

Generative AI and Occupational Change in Australia: A Task-Based CGE Analysis

James A. Lennox*

Janine M. Dixon*

This version: April 14, 2026

Abstract

This paper analyses the long-run labour-market and macroeconomic implications of generative AI (GenAI) using a task- and occupation-based computable general equilibrium (CGE) model of the Australian economy. Unlike partial-equilibrium task-based approaches or CGE studies representing automation as coarse labour–capital substitution, our framework models augmentation and automation as endogenous, competing task technologies and embeds this mechanism in a 45-industry, 97-occupation general-equilibrium structure.

We calibrate the model using separate occupational exposure scores for augmentation and automation potential, and report results from a reference simulation in which GenAI is fully and broadly adopted. Real GDP rises by 29.8%, while welfare rises by only 16.2%, reflecting a substantially more capital-intensive equilibrium—investment’s share of GDP increases from 34.5% to 40.7%. Direct task-level productivity improvements account for roughly one third of the GDP gain, with the majority of the aggregate effect arising through induced capital deepening and general-equilibrium reallocation. The qualitative patterns of these results are robust to alternative assumptions about the level of global export demand.

The dominant labour-market effect is a reconfiguration of task composition *within* occupations rather than large-scale displacement *between* them: approximately 46.7% of tasks are augmented and 9.1% automated, affecting roughly two thirds of all work, with gross task reallocation of a similar magnitude while net employment shifts across occupations remain modest. Net occupational employment shifts are moderate, with losses concentrated in lower-skilled white-collar roles. At the aggregate level, the adjustment is accompanied by a shift in factor incomes: the labour share of income falls, with capital capturing a commensurate increase.

Keywords: Generative artificial intelligence; Computable general equilibrium; Task-based production; Occupational reallocation; Augmentation and automation; Fréchet distribution

JEL classification: C68; J23; J24; O33

1 Introduction

Generative AI (GenAI)—most notably, large language models—exhibits the core characteristics of a general-purpose technology: pervasiveness, technical dynamism, and the capacity to catalyse complementary innovations (Bresnahan and Trajtenberg, 1995; Eloundou et al., 2024). Unlike earlier waves of automation that mainly affected routine tasks (Autor et al., 2003), GenAI’s applicability extends to non-routine cognitive work and the processing of unstructured data across diverse industries. A central economic distinction is whether GenAI is used to *augment* workers’ productivity in performing tasks or to *automate* tasks entirely (Brynjolfsson, 2023; Acemoglu, 2025). These features motivate the framework developed below.

Understanding the long-run structural implications of GenAI is a general-equilibrium problem: task-level cost changes alter factor demands, investment incentives, and the allocation of labour across occupations and industries. We address these interactions with a task-based computable general equilibrium (CGE) model calibrated to Australia, treating occupations as bundles of heterogeneous tasks and linking labour supply to educational qualifications. Within occupations, task-level heterogeneity is represented by correlated Fréchet draws, an extreme-value structure that yields closed-form mode shares and maps empirical exposure indices into the model’s key distributional parameters.

*Centre of Policy Studies, Victoria University, Melbourne, Australia. Corresponding author: james.lennox@vu.edu.au.

Experimental evidence documents substantial productivity gains from GenAI in specific applications—time savings of 26–40% for coding and writing tasks (Noy and Zhang, 2023; Cui et al., 2026)—but aggregate employment effects to date appear limited and heterogeneous (Humlum and Vestergaard, 2025; Brynjolfsson et al., 2025; Liu et al., 2025). Such patterns are consistent with early-stage adoption of a general-purpose technology (Crafts, 2021; Brynjolfsson et al., 2021) and provide limited guidance on how task-level productivity changes will translate into wages, employment, and industry structure once the technology has matured and diffused. Macroeconomic assessments of longer-run impacts differ widely—from aggregate productivity gains of around 0.66% over a decade (Acemoglu, 2025) to effects an order of magnitude larger under more optimistic micro evidence (Aghion and Bunel, 2024). The OECD has estimated annualised gains of 0.4–1.3%, depending on adoption conditions (Filippucci et al., 2025). Our interest lies not in the pace of technology development and adoption but in the long-run effects on wages, occupational employment, and industry structure.

Building on the task-based production literature (e.g. Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2022), we model each occupation as performing a continuum of tasks that can be executed in any one of three modes: with raw labour; with labour *augmented* by AI services; or *automated* using equipment and AI services. Augmentation and automation compete with raw labour as endogenous modes of task production, with the cost-minimising mode determined task by task, rather than imposed *ex ante*.¹ Assuming correlated Fréchet marginals for augmented and automated mode productivities delivers closed-form task-share weights and a CES-type unit cost for the occupational task composite, facilitating our embedding of endogenous mode choice within a multi-sector CGE model (Acemoglu et al., 2025; Cubas et al., 2024).

Our approach makes three contributions. First, we model augmentation and automation as endogenous, competing modes in a multi-sector general-equilibrium framework. Bekkers et al. (2025) impose automation shares exogenously and aggregate occupational exposure into broad skill groups, while Filippucci et al. (2024) model GenAI as sectoral productivity shocks derived off-model from detailed occupational exposure scores. Both approaches mask the occupational reallocation and industry adjustment mechanisms that our framework makes explicit. Second, we extend the task structure to previously automated activities, allowing GenAI to enhance equipment tasks analogously to occupational labour. Third, we calibrate the model using exposure indices from Jobs and Skills Australia (2025a). In addition to being specific to the Australian occupational classification, these data are distinguished from most published indices by providing separate scores for augmentation and automation potential at the occupational level, which allows us to calibrate distinct sets of distributional scale parameters for the two modes.

We calibrate the model using exposure scores for 97 Australian occupations (Jobs and Skills Australia, 2025a,b), embedded in a 45-industry framework that captures price feedbacks through AI services composites of ICT goods and services, electricity, and utilities. The reference case simulation yields a 29.8% long-run increase in real GDP, but the welfare gain is substantially smaller at 16.2% because of the higher investment required. Of the GDP gain, roughly one third is attributable to direct task-level productivity improvements; the remainder arises through capital deepening and general-equilibrium reallocation. Of tasks originally performed by labour, 46.7% are augmented and 9.1% are fully automated in the new equilibrium. Employment shifts between occupations are moderate, with losses concentrated in lower-skilled white-collar roles and gains more widespread across blue-collar occupations, while changes in task composition within occupations are large—equivalent to roughly two thirds of current work. We also assess the sensitivity of these results to the assumed outward shift in export demand curves arising from GenAI adoption abroad.

The remainder of the paper is organised as follows. Section 2 presents the model. Section 3 describes the calibration and data. Section 4 reports results: aggregate macroeconomic outcomes (Section 4.1), industry adjustment (Section 4.2), and occupational reallocation and wages (Section 4.3), followed by a robustness check on the export-demand assumption (Section 4.4). Section 5 concludes.

2 Model

The model is comparative static and represents the economy in long-run equilibrium.

¹Acemoglu (2025) distinguishes four channels through which AI can affect production: extensive-margin automation, new task complementarities, deepening of automation, and creation of new labour-intensive tasks. Task automation and augmentation in our framework correspond to the first two; we also allow for the third but abstract from the fourth.

There is a single representative household that consumes domestic and imported goods and services. Each household member chooses one of 66 qualifications ($q \in \mathbb{Q} \equiv 1, 2, \dots, N_q$) and, conditional on that choice, one of the 97 occupations ($o \in \mathbb{O} \equiv 1, 2, \dots, N_o$) to which she supplies a unit of labour. As the household utility function and the nested logit structure governing labour supply are standard, we summarise them briefly in Section 2.1; full mathematical details are given in [Lennox and Dixon \(2026\)](#). In addition to labour income, households receive income from financial assets, which earn a common economy-wide rate of return.

On the production side, 45 industries ($i \in \mathbb{I} \equiv \{1, 2, \dots, N_i\}$) combine intermediate inputs, capital and up to 97 types of occupational labour to produce corresponding goods. The key innovation is a task-based representation of occupational labour embedded within the value-added nest. Within each occupation–industry pair, a continuum of tasks can be performed using non-AI, AI-augmented or AI-automated modes. AI-using modes require an additional AI services input, modelled as a composite of domestic and imported ICT goods and services, electricity, and other utilities. The production structure is described in Section 2.2.

