. Even with full MPC capital recycling and a fully functioning new industry channel, the consequence is that the income share accruing to wage and salary earners declines substantially.

The implications compound over time. In a world where capital income grows at 93.4% (central) or 275.7% (high) while wage income grows at 23.5% or 39.4%, the wealth distribution widens significantly within a single model period. Over multiple periods, through the mechanism that Piketty identified — $r > g$, where r is the return on capital — this translates into an ever-larger wealth concentration, which in turn affects the MPC of capital income, the political economy of redistribution, and the demand dynamics explored in Section 6.1.

The question to be addressed though is whether this shift is necessary, or indeed even sustainable. Historical evidence on factor shares provides a partial guide but not a decisive one. The Kaldor fact — approximate stability of factor shares over long periods — held reasonably well for most of the twentieth century but has been visibly eroding since approximately 1980. Several plausible stabilising mechanisms operated historically: strong collective bargaining institutions, political economies of full employment, and the complementarity between capital and labour that characterised most pre-digital production technologies. None of these can be assumed to apply with the same force to AI deployment.

Two other longer-run stabilisers offer more comfort, though their pace and reliability are uncertain. First, capital accumulation itself eventually reduces the return to capital as the capital-to-output ratio rises and diminishing marginal returns set in — the neoclassical convergence mechanism. When the marginal return to capital falls toward the growth rate of the economy, excess savings decline and the pressure on factor shares partially corrects. Second, technological and economic obsolescence continuously erodes the value of existing capital vintages, preventing indefinite compounding. Neither mechanism can be assumed to operate quickly enough to prevent significant wealth concentration in the transition period, particularly if the winner-take-all dynamics associated with network effects and AI platform scaling allow returns in the leading AI assets to remain elevated even as aggregate capital accumulates. However, it is not inevitable that the increased concentration of wealth in the capital owning class would persist in the medium-term, even absent active Govenment intervention to redistribute wealth.

---

9. Limitations and Caveats

Several limitations of the present model should be acknowledged explicitly.

Static structure: The Leontief model assumes fixed technical coefficients. In reality, firms substitute between inputs in response to relative price changes, and the industry structure of the economy evolves over time. The model captures the first-round effects of TFP changes but not the dynamic industrial reorganisation that would follow.

Homogeneous labour: The model treats the labour force as a single factor. In reality, AI substitution is strongly skill-biased: cognitive routine tasks are more susceptible than manual dexterity tasks, and high-skill professional tasks more susceptible than low-skill personal service tasks (though this relationship is not monotonic). The distributional implications within labour income — wage divergence between AI-complementary and AI-substitutable workers — are not captured.

Closed financial system: The model does not capture the financial channels through which productivity gains are intermediated — equity market valuation changes, interest rate effects, wealth effects on consumption, and international capital flows. These channels can both amplify and dampen the real-economy effects modelled here.

No transition dynamics: The model compares a before and after state but does not model the path between them. Transition costs — including cyclical unemployment during reallocation, capital obsolescence write-downs, and the hysteresis effects of prolonged unemployment — could be substantial even if the end state is broadly benign.

New industry parameterisation: The innovation industry categories are necessarily schematic. The labour ratio, wage levels, and demand allocation are uncertain parameters that affect employment outcomes materially. The sensitivity analysis in Section 6.2 illustrates the range but does not constrain it.

Parameter uncertainty: All scenario parameters — TFP magnitudes, income distribution splits, price and income elasticities, demand saturation multiples — carry substantial uncertainty. The saturation multiple of 2.0 is a modelling choice; empirical validation across income levels and product categories would require substantial additional work. The scenarios should be read as illustrative of the structural relationships in the model rather than as predictions.

