Australian Government | Jobs and Skills Australia

Our Gen AI Transition

Implications for Work and Skills

Final Release

30 September 2025

Acknowledgement of Country

Jobs and Skills Australia acknowledges the Traditional Owners of Country throughout Australia and recognises the continuing connection to lands, waters and communities. We pay our respect to Aboriginal and Torres Strait Islander cultures, and to Elders past and present.

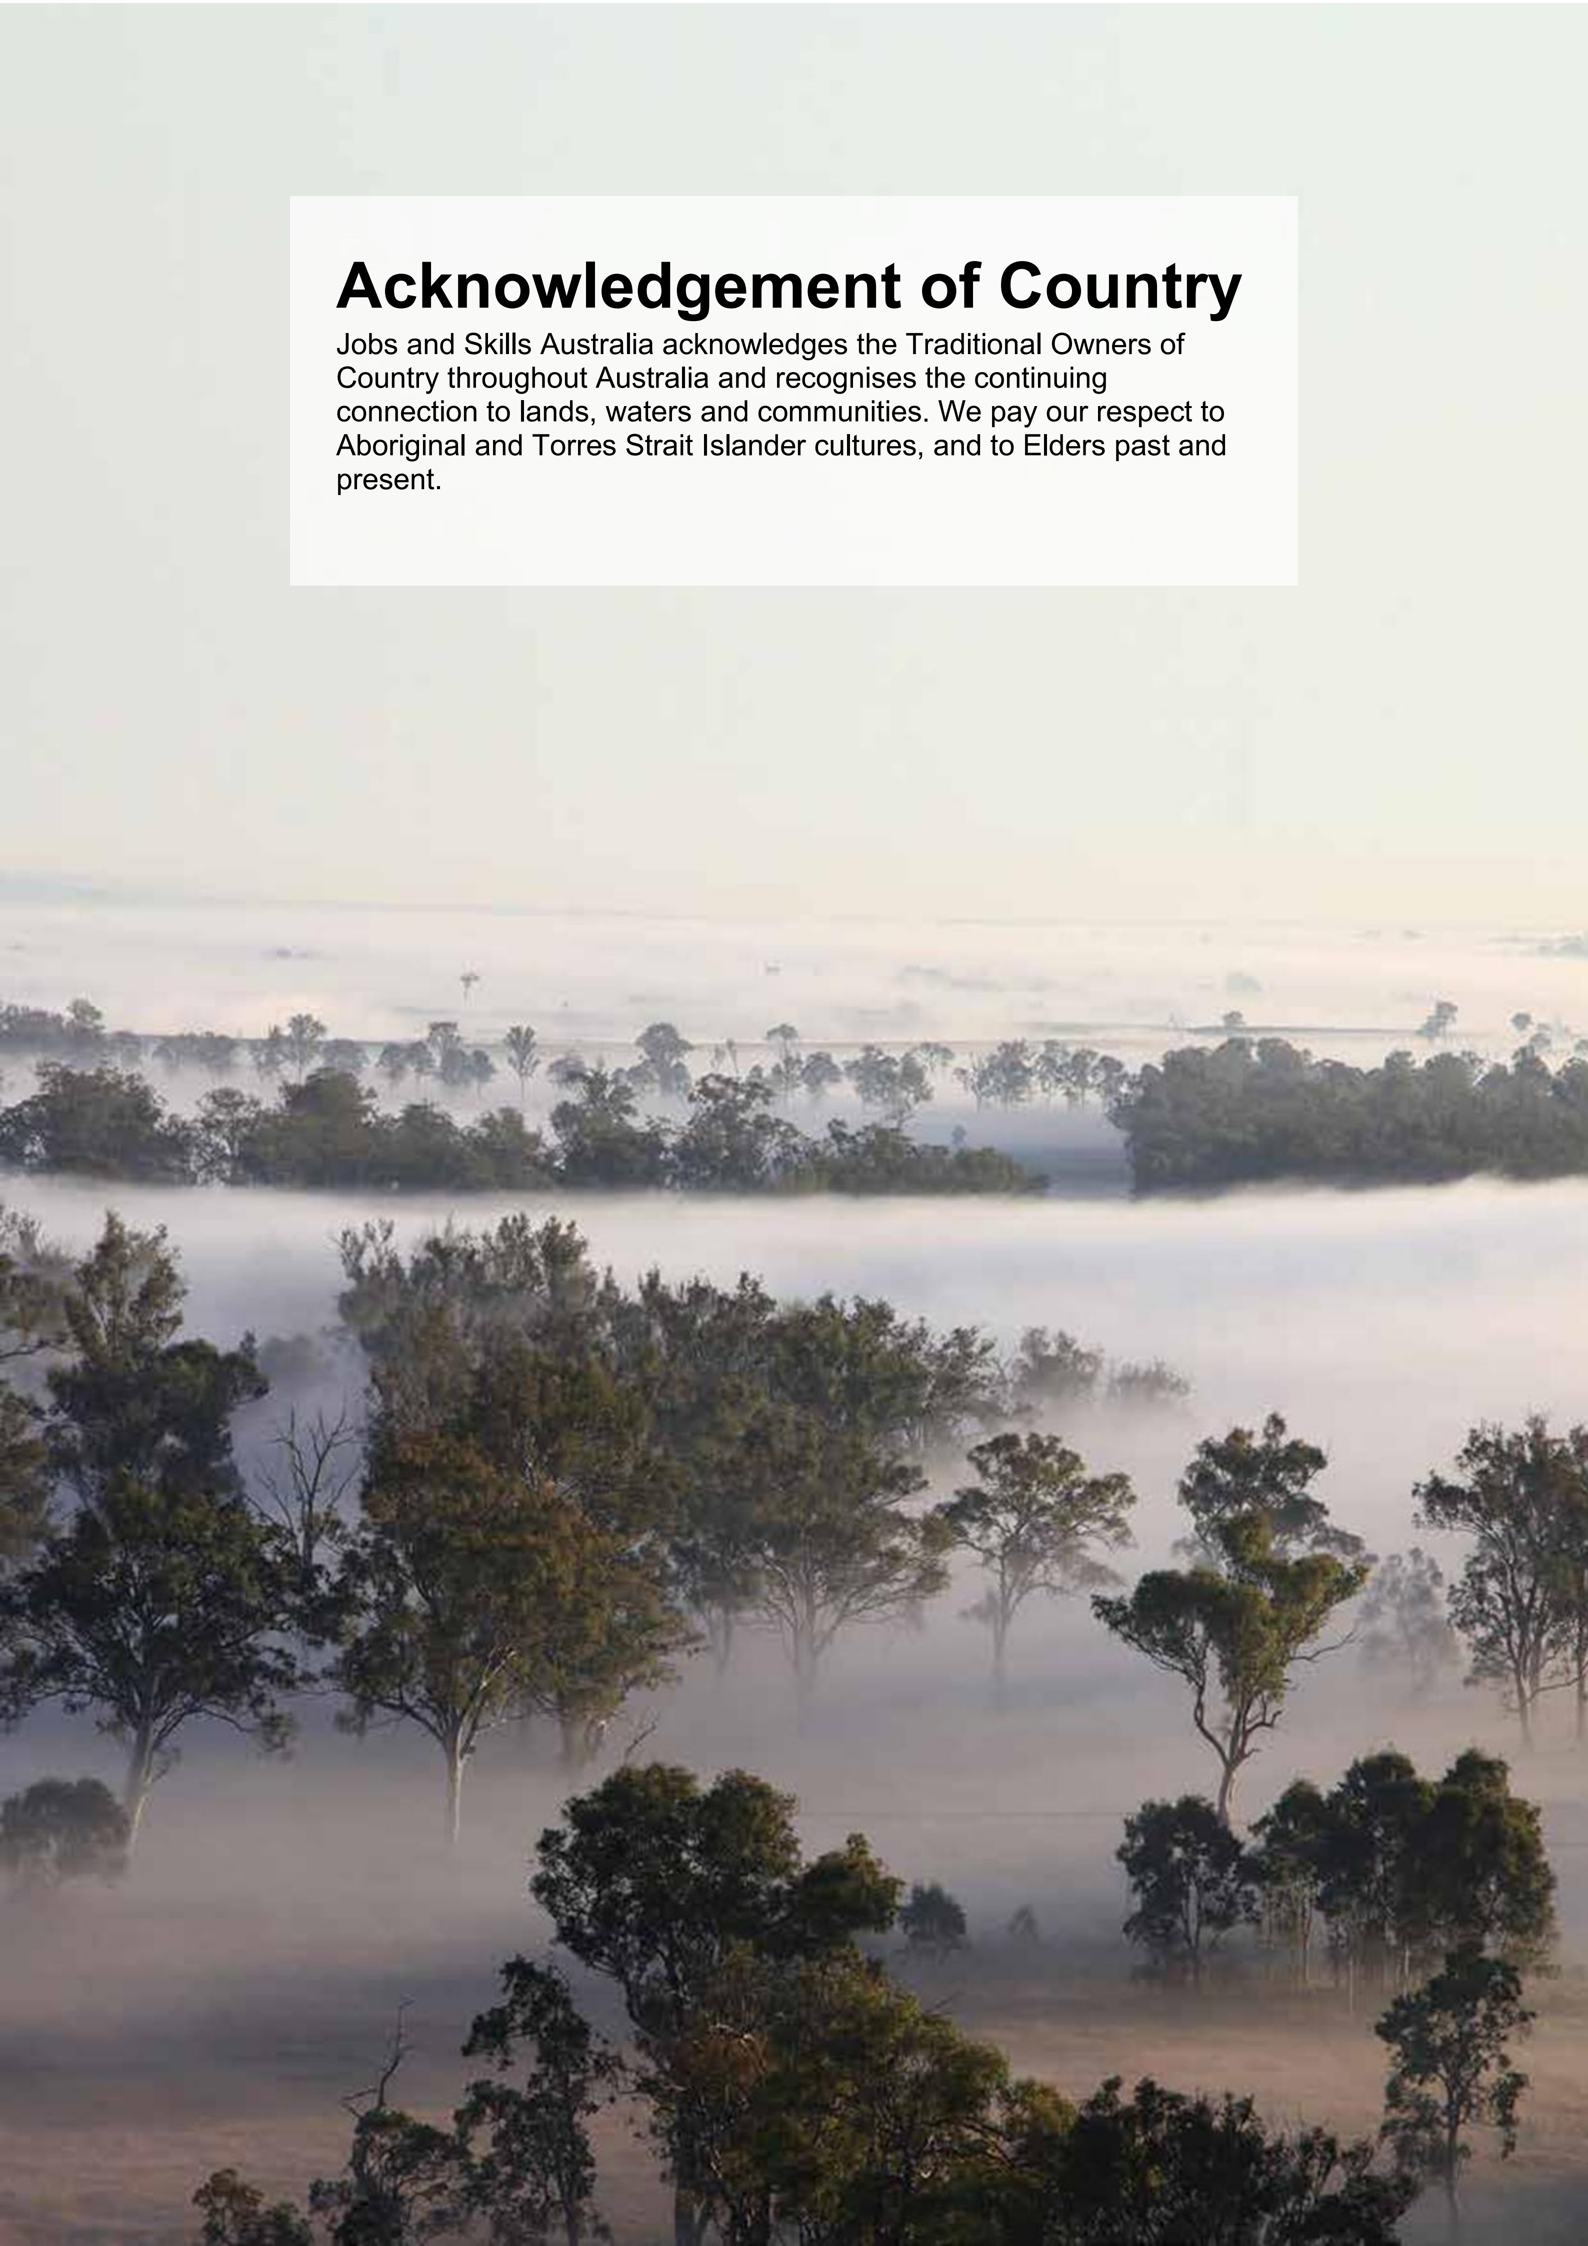

Contents

JSA acknowledges and thanks Commissioner Barney Glover, Executive Director David Turvey and Bjorn Jarvis for overseeing this study; Ishita Acharyya and Hudan Nuch for leading the project team; and team members Mohammad Seyani, Rachel Neumann, Paul O'Shea, Denis Leonov; Andrew Kruup; Kevin Matthew; Mohd Mohsin; Blane Simpson; Simon Madden, Ishani Sahama, Andrew Maurer, Darcy Bermingham and Brad Waters for their contributions. JSA acknowledges the use of AI tools in undertaking research and analysis for this paper.