Tasks differ in their suitability for augmentation and automation. We capture this heterogeneity by drawing task-level AI productivities from a correlated Fréchet distribution and assuming that firms allocate each task to the least-cost mode given equilibrium input prices. Aggregation over the task continuum then yields smooth, price-responsive mode shares and a tractable reduced-form representation of occupation-specific task costs and input requirements. This is developed in Section 2.3.

The model is closed as a long-run equilibrium system in which occupational wages, domestic goods prices, capital stocks, the exchange rate and the aggregate saving rate adjust to satisfy market-clearing and macroeconomic balance conditions. Closure assumptions are summarised in Section 2.4.

2.1 Household behaviour and labour supply

A representative household maximises a nested CES utility function over Armington composites of domestic and imported goods. Aggregate expenditure equals factor income net of savings; the savings rate adjusts endogenously so that the economy-wide rate of return on capital remains at its exogenous long-run level. The utility system and associated demand equations are given in [Lennox and Dixon \(2026\)](#).

The economy is endowed with a fixed mass of workers, each supplying one unit of labour inelastically; we abstract from potential responses of hours, participation, and equilibrium unemployment in order to focus on the long-run reallocation of labour across occupations and industries. Occupational wages act as relative prices that reallocate workers across qualifications and occupations via a two-level nested logit structure.² At the lower level, the share of workers holding a given qualification who choose a particular occupation is a logit function of that occupation’s wage relative to all others available to that qualification group; the responsiveness of these shares to wage differentials is governed by an elasticity set to 2. At the upper level, the population share acquiring each qualification responds in the same logit fashion to a wage index that aggregates the occupational wages accessible to that qualification group; this elasticity is set to 2/3, reflecting the greater persistence of educational investments relative to occupational choices. In both levels, fixed amenity and feasibility weights are calibrated to match baseline employment shares by qualification and by qualification–occupation cell, and zero weights rule out inadmissible matches (e.g. occupations requiring qualifications the worker does not hold). Full details are given in [Lennox and Dixon \(2026\)](#).

2.2 Production and task structure

Figure 1 shows the nesting structure of the production function for industry i .

Each industry combines intermediate inputs, occupational labour, equipment capital and non-equipment capital to produce output.³ At the top level, output is produced by combining intermediates with value added. Intermediates are Armington composites of domestic and imported goods, while value added combines non-equipment capital with a task composite. In simulations we set the top-level and intermediate substitution elasticities to zero (Leontief), so these nests organise demand flows but do not generate substitution at those levels. Figure 1 labels

²Qualifications take time and resources to acquire, and occupation-specific human capital accumulates on the job. The elasticity parameters should therefore be interpreted as capturing long-run steady-state allocations.

³The dwellings industry is an exception, using only structures capital and intermediate inputs.




Figure 1: Illustration of the nested dual cost function for industry i . Leaves are prices for goods ($P_{c,g,i}$), non-equipment capital ($r_{ne,i}$), equipment capital ($r_{eq,i}$), labour (w_o) and AI services (p_s) inputs. Intermediate nodes show prices for composites of these inputs. The relevant substitution elasticities are shown for CES sub-nests. The two levels of nesting below $P_{t,i}$ represent the task-based structure. The elasticity of substitution between all individual tasks is ϵ , while modes are perfectly substitutable within tasks. The lowest level shows the Leontief technologies for augmented and automated task modes.

each nest with the price of the corresponding composite; the unit cost of the task composite $P_{t,i}$ is derived in Section 2.3. Standard CES demand and cost equations for the upper nests are given in [Lennox and Dixon \(2026\)](#).

The task composite T_i aggregates ‘equipment tasks’ (previously automated activities) and occupational tasks (tasks initially performed by occupational labour):

$$T_i = \left[\tilde{\varphi}_i^{1/\epsilon} \int_0^1 \tilde{x}_i(\tau)^{(\epsilon-1)/\epsilon} d\tau + \sum_{o \in \mathbb{O}} \varphi_{o,i}^{1/\epsilon} \int_0^1 x_{o,i}(\tau)^{(\epsilon-1)/\epsilon} d\tau \right]^{\epsilon/(\epsilon-1)} \quad (1)$$

where $\tilde{x}_i(\tau)$ denotes equipment task outputs, $x_{o,i}(\tau)$ denotes occupational task outputs, and ϵ is the elasticity of substitution between tasks. The weights $\tilde{\varphi}_i$ and $\varphi_{o,i}$ capture industry-specific differences in the quantities of each task type required. Unlike models with a fixed automation cutoff, mode assignment here is not determined by a threshold on the task index τ but by the realised productivity draw relative to input price ratios. Consequently the integrals in (1) run from 0 to 1 for all modes, with effective mode shares determined endogenously through the productivity distributions developed below.

Occupational tasks may be performed without AI, with labour augmented by AI services, or with equipment automated using AI services. The task production function for occupation o in industry i is:

$$x_{o,i}(\tau) = x_{o,i}^n(\tau) + \psi_o^g(\tau) x_{o,i}^g(\tau) + \psi_o^a(\tau) x_{o,i}^a(\tau) \quad (2)$$

where the superscripts n, g, and a denote non-AI, augmented, and automated modes respectively. The task-specific productivity parameters $\psi_o^g(\tau)$ and $\psi_o^a(\tau)$ capture variation across tasks in the effectiveness of the two AI-using modes.

Production in each mode combines primary inputs with AI services in fixed proportions:

$$\begin{aligned} x_{o,i}^n(\tau) &= A_{o,i}^L n_{o,i}^n(\tau) \\ x_{o,i}^g(\tau) &= \min \left\{ A_{o,i}^L n_{o,i}^g(\tau), A_{o,i}^G s_{o,i}^g(\tau) \right\} \\ x_{o,i}^a(\tau) &= \min \left\{ A_i^K k_{o,i}^a(\tau), A_{o,i}^S s_{o,i}^a(\tau) \right\} \end{aligned} \quad (3)$$

where labour, equipment and AI services are denoted by l, k and s respectively. The A^* terms are factor-augmenting productivity parameters that allow for industry- and occupation-specific calibration; they are not varied in our simulations. In the augmented mode, labour remains the primary input and AI services raise worker productivity. In the automated mode, equipment replaces labour as the primary input and AI services raise the productivity of that equipment. We therefore distinguish between *equipment*, which refers to the dedicated capital used in automated tasks, and *AI services*, which refers to inference compute and related ICT inputs consumed in proportion to AI-enabled task volume.

AI services are not produced as a separate industry. Instead, they are represented as a fixed-proportions composite of existing domestic and imported goods, with cost shares concentrated in ICT services and goods, electricity, and other utilities, as reported in Table 2. This specification allows AI use in production to feed back into the broader economy through higher demand for these underlying inputs, rather than treating AI as a pure efficiency shifter.