Globally homogenous deflation: The model treats Australia as an analogue of a typical advanced economy and assumes AI-driven productivity gains occur globally at broadly similar rates. Consequently, imports are deflated at the same rate as domestic production. Currency and terms-of-trade effects, which would matter in a small-open-economy framing, are excluded — consistent with the closed-financial-system limitation. As discussed in Section 3.5, the terms-of-trade effects of AI deployment are ambiguous in sign and cannot be quantified without assumptions that carry no strong empirical basis. The GDP projections should accordingly be interpreted as real output measures at domestic prices, not as real purchasing power measures.

No intermediate inputs to new industries: Real-world emergent industries (digital content, e-commerce, fintech, streaming, biotech) all have substantial intermediate purchases — technology, real estate, professional services, advertising. However, given the theoretical new industries have no known characteristics, the model treats the two innovation residuals as purely value-added, with no I-O propagation. This understates the multiplier effect through the rest of the economy and, all else equal, understates the GDP impact from the new-industry channel. In the high scenario, where new industries reach 54% of employment, this is a non-trivial simplification. Their gross output equals their household consumption value, with no induced demand for inputs from the rest of the economy.

Fixed wages in new industries: Wages in the two innovation industry categories are held fixed regardless of the scale to which these sectors grow. Under scenarios where new industries absorb a large share of total employment, labour-market clearing would normally drive wage rates up; the model's fixed-wage assumption is conservative in employment terms (more jobs at a constant wage rather than fewer at a higher one) but means the absolute labour income flowing to new-industry workers may be understated, particularly under the high scenario.

---

10. Conclusion

The model presents a coherent quantitative account of how AI-driven productivity improvements would flow through the economy. The headline finding is more benign than initial intuition might suggest: under base-case parameters, employment falls only modestly in the low and central scenarios and rises in the high scenario, while real GDP grows substantially and real household incomes rise for both workers and capital owners.

This optimistic view needs to be heavily qualified by three observations though:

First, the benign outcome is parameter-dependent in important ways. It assumes that capital owners recycle their substantially enlarged incomes into consumption at a rate approaching 100% — a strong assumption given historical savings behaviour among high-income households, and one that becomes less plausible as capital income concentration intensifies. It also assumes that fiscal policy maintains a broadly non-contractionary stance. If either condition fails, employment outcomes deteriorate sharply.

Second, and more fundamentally, the benign outcome requires a new industry channel of substantial size. In the central scenario, 12.3% of total employment must be in industries that do not currently exist at scale; in the high scenario, more than half. Standard economic theory predicts this will happen — rising incomes create demand for new goods and services, and entrepreneurial activity responds. But economic theory predicts *that* new industries emerge; it does not guarantee *when*, at what pace, or at what wage levels for displaced workers. The orthodox outcome is possible, but it is not inevitable.

Third, regardless of employment outcomes, the distribution of income shifts persistently and substantially toward capital. Real wages grow, but capital income grows two to three times faster, compressing the labour share of national income across all scenarios. The political economy implications of this shift for tax systems, for redistribution mechanisms, and for the legitimacy of market-determined income allocation, are significant, and are not resolved by the observation that employment may be broadly maintained.

The model, in replicating orthodox economic outcomes under appropriate conditions, has clarified rather than resolved the central question. The AI transition may follow the historical pattern of disruptive general-purpose technologies — creative destruction followed by new industry creation, with broadly stable employment and rising real incomes. Whether it does depends on outcomes around capital income deployment, fiscal policy, labour market institutions, and the social and political conditions under which new industries are cultivated, that the model itself assumes, but is unable to give any insight into.

Attitudes of AI commentators seem to have hardened into three camps, all with arguably extreme views. The dystopians imagine a "Terminator" world of machine dominance and human civilisation decline, if not collapse. The orthodox economic reaction is that AI is a nothing-burger - it is just the next in a long series of significant, but not world-defining, productivity advances that will play out over decades and produce historically aligned outcomes. The techno-optimists picture a world of abundance and human flourishing in a world where AI is a ubiquitous servant of humanity.