An appropriate reference for this publication is:
Jobs and Skills Australia 2025, *Our Gen AI Transition: Implications for Work and Skills, Final Release*, Canberra

Foreword

The final release of *Our Gen AI Transition*, on 30 September 2025, has three *parts.*

The first part contains an additional analysis paper (**Analysis Paper F – Workers' exposure & experiences**) to be read in conjunction with the five papers released in ***Our Gen AI Transition: Analysis Papers*** on 2 September 2025. This final analysis paper explores how the implications of Gen AI vary across the workforce, including exploring exposure, experiences and perspectives, through cohort and intersectional lenses.

The second part (**Technical Notes**) provides the technical detail of how JSA undertook the key pieces of analysis in the overarching report and analysis papers. The technical notes also provide guidance on interpreting and using the analytical outputs from the Study.

Lastly, as part of this 30 September release, we have also updated ***Our Gen AI Transition: Case Studies*** to include two additional case studies to the eight released on 2 September 2025. The two additional case studies consider Gen AI experiences and complexities in non-market settings and in leveraging co-design.

Data – available online

The final release includes a full Chart and Data Pack for the reports.

The Study's full published data suite provides a range of additional outputs to support analysis. These include bespoke tables, interactive tables and highcharts.

- **Bespoke tables** present 17+ tables with detailed analytical results for automation, augmentation, adaptation, mobility, skill change, and entry-level analysis. Tables include results on cohorts, including age, gender, disability, employment type, region, industry, and skill level. The tables compare across occupations and sectors.
- **Interactive tables and charts (highcharts)** allow users to filter and view data by any of the Study's topics and occupation and industry. These tables and charts include indicators for AI exposure, adaptation, mobility, and entry-level job shares. The interactive format is designed to support further exploration of the results.

Some underpinning data used in this study has not been published (e.g. due to agreed undertakings with data custodians).

F: Workers' exposure and experiences

Implications of exposure and adoption for different workers

Adaptability across the workforce

How sections of the workforce experience technology

1. Introduction

This paper considers how the implications of Gen AI vary across the workforce. To the extent that cohorts and sections of the workforce experience different impacts of Gen AI, the transition to a Gen AI-enabled economy could help to resolve – or exacerbate – existing inequities in the labour market.

The analysis in this study shows that exposure and adoption are currently uneven across occupations and industries. This in turn means different impacts for different groups in the workforce, given the current composition of the workforce (Section 2).

The adoption and use of Gen AI is not necessarily a positive or negative development for workers, and its impact depends on how and where it is used. A key determinant of whether individual workers gain from the use of Gen AI is their ability to adapt to changes in their occupations and workplaces, including the capability to use digital and AI tools. At a foundational level, this includes baseline capabilities and access to technologies – both of which vary across the workforce (Section 3).

Workers may also have different experiences of technology based on their personal characteristics and circumstances (Section 4). People may be affected differently within the workplace, occupation or industry. This can arise due to biases in the design of AI technologies and their applications, or the ways in which they are used.

While it is critical that each of the issues above are considered through an intersectional lens, the availability of appropriate data – beyond what is presented in Section 2 – is a limiting factor. For example, online job advertisements data show which skills and occupations are in demand, but do not provide an indication of the personal characteristics of workers hired. Online worker profiles data also tends to lack useful demographic markers.

As Australia progresses further in its AI transition, more intersectional data will become available, including in official statistics and the large integrated datasets that JSA uses heavily in its research. This will need to continue to be supplemented by qualitative research, like the case studies of this Study, to also understand the different experiences of different groups of people.

It is critical that there is an increasing intersectional focus in quantitative and qualitative analysis, to better understand the implications of Gen AI for different groups of people in the labour market.

2. Implications of exposure and adoption for different workers

Exposure worker cohorts

The effects of Gen AI are likely to vary across different groups of workers. In this section, we discuss three priority cohorts in depth: Females, First Nation People, and People with Disability, and identify the opportunities and challenges associated with Gen AI adoption.