With task modes being perfect substitutes, each task is allocated to the mode with lowest unit cost. Given the wage w_o , equipment rental $r_{eq,i}$, and AI services price p_s , the unit cost for occupational task τ is:

$$\begin{aligned} p_{o,i}(\tau) &= \min \left\{ \omega_{o,i}^n, \frac{\omega_{o,i}^g}{\psi_o^g(\tau)}, \frac{\omega_{o,i}^a}{\psi_o^a(\tau)} \right\} \\ \omega_{o,i}^n &\equiv w_o / A_{o,i}^L, \quad \omega_{o,i}^g \equiv w_o / A_{o,i}^L + p_s / A_{o,i}^G, \quad \omega_{o,i}^a \equiv r_{eq,i} / A_i^K + p_s / A_{o,i}^S. \end{aligned} \quad (4)$$

The cost ratios that determine mode thresholds are

$$\hat{\omega}_{o,i}^g \equiv \omega_{o,i}^g / \omega_{o,i}^n = 1 + \frac{p_s A_{o,i}^L}{A_{o,i}^G w_o}, \quad \hat{\omega}_{o,i}^a \equiv \omega_{o,i}^a / \omega_{o,i}^n = \frac{r_{eq,i} A_{o,i}^L}{A_i^K w_o} + \frac{p_s A_{o,i}^L}{A_{o,i}^S w_o}. \quad (5)$$

Equipment tasks—activities already performed by automated equipment prior to GenAI—may also be enhanced by AI services embodied in that equipment.⁴ The production function for equipment task τ in industry i is

$$\begin{aligned} \tilde{x}_i(\tau) &= \tilde{x}_i^n(\tau) + \tilde{\psi}_i^g(\tau) \tilde{x}_i^g(\tau) \\ \tilde{x}_i^n(\tau) &\equiv A_i^K \tilde{k}_i^n(\tau) \\ \tilde{x}_i^g(\tau) &\equiv \min \left\{ A_i^K \tilde{k}_i^g(\tau), \tilde{A}_i^S \tilde{s}_i^g(\tau) \right\} \end{aligned} \quad (6)$$

where $\tilde{k}_i^n(\tau)$ and $\tilde{k}_i^g(\tau)$ are equipment inputs in the non-AI and AI-enhanced modes, $\tilde{s}_i^g(\tau)$ is AI services input, A_i^K is an equipment productivity parameter common to both modes, and $\tilde{\psi}_i^g(\tau)$ is the task-specific productivity draw. Unlike automation of occupational tasks, the AI-enhanced equipment mode does not substitute a new primary input: it raises the productivity of equipment already in use, requiring AI services in fixed proportions. In our main simulations we set AI enhancement of equipment tasks to zero, since our empirical exposure estimates are limited to occupational tasks; a supplementary simulation illustrates the effects of allowing this channel.

The unit cost per equipment task is

$$\begin{aligned} \tilde{p}_i(\tau) &= \min \left\{ \tilde{\omega}_i^n, \frac{\tilde{\omega}_i^g}{\tilde{\psi}_i^g(\tau)} \right\} \\ \tilde{\omega}_i^n &\equiv r_{eq,i} / A_i^K, \quad \tilde{\omega}_i^g \equiv r_{eq,i} / A_i^K + p_s / \tilde{A}_i^S. \end{aligned} \quad (7)$$

⁴We label the AI-using mode for equipment tasks g, reserving a for tasks that switch their primary input from labour to equipment. The AI-using mode here augments equipment already in use rather than automating anything further.

Note that $\tilde{\omega}_i^n$ equals the equipment component of $\omega_{o,i}^a$ and $\tilde{\omega}_i^g$ has the same structure—equipment plus AI services—but with equipment-specific AI productivity $\tilde{A}_i^s$. An equipment task is AI-enhanced if and only if $\tilde{\psi}_i^g(\tau) > \hat{\omega}_i^g \equiv \tilde{\omega}_i^g / \tilde{\omega}_i^n$, i.e. the AI productivity gain exceeds the ratio of combined (equipment plus AI) to pure-equipment unit costs.

Cost minimisation over the task composite (1) yields demand for each individual task output. Letting $P_{t,i}$ denote the unit cost of the task composite, the demand for occupational task τ of type (o, i) is

$$x_{o,i}(\tau) = \varphi_{o,i} \left(\frac{P_{t,i}}{p_{o,i}(\tau)} \right)^\epsilon T_i, \quad (8)$$

and the demand for equipment task τ in industry i is

$$\tilde{x}_i(\tau) = \tilde{\varphi}_i \left(\frac{P_{t,i}}{\tilde{p}_i(\tau)} \right)^\epsilon T_i. \quad (9)$$

Since unit costs $p_{o,i}(\tau)$ and $\tilde{p}_i(\tau)$ depend on task-specific productivity draws, these demand schedules are heterogeneous across tasks and must be integrated over the productivity distributions to obtain aggregate factor demands.

2.3 Task productivity distributions and reduced-form aggregation

To derive aggregate mode shares and task share weights, we specify distributions for task productivities $\psi_o^g(\tau)$ and $\psi_o^a(\tau)$. Our data suggest that these productivities are positively correlated within occupations: tasks with high potential for augmentation also tend to have high potential for automation.⁵ To capture this pattern, we assume ψ_o^g and ψ_o^a follow a correlated Fréchet distribution—Fréchet marginals linked by a Gumbel copula.⁶

The joint CDF of productivities is:

$$F(\bar{\psi}_o^g, \bar{\psi}_o^a) = \exp \left[- \left(\left(\frac{\kappa_{o,i}^g}{\bar{\psi}_o^g} \right)^{\alpha_o \theta_o} + \left(\frac{\kappa_{o,i}^a}{\bar{\psi}_o^a} \right)^{\alpha_o \theta_o} \right)^{1/\theta_o} \right] \quad (10)$$

where $\kappa_{o,i}^g$ and $\kappa_{o,i}^a$ are occupation-specific scale parameters, α_o is the shape parameter (smaller values imply heavier tails), and $\theta_o \in [1, \infty)$ governs the strength of upper-tail dependence (Kendall's tau equals $1 - 1/\theta_o$). In our reference case, α_o and θ_o are set to common values across occupations.

This distributional specification yields closed-form expressions for the endogenously determined task share weights that enter the reduced-form CES production function. The weights for tasks performed without AI are

$$\Gamma_{o,i}^n(\omega_{o,i}) = \exp \left[- \left(\left(\frac{\kappa_{o,i}^g}{\hat{\omega}_{o,i}^g} \right)^{\alpha_o \theta_o} + \left(\frac{\kappa_{o,i}^a}{\hat{\omega}_{o,i}^a} \right)^{\alpha_o \theta_o} \right)^{1/\theta_o} \right] \quad (11)$$

where $\hat{\omega}_{o,i}^g$ and $\hat{\omega}_{o,i}^a$ are the cost ratios defined in (5). Augmented and automated task share weights involve the Gamma and incomplete Gamma functions evaluated at these cost ratios, with the copula parameter θ_o mediating the interaction between the two competing AI modes. Derivations and the full expressions—including the analogous but simpler univariate structure for equipment tasks—are given in [Lennox and Dixon \(2026\)](#).

Using these weights, the explicit task production function (2) can be expressed as the following reduced-form CES production function:

$$T_i = \left[\tilde{\varphi}_i^{1/\epsilon} \sum_{m \in \{n,g\}} (\tilde{\Gamma}_i^m)^{1/\epsilon} (\tilde{x}_i^m)^{(\epsilon-1)/\epsilon} + \sum_{o \in \mathbb{O}} \varphi_{o,i}^{1/\epsilon} \sum_{m \in \{n,g,a\}} (\Gamma_{o,i}^m)^{1/\epsilon} (x_{o,i}^m)^{(\epsilon-1)/\epsilon} \right]^{\epsilon/(\epsilon-1)} \quad (12)$$

⁵The underlying task-level analysis suggests substantial overlap in the tasks that are augmentable and automatable within occupations (Jobs and Skills Australia, personal communication, May 2025.).

⁶The Fréchet distribution with scale κ and shape α is a heavy-tailed continuous probability distribution on $(0, \infty)$ with CDF $F_X(x) = \exp[-(\kappa/x)^\alpha]$ and PDF $f_X(x) = (\alpha/\kappa)(\kappa/x)^{1+\alpha} \exp[-(\kappa/x)^\alpha]$. The two-dimensional Gumbel copula $C(U, V) = \exp[-(-\ln U)^\theta + (-\ln V)^\theta]^{1/\theta}$ exhibits upper-tail dependence, capturing the tendency for large values of both variables to occur jointly. A joint CDF with given marginals is obtained by substituting their CDFs into the copula.

in which the weights $\Gamma_{o,i}^m$ are endogenous functions of technology parameters and equilibrium prices. The corresponding dual cost function—the unit cost of the task composite $P_{t,i}$ that enters the value-added nest in Figure 1—takes the CES form

$$P_{t,i} = \left[\tilde{\varphi}_i \sum_{m \in \{n,g\}} \tilde{\Gamma}_i^m (\tilde{\omega}_i^m)^{1-\epsilon} + \sum_{o \in \mathbb{O}} \varphi_{o,i} \sum_{m \in \{n,g,a\}} \Gamma_{o,i}^m (\omega_{o,i}^m)^{1-\epsilon} \right]^{1/(1-\epsilon)} \quad (13)$$

where $\omega_{o,i}^m$ and $\tilde{\omega}_i^m$ are the mode-specific unit factor costs defined in equations (4) and (7). This is the price that appears as the cost of the task composite in the value-added nest. Because the weights $\Gamma_{o,i}^m$ respond endogenously to relative prices, $P_{t,i}$ is more price-elastic than a conventional CES cost function with fixed distribution parameters. Factor demands then follow from Shephard’s lemma: labour is used in the non-AI and augmented modes, equipment in the automated occupational and equipment task modes, and AI services in all AI-using modes (Lennox and Dixon, 2026).