This exercise poses sharp questions for all three: Essentially, humans are not mere passengers on this journey. Decisions made by individuals and Governments will profoundly influence the destination. The dystopians need to recognise that with the right policy settings AI can be a boon for humanity. Orthodox economists need to recognise that there are very real arguments for why the scale of change will disrupt traditional automatic stabilisers. Techno-optimists need to address the reality of demand destruction from labour displacement, which could undermine the entire project, and they need to offer solutions beyond vague notions like "universal high income". Real policy responses are very likely to be needed, and the techno-optimists are at least being proactive in floating novel ideas. However, the scope of what is required has high potential to break conventional economic and political economy frameworks. And responses are likely to be needed faster than conventional political processes typically operate.

Appendix: Key Model Parameters

Table A1: TFP Scenario Assumptions (Labour and Capital Productivity Increases)

Industry Group	Low	Central	High
Agriculture	7.5%	15.0%	45.0%
Coal, oil & gas	2.5%	5.0%	15.0%
Metals & minerals	2.5%	5.0%	15.0%
Extraction support	25.0%	50.0%	95.0%
Food & clothing manufacturing	12.5%	25.0%	75.0%
Construction materials	12.5%	25.0%	75.0%
Printing & entertainment	2.5%	5.0%	15.0%
Pharmaceuticals	25.0%	50.0%	95.0%
Motor vehicles	25.0%	50.0%	95.0%
Equipment manufacturing	12.5%	25.0%	75.0%
Electricity	5.0%	10.0%	30.0%
Water & waste	12.5%	25.0%	75.0%
Residential construction	20.0%	40.0%	95.0%
Non-residential construction	20.0%	40.0%	95.0%
Civil construction	20.0%	40.0%	95.0%
Construction services	20.0%	40.0%	95.0%
Wholesale trade	37.5%	75.0%	95.0%
Retail trade	10.0%	20.0%	60.0%
Accommodation	10.0%	20.0%	60.0%
Food & beverage services	2.5%	5.0%	15.0%
Road transport	40.0%	80.0%	95.0%
Other transport	10.0%	20.0%	60.0%
Courier & logistics	30.0%	60.0%	95.0%
Computer & internet services	7.5%	15.0%	45.0%
Finance & insurance	25.0%	50.0%	95.0%
Rental & hiring	25.0%	50.0%	95.0%
Imputed & actual rent	0.0%	0.0%	0.0%
Property & agency services	7.5%	15.0%	45.0%
Professional & technical services	30.0%	60.0%	95.0%
Cleaning & maintenance	15.0%	30.0%	90.0%
Public administration	12.5%	25.0%	75.0%

Defence	12.5%	25.0%	75.0%
Public order & safety	5.0%	10.0%	30.0%
Primary & secondary education	25.0%	50.0%	95.0%
Tertiary education	25.0%	50.0%	95.0%
Health care	25.0%	50.0%	95.0%
Aged & social services	12.5%	25.0%	75.0%
Automotive repair	25.0%	50.0%	95.0%
Personal services	2.5%	5.0%	15.0%

Note: Exogenous demand adjustments (decarbonisation, MaaS substitution, health improvement effects) applied separately and not shown in this table.

Table A2: Key Macro Parameters (Base Case)

Parameter	Value
Marginal propensity to consume — wage income	100%
Marginal propensity to consume — capital income	100%
Labour share multiplier (TFP to labour)	50%
Share of residual TFP competed away to prices	50%
Residual share of TFP to capital	50%
AI/robotics payback period	1 year
AI/robotics depreciation rate	20%
Share of innovation demand to human-centric industries	50%
Demand saturation multiple (M)	2.0
Average wage — human-centric innovation	A$58,300 per FTE
Average wage — AI-augmented innovation	A$135,000 per FTE

Model calibrated to ABS 2023–24 Input-Output Tables, 115 industries. All monetary values in 2023–24 Australian dollars. Results derived from Input-Output Model v22.