Box 1 Women have higher exposure to augmentation and automation

Based on the current gender composition of occupations and their respective exposure to Gen AI, women face greater exposure to both automation and augmentation (Figure 1). Women are overrepresented in highly automatable roles (Figure 2). However, there are many occupations that are growing and have strong potential for augmentation, more than automation, including a range of female-dominated occupations.

Figure 1: On average, women face greater augmentation and automation

Weighted average automation and augmentation scores by gender

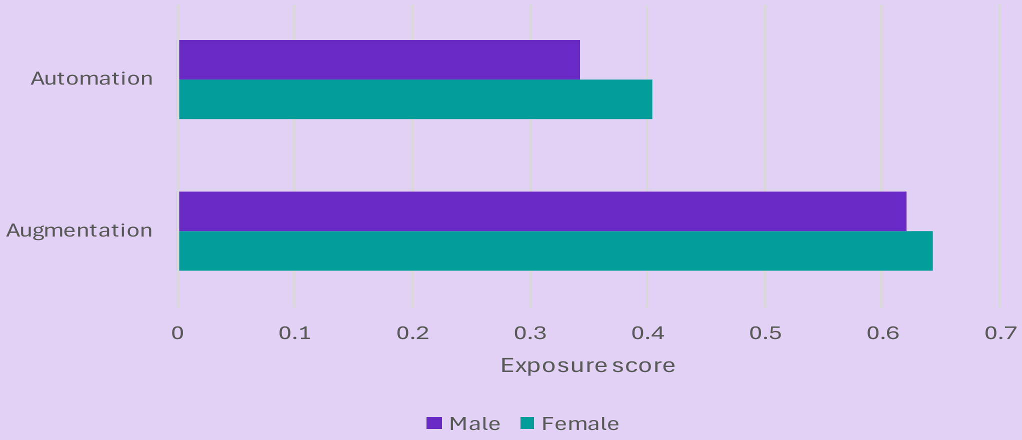

Source: JSA Analysis; ANZSCO (v1.3), Census 2021 (TableBuilder).

Figure 2: Augmentability is similar and automatability is higher for the female workforce

Female and male workforce exposure to automation and augmentation in current occupations

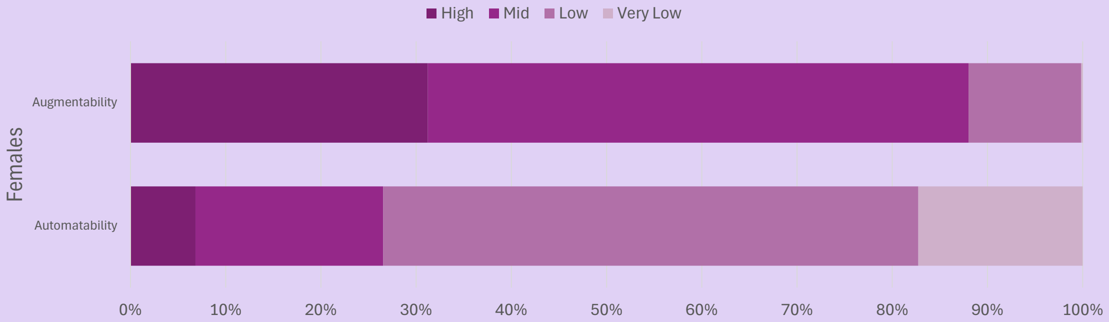

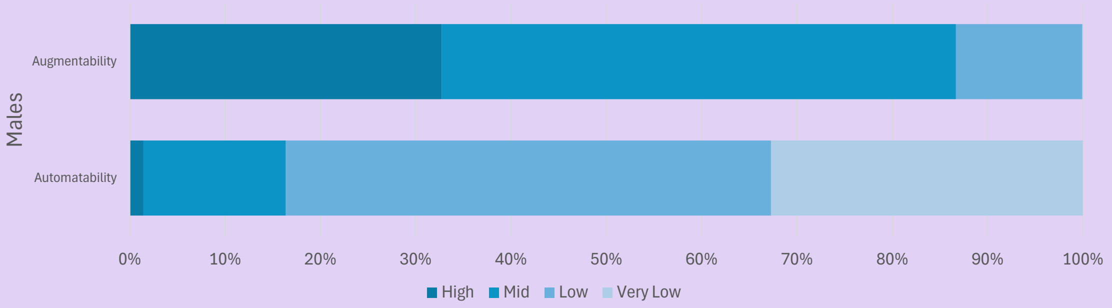

Source: JSA analysis of ANZSCO (v1.3), Census 2021 (Tablebuilder)