2.4 Closure

The model is solved under a long-run closure in which factor supplies and macroeconomic balances adjust to support a new comparative-static equilibrium. The economy-wide net rate of return on capital is fixed exogenously at its long-run level, and the capital stock adjusts to satisfy this required return. The aggregate saving rate then adjusts to finance the corresponding level of investment, given a fixed trade-balance-to-GDP ratio.⁷ Capital is mobile in the long run, so industry capital stocks expand or contract to meet demand at the common required rate of return. Rental prices therefore move with endogenous investment-good prices rather than with capital scarcity per se.

Market-clearing operates in the usual way. In labour markets, demand for each occupation-specific labour type must equal the supply generated by the household’s qualification and occupation choices, with occupational wages adjusting to clear these markets. In product markets, domestic output is allocated across intermediate use, household demand, investment, government and exports, while domestic goods prices adjust to equate supply and demand for each locally produced good. Imports are differentiated from domestic goods through Armington aggregation, so domestic and imported varieties are imperfect substitutes rather than a single pooled commodity.

AI services enter the closure only through their derived demands. Because they are defined as a fixed-proportions composite of existing domestic and imported goods rather than as a separate produced commodity, increased AI use raises demand for the underlying ICT, utility and related inputs that comprise the bundle. This creates additional general-equilibrium feedbacks through the corresponding goods markets and through the prices of AI-using production modes.

External balance is determined by downward-sloping export demand schedules with exogenous demand shifters. In the simulations, these shifters are moved outward to capture stronger foreign demand associated with global GenAI adoption and lower international trade costs. The real exchange rate adjusts to maintain a fixed trade-balance-to-GDP ratio, so export volumes, export prices and the terms of trade are all endogenous. Real GDP and its associated price deflators are computed as Divisia indices along the GEMPACK solution path (Dixon et al., 1982; Horridge, 2000). Further details of the market-clearing conditions and macro closure are given in Lennox and Dixon (2026).

3 Calibration and Data

This section describes the calibration procedure and data sources. The baseline database, exposure measures, key parameter choices, and AI services cost structure follow Lennox and Dixon (2026) but are presented more compactly.

⁷The model does not distinguish between domestic and foreign ownership of capital; households hold claims on aggregate capital income at the common rate of return.

3.1 Data and baseline projection

Australian input–output and employment data for 2019 are projected to 2050 using the Victoria University Employment Forecasting (VUEF) model (Dixon, 2022). These projections focus on the evolution of qualifications in the labour force and the allocation of qualified labour among occupations and industries, providing a natural starting point for long-run comparative-static simulations. The baseline does not incorporate impacts of GenAI.⁸ We interpret the comparative-static simulations as characterising differences between a no-GenAI and a GenAI long-run equilibrium around the 2050 projected economy.

VUEF is calibrated using ABS input–output and national accounts data, and provides the industry structure of intermediate demands, value added, and final demands required to calibrate production function parameters and market-clearing conditions. Employment by occupation and industry is drawn from the ABS Labour Force Survey and Census, covering 97 three-digit ANZSCO occupations and 121 production industries. To simplify exposition, we suppress some detailed VUEF features less relevant to our analysis, including margin demands and the distinction between public and private consumption.⁹ We aggregate the projected database from 121 to 45 industries.

To distinguish capital used in automated tasks from non-equipment capital, we split industry capital and investment into equipment and non-equipment categories. ‘Equipment’ is defined using four producing industries and competing imports: (i) professional, scientific, computer and electronic equipment manufacturing; (ii) electrical equipment manufacturing; (iii) domestic appliance manufacturing; and (iv) specialised and other machinery and equipment manufacturing. Transport equipment is deliberately excluded, as the aim is to focus on computers and other machinery plausibly used to automate occupational tasks.

3.2 Parameters

The shape parameter α_o and dependence parameter θ_o are set to common values across occupations in the reference case (see Table 1). These values imply moderately heavy tails and moderate upper-tail dependence. Given the limited empirical basis for these parameters, sensitivity analysis in Lennox and Dixon (2026) explores a wide range: $\alpha_o \in \{3, 5, 7\}$ and $\theta_o \in \{5/3, 5/2, 5\}$.

The elasticity of substitution between tasks ϵ is set to a value consistent with estimates in the task-based literature (e.g. Humlum and Vestergaard, 2025; Acemoglu and Restrepo, 2022) (see Table 1). The elasticity between capital and the task composite σ_v is set similarly. At higher levels of the nesting structure we set $\sigma_y = 0$ and $\sigma_n = 0$, so the reference model has Leontief aggregation between value added and intermediates and across broad intermediate composites. A common Armington elasticity σ_c is applied to all goods in all uses.

The elasticity of substitution in the top-level consumption composite σ_u and the specification of the investment goods composite (Leontief) are reported in Table 1.

Workers are allocated across occupations and qualifications using modest long-run elasticities (Table 1). Export demand elasticities differ by commodity group (Table 1). The reference simulation shifts all export-demand schedules outward by 10%, reflecting the assumption that GenAI adoption is global rather than uniquely Australian. Sensitivity to this assumption is the only robustness exercise in the present paper.

3.3 GenAI exposure scores

Task productivity distributions are calibrated using task-based measures of generative AI exposure for 97 three-digit ANZSCO occupations (Jobs and Skills Australia, 2025a). These measures build on the methodology developed by the International Labour Organization (ILO) for assessing generative AI exposure at the task level (Gmyrek et al., 2023). The approach evaluates the technical feasibility of applying generative AI to detailed occupational tasks rather than assigning exposure directly at the occupation level. In the ILO framework, task descriptions associated with each occupation are assessed using large language models to determine whether generative AI could plausibly perform the task or assist a human worker in carrying it out. Each task is assigned a

⁸A more limited version of the task-based features presented here is incorporated in VUEF and is used to generate dynamic projections of GenAI impacts in Jobs and Skills Australia (2025b).

⁹Our simplifications include collapsing transport and other margin demands into the general intermediate inputs nest, and combining public and private consumption.

Table 1: Summary of key model parameters

Parameter	Description	Reference value	Sensitivity range
<i>Task productivity distributions</i>			
α_o	Fréchet shape (tail thickness)	5	{3, 5, 7}
θ_o	Gumbel dependence	2.5	{5/3, 5/2, 5}
$\kappa_{o,i}^g, \kappa_{o,i}^a$	Fréchet scale (occupation-specific)	Calibrated	—
<i>Production nests</i>			
ϵ	Task substitution elasticity	0.5	{0.25, 0.5, 0.75}
σ_v	Capital–task composite elasticity	0.5	{0.25, 0.5, 0.75}
σ_y	Output (VA vs. intermediates)	0	—
σ_n	Intermediate goods composite	0	—
σ_c	Armington (domestic vs. imported)	2	—
σ_u	Consumption composite	0.9	—
σ_i	Investment goods composite	0	—
<i>Labour supply</i>			
σ_o	Occupational choice elasticity	2	—
σ_q	Qualification choice elasticity	2/3	—
<i>Trade and external closure</i>			
η_g	Export demand elasticity (by commodity group)	4 (goods & business services) 2 (consumer services)	—
$\Delta \bar{E}_g$	Export demand shift	+10%	{0, 10, 20}

score between zero and one representing its potential exposure to generative AI, and occupation-level measures are obtained by aggregating these task-level scores (Gmyrek et al., 2023).

Jobs and Skills Australia (JSA) adapts this framework to the Australian labour market using the Australian and New Zealand Standard Classification of Occupations (ANZSCO). Task descriptions for each ANZSCO four-digit occupation are evaluated using a similar procedure, producing task-level exposure measures subsequently aggregated to the occupational level. In contrast to the original ILO implementation, the JSA methodology distinguishes explicitly between two channels through which generative AI may affect work tasks: augmentation, where AI tools assist human workers in performing a task more effectively, and automation, where the task may potentially be carried out by AI systems with minimal human input. The resulting occupational indices therefore report separate augmentation and automation exposure scores for each occupation (Jobs and Skills Australia, 2025a). These scores measure the technical potential for generative AI to affect tasks within an occupation rather than the probability of adoption.

Figure 2 shows empirical cumulative distribution functions (ECDFs) of augmentation and automation exposure across occupations. Solid lines report employment-weighted ECDFs; dotted lines treat each occupation equally.

Three features of the distributions are economically relevant. First, augmentation exposure is systematically higher than automation exposure: the employment-weighted augmentation mean (0.638) exceeds the automation mean (0.383) by roughly 26 percentage points, a difference that carries through to the dominance of augmented over automated task shares in the model simulations. Second, automation exposure is mildly right-skewed—the employment-weighted mean (0.383) exceeds the median (0.357)—indicating that a relatively small set of highly exposed occupations raises the average. Third, augmentation exposure exhibits the opposite pattern, with the mean (0.638) lying below the median (0.661), consistent with a small number of low-exposure occupations pulling down an otherwise broadly high level of exposure. The divergence between weighted and unweighted ECDFs for both series indicates that larger occupations tend to have higher exposure scores, amplifying their aggregate importance.

We interpret these exposure scores as describing the long-run technological potential of each mode in isolation and in partial equilibrium. This interpretation abstracts from competition between augmentation and automation, AI services costs, and relative input price changes—factors that are all accounted for when we simulate the model. Specifically, we calibrate the scale parameters $\kappa_{o,i}^g$ and $\kappa_{o,i}^a$ so that, in the absence of the competing mode and at zero AI services cost, the share of tasks suitable for augmentation or automation matches the corresponding exposure score. Concretely, given a target augmentation exposure $\bar{e}_o^g$, we solve $\Gamma_{o,i}^n(\omega_{o,i})|_{\kappa_{o,i}^a=0, p_s=0} = 1 - \bar{e}_o^g$ for $\kappa_{o,i}^g$ (and analogously for $\kappa_{o,i}^a$ using $\bar{e}_o^a$). This single-mode calibration isolates the scale parameters from the




Figure 2: Empirical cumulative distribution functions of augmentation and automation exposure across occupations, unweighted and weighted by initial occupational employment. Vertical lines indicate the employment-weighted mean and median for each series.

copula interaction and from general-equilibrium price effects, both of which operate when the full model is solved. Full occupation-level scores are given in [Lennox and Dixon \(2026\)](#).

3.4 AI services cost structure

We represent AI services as a composite of existing modelled domestic and imported goods rather than as a fully-fledged production industry. This simplifying device reflects the considerable uncertainty surrounding both the current cost structure of GenAI and its future evolution. Algorithmic efficiency gains alone have reduced inference costs on AI benchmarks by roughly a factor of three per year ([Gundlach et al., 2025](#)), and specialised smaller language models may out-compete very large generalist models ([Wang et al., 2024](#)), further lowering costs. Although graphics processing unit (GPU)-based data centres are particularly energy- and water-intensive, improvements in chip design and cooling technologies are likely to mitigate these costs over time ([OECD, 2025](#)). The assumed cost structure is summarised in Table 2.

We distinguish between ‘upstream’ and ‘downstream’ AI activities. Upstream activities encompass fundamental AI research, the development of large language models, and internationally marketed AI services. Downstream activities comprise the adaptation and application of these products to firm-specific needs, as well as the compute and storage required to support inference. We assume that upstream activities remain dominated by large foreign firms, while downstream services are domestically provided. Downstream AI services also require imported inputs, notably GPUs and other ICT equipment; as we do not model stocks of AI-specific capital, these imported investment goods are treated as intermediate inputs.¹⁰ A more complete treatment of AI services supply is left for future work.

¹⁰This treatment slightly understates GDP, as expenditures on AI-related ICT capital are omitted from aggregate investment and the associated returns are omitted from domestic income.

Division	Subdivisions	Industry name	Share within block (%)	Share of total (%)	Origin
<i>Upstream AI development and training (20% of total)</i>					
M	70	Computer Systems Design & Services	100	20.0	Imported
<i>Downstream services (10% of total)</i>					
M	70	Computer Systems Design & Services	100	10.0	Domestic
<i>Downstream compute and storage (70% of total)</i>					
J	58–59	Internet & Telecommunications	30	21.0	Domestic
D	26	Electricity	18	12.6	Domestic
D	27–29	Other Utilities	2	1.4	Domestic
C	24	Machinery & Equipment Manufacturing	50	35.0	Imported
Total				100.0	

Table 2: Cost structure of AI services represented as a composite of existing modelled industries.

4 Results

The results below describe a long-run equilibrium in which GenAI—characterised by the occupational exposure scores and modelling assumptions described in Section 3—is fully integrated into production. The model captures the economic structure implied by the technology’s technical potential, abstracting from diffusion dynamics and institutional constraints. We present aggregate, industry, and occupational results for the reference case ($\alpha_o = 5$, $\theta_o = 2.5$, +10% export demand shift), and conclude with a robustness check on the export-demand assumption.

4.1 Reference case simulation

The reference case simulation yields a 29.8% increase in real GDP but only a 16.2% increase in welfare, as measured by the long-run change in the real consumption composite. The gap reflects a crucial distinction between output and living standards under capital-deepening growth: in the new equilibrium, investment’s share of GDP rises from 34.5% to 40.7%, so much of the extra output is absorbed as the additional saving required to sustain the higher capital stock rather than flowing to households as consumption. Households nonetheless benefit from the resulting productivity gains, primarily through lower consumer prices—which raise real wages given nominal wage adjustments—and, to a lesser extent, through increases in consumption. The reference case therefore does not describe a pure productivity windfall; it describes a structural transformation in which the economy produces more output but must invest substantially more to do so.

Table 3 decomposes the GDP gain on both the income and expenditure sides. On the income side, task-level technical change—the direct productivity gains from augmentation and automation—accounts for roughly one third of the total (10.1 pp). The majority of the aggregate output gain therefore does not arise from the direct substitution of AI for labour, but from the capital deepening that AI-driven automation necessitates. More intensive use of non-equipment capital contributes 12.3 pp, equipment in previously automated tasks 1.2 pp, and equipment in newly automated occupational tasks 6.1 pp. Equipment capital deepening is driven directly by task automation—in which equipment displaces labour—and reinforced by substitution higher in the production structure: GenAI raises wages relative to the prices of equipment and AI services, encouraging further automation, while equipment and AI services prices in turn rise relative to non-equipment capital, inducing broader capital deepening through the value-added nest. Compositional shifts across industries amplify these within-industry mechanisms. With a fixed workforce, labour reallocation contributes negligibly to aggregate GDP, although the underlying gross flows are large: the value of labour shifting from non-AI to augmented tasks alone (measured as the gross value of reallocated labour services) is roughly equal to the GDP gain.

On the expenditure side, consumption contributes only 10.8 pp of the 29.8% GDP gain, while investment contributes 16.5 pp and net trade 2.5 pp. The rising import intensity of production is driven primarily by demand for machinery and equipment, used both as investment goods and to deliver AI services. With the trade-balance-to-GDP ratio held fixed, export volumes must expand, and by more than import volumes, since finite export demand elasticities imply a terms-of-trade deterioration that moderates the net trade gain. The higher household saving rate required to finance the elevated investment share is the proximate mechanism linking the GDP and welfare outcomes: more output is produced, but a larger fraction is immediately reinvested rather than consumed.

Table 3: Income- and expenditure-side decomposition of real GDP in the reference case GenAI simulation

Income-side decomposition		Expenditure-side decomposition	
Component	pp contribution	Component	pp contribution
Labour (composition)	0.0	Real consumption	10.8
Non-AI tasks	-31.3	Real investment	16.5
Augmented tasks	31.3	Other capital goods	10.3
Capital	19.7	Equipment goods	6.2
Non-equipment capital	12.3	Net trade	2.5
Equipment: previously automated tasks	1.2	Real exports	11.4
Equipment: occupational tasks	6.1	Real imports	-8.9
Technical change (GenAI)	10.1		
Real GDP (income)	29.8	Real GDP (expenditure)	29.8

Notes: Each entry gives a contribution in percentage points to the overall change in GDP. Figures are rounded; individual entries may not sum to subtotals or totals.

The labour share of income falls from 60.0% to 52.4% of GDP, with the capital share rising commensurately—a shift that benefits asset owners disproportionately and is the principal distributional channel through which GenAI reshapes factor incomes.

4.2 Industry adjustment

Task adoption differs substantially across industries, primarily because occupational composition varies: industries that employ more professionals and managers inherit their higher augmentation exposure, while those relying on clerical and routine workers inherit higher automation exposure. Figure 3 plots augmented against automated task shares for each industry in the reference case, with industries identified by their ANZSIC Division letter. Across all industries, augmentation shares exceed automation shares, reflecting the systematically higher augmentation exposure in the source data. At the same time, equilibrium price adjustments—declining equipment prices relative to wages—work in the opposite direction, partially offsetting this calibrated advantage and generating the mode competition visible in the figure. Heterogeneity across industries is nonetheless considerable: professional and business services Divisions (M, K, L) are strongly augmentation-dominant, while goods-producing Divisions (B, C, D) and some services with more routine content (F, G, N) sit closer to the automation axis. Crucially, this dispersion arises endogenously from occupational composition interacting with general-equilibrium input price adjustments, not from any industry-specific technology assumptions.

Figure 4 decomposes industry-level changes in value added into contributions from technical change and factor inputs. All industries experience positive long-run gains, but there is substantial heterogeneity in both magnitude and source. Capital-intensive and export-oriented industries exhibit the strongest increases, reflecting a combination of capital deepening and demand expansion. Mining (Division B) benefits from both its initial factor intensities and the outward shift in export demand. Industries supplying AI-related services, such as Internet and Telecommunications (Division J), also expand strongly. The very large percentage gain for Machinery and Equipment Manufacturing reflects its small initial base combined with strong demand for equipment investment; this result is sensitive to the assumption that the equipment needed to automate occupational tasks has a similar composition to equipment currently used for other purposes. The relative contribution of technical change varies considerably and is not simply proportional to overall value-added gains: in several consumption-oriented industries (Divisions P to S) the technical-change contribution is large relative to value added, yet overall gains are modest, because these industries face relatively weak demand growth. This is an instance of Baumol’s cost disease (Baumol, 1967): when productivity gains are concentrated in sectors facing inelastic or saturating demand, resources shift toward lower-productivity activities, moderating aggregate growth.

Figure 5 attributes value-added changes to demand-side and input–output effects using the Leontief shift-share decomposition described in Appendix A. For most industries, the dominant demand-side driver is the general expansion of economic activity. Changes in the composition of final demand make significant positive contributions to export-oriented industries (Division B) and investment-intensive industries (Division C machinery, Division E construction), while consumption-oriented industries (Divisions O, P, Q, R, W) face significant negative compositional effects—the counterpart, on the demand side, of the Baumol pattern described above. For industries




Figure 3: Automated versus augmented task shares by industry in the reference case. Each point is labelled by its ANZSIC Division letter.




Figure 4: Industry-level changes in value added, decomposed into contributions from technical change and factor inputs. Bar widths are proportional to initial industry value added.

supplying equipment and AI services, import substitution and input intensity effects are also important: in Machinery and Equipment Manufacturing the dominant driver is the compositional shift toward equipment investment, while in Computer Systems Design and Services (Division M) rising AI services intensity makes a strong positive contribution, partly offset by intermediate import substitution.




Figure 5: Industry-level changes in value added, decomposed into contributions from demand and supply-side (input–output) effects. Bar widths are proportional to initial industry value added.

4.3 Occupational task adoption, reallocation, and wages

The industry-level patterns reflect underlying differences in occupational task adoption. Figure 6 plots augmented against automated task shares for all 97 three-digit ANZSCO occupations, coloured by major group. Clerical and administrative workers (group 5) have the highest automation shares and sit furthest above the diagonal, reflecting both their high automation exposure scores and the model’s endogenous response to the relatively lower wages in that group, which makes automation cost-competitive. Managers and professionals (groups 1 and 2) are strongly augmentation-dominant, clustered at high augmentation and low automation. Blue-collar groups (technicians and trades, machinery operators and drivers, labourers; groups 3, 7, 8) show moderate augmentation and low automation, consistent with the physical-task content of their work limiting AI applicability in the current technological regime. These patterns reflect the interaction between occupation-specific exposure scores and the general-equilibrium cost thresholds: across all occupations, 46.7% of tasks are augmented and 9.1% are automated in the reference case—a ratio that echoes the distributional asymmetry visible in Figure 2.

These task-level patterns translate into moderate but heterogeneous employment shifts. Figure 7 compares modelled long-run employment changes due to GenAI with observed employment changes over the past decade. Net reallocation is largest in blue-collar occupations (up to +13%) and most negative in lower-skilled white-collar occupations (down to –22%), where automation exposure is highest. Relative to historical experience these changes are moderate: employment declined by more than 15% in seven occupations and increased by more than 50% in twenty-one occupations in the decade to 2024. A notable exception is delivery drivers, a blue-collar occupation that has grown strongly in recent years but is modelled as contracting under GenAI, consistent with its high automation exposure score and the economy-wide compression of wages relative to equipment rental costs.

Table 4 decomposes employment changes by major ANZSCO group. The first five columns decompose changes in industry labour productivity: the Non-AI and Augmented column pairs each capture task-mode substitution at the extensive margin (tasks shifting between modes) and the intensive margin (changes in how intensively remaining tasks in each mode are used), while Other Input Substitution captures changes in labour demand arising from substitution between the task composite and non-task capital as relative prices adjust. The remaining columns decompose output effects into an industry output-share component and an aggregate demand scale effect.

The key finding is that within-occupation task reallocation dominates between-occupation employment shifts by a wide margin. Extensive-margin reallocation from non-AI to augmented or automated modes generates labour savings of approximately 11.4 million worker equivalents—roughly two thirds of all work currently performed. This effect is modestly reinforced on the intensive margin, as tasks remaining in non-AI mode are used less intensively; declining equipment prices simultaneously increase competitive pressure from automation, while some




Figure 6: Augmented versus automated task shares by occupation in the reference case, coloured by ANZSCO major group. Each point represents one three-digit occupation. Ellipses show the one-standard-deviation interval for each group.




Figure 7: Modelled employment growth due to GenAI (vertical axis) versus observed employment growth 2014–24 (horizontal axis) across 97 three-digit ANZSCO occupations. Bubble sizes reflect 2019 employment.

displaced work is reallocated to augmented tasks within the same occupational groups. These large gross reallocations are primarily offset by growth in aggregate demand, though input substitution and demand composition effects are large relative to the net employment changes across occupations. The overall total is zero by construction, with aggregate labour supply fixed. Between-occupation reallocation amounts to at most a few percentage points of total employment, while within-occupation task restructuring touches the majority of work—confirming that the primary adjustment burden falls on workers adapting to changed task mixes and skill demands within their current role, not on the economy’s capacity to move workers between occupations.

Table 4: Decomposition of occupational employment changes by major ANZSCO group (’000 worker equivalents). The total workforce comprises 17.3 million persons.

	Labour productivity						Industry output			Total
	Non-AI		Augmented		Other input substitution	Subtotal	Output shares	Output scale	Subtotal	
	Extensive margin	Intensive margin	Extensive margin	Intensive margin						
Managers	-2398	-193	1691	-69	20	-949	57	934	992	42
Professionals	-4470	-345	3121	-130	31	-1792	17	1696	1713	-78
Tech. & Trades	-1199	-130	847	-34	3	-513	81	556	638	124
Services	-603	-65	453	-16	3	-228	-82	307	224	-4
Clerical & Admin	-1303	-54	740	-22	10	-629	3	438	441	-187
Sales	-401	-29	256	-9	5	-177	-1	155	154	-23
Operators & Drivers	-531	-66	348	-14	4	-259	67	267	334	75
Labourers	-269	-47	190	-7	1	-133	10	174	185	52
Total	-11178	-933	7650	-305	82	-4684	154	4530	4684	0

Wage outcomes reinforce this picture. Figure 8 plots modelled real wage changes against full-time weekly wages in 2025. The correlation between wage levels and wage changes is weak (16%) and dispersion is highest among lower-wage occupations. Clerical and administrative occupations experience the weakest wage gains—with Keyboard Operators seeing a slight decline—while blue-collar occupations experience the strongest. Applying the modelled wage changes to 2025 wages, the Gini index for occupational wages rises from 0.155 to 0.163—a modest increase in dispersion.

The near-universal real wage increase reflects the economy-wide productivity gain: lower production costs reduce the price level relative to nominal wages, raising real purchasing power even in occupations with little direct AI exposure. The non-linear mapping from calibrated exposure scores to equilibrium task-mode shares—via productivity draws relative to input price ratios—generates wage responses that are not proportional to initial exposure. Occupations where augmentation dominates benefit additionally from the productivity premium of AI-assisted work; those where automation is more prevalent face competitive displacement of labour from many tasks, which moderates wage growth even as aggregate productivity rises. The more consequential distributional implication, however, arises through factor shares: as documented in Table 3, the labour share of income falls from 60.0% to 52.4% of GDP, with the capital share rising commensurately. This shift—equivalent to a transfer of roughly 7.6 percentage points of GDP from labour to capital—benefits asset owners disproportionately and likely dominates the modest widening of occupational wage dispersion as the primary channel through which GenAI reshapes the income distribution.

4.4 Export-demand sensitivity

In addition to the domestic technology shocks informed by the exposure indices, we allow for a uniform outward shift in all export demand curves to reflect GenAI adoption abroad. The reference case assumes a +10% shift, which we consider conservative: for incomplete-adoption scenarios, [Bekkers et al. \(2025\)](#) find real exports rising on the order of 10% for primary goods and 30% for other services—Australia’s main export categories (agricultural and mineral commodities, inbound tourism, and tertiary education). This assumption determines the terms-of-trade environment: under domestic-only adoption (0% shift), Australia faces favourable terms of trade as trading partners’ costs are unchanged; under global adoption, cost reductions are shared and Australia must compete on a shifting frontier. Table 5 reports aggregate outcomes across the three scenarios.




Figure 8: Modelled real wage changes due to GenAI (vertical axis) versus full-time weekly wages in 2025 (horizontal axis) across 97 three-digit ANZSCO occupations. Bubble sizes reflect 2019 employment. Wage data from ABS Employee Earnings and Hours, May 2025.

Table 5: Sensitivity of key aggregate outcomes to the export-demand shift assumption.

Measure	Component	Percentage point contribution to GDP			Share of total GDP gain (%)		
		0%	10%(Ref)	20%	0%	10%(Ref)	20%
GDP(E)	Consumption	9.9	10.8	11.7	36.2	36.2	36.0
	Investment structures	9.5	10.3	11.3	34.6	34.5	34.5
	Investment equipment	5.2	6.2	7.4	19.0	20.9	22.6
	Exports	10.2	11.4	12.9	37.2	38.3	39.5
	Imports	−7.4	−8.9	−10.6	−27.0	−29.9	−32.6
	Net exports	2.8	2.5	2.2	10.2	8.4	6.9
	Real GDP(E)	27.4	29.8	32.6	100.0	100.0	100.0
GDP(I)	Technical change	9.9	10.1	10.4	36.1	33.9	31.9
	Equipment	6.1	7.3	8.7	22.5	24.6	26.8
	Structures	11.3	12.3	13.5	41.4	41.4	41.4
	Real GDP(I)	27.3	29.8	32.6	100.0	100.0	100.0

The qualitative pattern is robust: augmentation dominates automation across industries and occupations, capital deepening accounts for most of the GDP gain, and the broad occupational adjustments are preserved across scenarios. The key difference lies in the investment–consumption split. Under domestic adoption (0% shift), Australia experiences a terms-of-trade improvement that raises consumption relative to GDP; under larger global shifts, this advantage diminishes. As a result, the gap between GDP and consumption gains documented in the reference case—driven by higher investment requirements—is smaller when adoption is not global, even though the GDP gain itself is reduced without the export stimulus.

5 Conclusion

This paper develops a task-based CGE framework to analyse the long-run economic effects of generative AI. The model explicitly represents competition between non-AI, AI-augmented, and AI-automated task modes within occupations, allowing adoption decisions at the task level to propagate through industry structure, capital deepening, and labour demand in general equilibrium. By embedding this mechanism in a 45-industry, 97-occupation model calibrated using separate augmentation and automation exposure scores, the framework provides a coherent account of how GenAI reshapes production, occupational employment, and relative wages.

Two main conclusions emerge. First, large aggregate gains from GenAI need not arise primarily from direct labour productivity improvements. In the reference case, real GDP increases by 29.8%, but around two thirds of this gain reflects capital deepening and general-equilibrium reallocation triggered by task-level adoption rather than direct technical change (10.1 pp). Welfare rises by a materially smaller 16.2%, because the more capital-intensive structure of the GenAI-equipped economy requires a permanently higher investment share—rising from 34.5% to 40.7% of GDP. The welfare–GDP divergence, together with the associated shift in factor income toward capital (the labour share falls from 60.0% to 52.4%), suggests that the distributional consequences of GenAI may be as significant as its aggregate productivity effects.

Second, labour-market adjustment occurs primarily within occupations rather than between them. Extensive-margin task reallocation affects roughly two thirds of all work in the economy, even though net changes in occupational employment and wages remain moderate. In the reference case, 46.7% of tasks are performed in augmented mode and 9.1% in automated mode, with the remainder non-AI; most work is therefore reorganised around AI-assisted production rather than eliminated. Where net employment shifts do occur, they favour blue-collar occupations and disadvantage lower-skilled white-collar roles—in particular clerical and administrative workers (group 5), which combine high automation exposure with a large initial employment share—but these movements are modest relative to the scale of within-occupation restructuring and to historical patterns of structural change. The primary adjustment burden thus falls on workers adapting to new task compositions within their occupation; the nature of work is changed at least as much as its quantity.

These findings should be interpreted with care. The model is comparative static and characterises differences between two long-run equilibria; it does not capture the transition path, diffusion lags, adjustment costs, or short-run labour-market disruption. Labour supply is fixed, and there is no explicit treatment of unemployment, hours worked, participation, or search frictions. The analysis also assumes a fixed global rate of return to capital. While this assumption is consistent with standard neoclassical modelling, global savings rates may not adjust sufficiently to sustain the increase in the investment share implied by the simulations, from 34.5% to 40.7%. Australia has a relatively small technology sector and will most likely rely heavily on foreign direct investment and imported AI hardware, with the returns on FDI flowing to foreign owners. What matters more, however, is the implicit assumption that, in the long run, the ratio of the net international investment position to GDP is unchanged—which amounts to assuming that Australians will increase their rate of saving commensurate with the increased capital intensity of the economy.¹¹

A further set of limitations concerns the exposure data on which the calibration rests. The JSA scores are derived from LLM-based assessments of task descriptions—a method whose accuracy and consistency have not been systematically validated against observed adoption outcomes—so the indices may mischaracterise the relative potential for augmentation and automation across occupations, or the underlying task descriptions may not adequately capture what workers actually do. More fundamentally, the analysis characterises GenAI entirely in terms of its exposure to existing occupational tasks, and so cannot account for the emergence of entirely new tasks or occupations (Acemoglu and Restrepo, 2018; Acemoglu, 2025), nor for the possibility that GenAI accelerates technological progress itself (Jones, 2025). Even modest increases in the rate of idea generation could compound over time and potentially dwarf the static reallocation effects examined here.

Several directions for future work follow naturally. Extending the framework to a dynamic setting would allow the costs of transition—not just the long-run equilibrium—to be quantified. Enriching the labour-market block to allow for unemployment, participation responses, and retraining would better capture the within-occupation adjustment frictions identified here. Modelling AI services supply more explicitly, including endogenous inference cost reductions, would sharpen the price feedbacks currently captured through a simplified composite.

Overall, the paper’s main message is that generative AI has the potential to produce large long-run gains in output, but that these gains are unlikely to arrive as a simple productivity dividend. They are mediated by investment, trade, relative prices, occupational reallocation, and the distinction between augmentation and automation. A useful contribution of the task-based CGE approach is precisely that it makes those channels explicit and quantifiable.

A Leontief decomposition

This appendix states the decomposition used to attribute industry value-added changes to demand-side drivers in Figure 5. A full derivation is given in Lennox and Dixon (2026, Appendix C).

Let $\mathbf{y}$ denote the vector of industry gross outputs, $\mathbf{F}_d$ the matrix of domestic final demand quantities (columns indexed by demand category h : consumption, exports, equipment investment, non-equipment investment), and $\mathbf{A}_d$ the matrix of domestic input–output coefficients. Market clearing gives $\mathbf{y} = \mathbf{A}_d \mathbf{y} + \mathbf{F}_d \mathbf{1}$, and the Leontief inverse $\mathbf{L} \equiv (\mathbf{I} - \mathbf{A}_d)^{-1}$ delivers $\mathbf{Y} = \mathbf{L} \mathbf{F}_d$, where element (i, h) is the output of industry i required to satisfy demand category h .

In the comparative-static solution, both $\mathbf{L}$ and $\mathbf{F}_d$ change between equilibria (the IO coefficients $\mathbf{A}_d$ are endogenous, reflecting import substitution and changing AI services intensity). Taking the total differential, $d\mathbf{Y} = d\mathbf{L} \mathbf{F}_d + \mathbf{L} d\mathbf{F}_d$. The first term captures supply-side effects (import substitution and input intensity changes); the second captures demand-side effects. Writing $\mathbf{F}_d = (\mathbf{S}_f \odot \mathbf{S}_c) \mathbf{D} \mathbf{Y}$ —where $\mathbf{S}_f$, $\mathbf{S}_c$, $\mathbf{D}$, and $\mathbf{Y}$ capture, respectively, final import shares, commodity composition within each demand category, category shares in GDP, and aggregate GDP—the demand-side differential yields four additive components: substitution against final imports, compositional shifts within demand aggregates, shifts in demand category shares, and aggregate scale. These

¹¹For context, Australia’s NIIP has been relatively stable at around –50% of GDP for three decades. There are nonetheless reasons to expect long-run saving to adjust. If AI adoption raises the capital share of income globally—as our results suggest it does domestically—standard neoclassical theory implies that the equilibrium saving rate rises endogenously: in the Ramsey model, the steady-state saving rate is increasing in the capital share, with no change in preferences required. Australia’s compulsory superannuation system provides a further automatic channel, since higher labour income translates mechanically into higher mandated saving.

four components are plotted in Figure 5.

References

- Acemoglu, D. (2025). The simple macroeconomics of AI. *Economic Policy*, 40(121):13–58.
- Acemoglu, D. and Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of labor economics*, volume 4, pages 1043–1171. Elsevier.
- Acemoglu, D., Kong, F., and Restrepo, P. (2025). Tasks at work: comparative advantage, technology and labor demand. *Handbook of Labor Economics*, 6:1–114.
- Acemoglu, D. and Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American economic review*, 108(6):1488–1542.
- Acemoglu, D. and Restrepo, P. (2022). Tasks, automation, and the rise in US wage inequality. *Econometrica*, 90(5):1973–2016.
- Aghion, P. and Bunel, S. (2024). AI and growth: Where do we stand? Technical report, Federal Reserve Bank of San Francisco. Working Paper.
- Autor, D. H., Levy, F., and Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly journal of economics*, 118(4):1279–1333.
- Baumol, W. J. (1967). Macroeconomics of unbalanced growth: the anatomy of urban crisis. *The American economic review*, 57(3):415–426.
- Bekkers, E., Humphreys, L., Kalachyhin, H., Wilczynska, K., and Zhao, D. (2025). Through the looking glass: Artificial intelligence, international trade, and economic growth in the long run. Technical report, WTO Staff Working Paper.
- Bresnahan, T. F. and Trajtenberg, M. (1995). General purpose technologies ‘engines of growth’? *Journal of econometrics*, 65(1):83–108.
- Brynjolfsson, E. (2023). The turing trap: The promise & peril of human-like artificial intelligence. In *Augmented education in the global age*, pages 103–116. Routledge.
- Brynjolfsson, E., Chandar, B., and Chen, R. (2025). Canaries in the coal mine? six facts about the recent employment effects of artificial intelligence. *Stanford Digital Economy Lab*. Published August.
- Brynjolfsson, E., Rock, D., and Syverson, C. (2021). The productivity j-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics*, 13(1):333–372.
- Crafts, N. (2021). Artificial intelligence as a general-purpose technology: an historical perspective. *Oxford Review of Economic Policy*, 37(3):521–536.
- Cubas, G., Silos, P., and Soini, V. (2024). Risk and the allocation of talent in the Roy model. *Economics Letters*, 236:111623.
- Cui, K. Z., Demirer, M., Jaffe, S., Musolff, L., Peng, S., and Salz, T. (2026). The effects of generative ai on high-skilled work: Evidence from three field experiments with software developers. *Management Science*.
- Dixon, J. (2022). *Forecasting labour shortages in the Australian Economy*. Victoria University.
- Dixon, P. B., Parmenter, B. R., Sutton, J., and Vincent, D. P. (1982). *ORANI: A Multisectoral Model of the Australian Economy*. North-Holland, Amsterdam.
- Eloundou, T., Manning, S., Mishkin, P., and Rock, D. (2024). Gpts are gpts: Labor market impact potential of llms. *Science*, 384(6702):1306–1308.

- Filippucci, F., Gal, P., Laengle, K., and Schief, M. (2025). Macroeconomic productivity gains from artificial intelligence in g7 economies. *OECD Artificial Intelligence Papers*.
- Filippucci, F., Gal, P., and Schief, M. (2024). Miracle or myth? assessing the macroeconomic productivity gains from artificial intelligence. Technical report, OECD Publishing.
- Gmyrek, P., Berg, J., and Bescond, D. (2023). Generative AI and jobs: A global analysis of potential effects on job quantity and quality. *ILO Working paper*, 96.
- Gundlach, H., Lynch, J., Mertens, M., and Thompson, N. (2025). The price of progress: Algorithmic efficiency and the falling cost of ai inference. *arXiv preprint arXiv:2511.23455*.
- Horridge, M. (2000). Orani-g: A general equilibrium model of the australian economy. Technical Report OP-93, Centre of Policy Studies and the Impact Project, Monash University, Melbourne.
- Humlum, A. and Vestergaard, E. (2025). Large language models, small labor market effects. Technical report.
- Jobs and Skills Australia (2025a). Our Gen AI transition: Exposures, adaptation, dynamism. Website. Available online; accessed 17 December 2025.
- Jobs and Skills Australia (2025b). Our Gen AI transition: Implications for work and skills. Final Overarching Report. Final release, 30 September 2025.
- Jones, B. (2025). Artificial intelligence in research and development. Technical report, National Bureau of Economic Research.
- Lennox, J. A. and Dixon, J. M. (2026). Occupations, tasks and generative AI: A computable general equilibrium analysis. Working paper, Centre of Policy Studies, Victoria University, Melbourne, Australia.
- Liu, Y., Wang, H., and Yu, S. (2025). Labor demand in the age of generative ai: Early evidence from the us job posting data. *Available at SSRN 5504741*.
- Noy, S. and Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654):187–192.
- OECD (2025). Ways to minimise water use related to AI operations. Policy and technical options for reducing water consumption in AI data centers.
- Wang, F., Zhang, Z., Zhang, X., Wu, Z., Mo, T., Lu, Q., Wang, W., Li, R., Xu, J., Tang, X., et al. (2024). A comprehensive survey of small language models in the era of large language models: Techniques, enhancements, applications, collaboration with llms, and trustworthiness. *ACM Transactions on Intelligent Systems and Technology*.