This aligns with some previous findings in Australian research. In the financial services sector, for example, female employment is high in occupations that are projected to see low growth, including Accounting Clerks, Bookkeepers, Payroll Clerks, and Bank workers, all of which face significant exposure to Gen AI (FSO, 2024). The FSO has also pointed to occupations in other sectors, such as Medical Administrative Assistants as well as Legal Assistants as being particularly exposed (Mandala, 2024).

This also aligns with ILO findings that occupations with the highest automation exposure employ more women than men (even if automation exposure is generally low overall). Moreover, critical occupations and industries that are female-dominated could need more investment in Gen AI technologies and skills to support already stretched workforces, with opportunities from these technologies (Gmyrek, Berg, & Bescond, 2023). Faster and more informed adoption of Gen AI in the female-dominated care economies, for example, could improve job quality and productivity in these sectors, and support them to meet increasing demand and existing unmet demand.

Looking across the most common occupations for different groups of people in the labour market, we can see that they include a range of automation and augmentation exposures, including for women (Table 1), young workers (Table 2), First Nations workers (Table 3) and workers living in remote and very remote areas in Australia (Table 4).

Table 1: Exposure varies across the most common occupations for women

Gen AI exposure for most common occupations for women; share of female workers (%)

Occupation	Augmentation Level	Automation Level	Percentage of total cohort workforce
Sales Assistants (General)	Medium	Low	5.9
Registered Nurses	Medium	Low	4.0
General Clerks	Medium	High	3.7
Aged and Disabled Carers	Low	Very Low	3.0
Child Carers	Medium	Very Low	2.7
Receptionists	Medium	Medium	2.6
Primary School Teachers	Medium	Low	2.4
Office Managers	High	Low	1.7
Retail Managers	High	Low	1.7
Secondary School Teachers	Medium	Low	1.7

Source: JSA analysis of PLIDA data, 2025 (Skills Tracker project), JSA Data on Occupation Mobility (DOM) 2022-23

Table 2: Exposure varies across the most common occupations for young people

Gen AI exposure for most common occupations for younger workers (15-24 years); share of younger workers (15-24 years) (%)

Occupation	Augmentation Level	Automation Level	Percentage of total cohort workforce
Sales Assistants (General)	Medium	Low	14.8
Kitchenhands	Medium	Very Low	3.5
Checkout Operators and Office Cashiers	Medium	Medium	3.3
Waiters	Low	Low	3.3
Bar Attendants and Baristas	Low	Very Low	2.8
Child Carers	Medium	Very Low	2.4
Fast Food Cooks	Medium	Low	2.2
Receptionists	Medium	Medium	2.1
Carpenters and Joiners	Medium	Very Low	1.9
Storepersons	Medium	Low	1.8

Source: JSA analysis of PLIDA data, 2025 (Skills Tracker project), JSA Data on Occupation Mobility (DOM) 2022-23

Table 3: Exposure varies across the most common occupations for First Nations people

Gen AI exposure for most common occupations for First Nations workers; share of First Nations workers (%)

Occupation	Augmentation Level	Automation Level	Percentage of total cohort workforce
Sales Assistants (General)	Medium	Low	5.4
Aged and Disabled Carers	Low	Very Low	2.8
General Clerks	Medium	High	2.5
Education Aides	Medium	Low	2.3
Truck Drivers	Medium	Very Low	2.2
Welfare Support Workers	Medium	Low	2.2
Child Carers	Medium	Very Low	1.9
Commercial Cleaners	Low	Very Low	1.8
Receptionists	Medium	Medium	1.6
Drillers, Miners and Shot Firers	Medium	Very Low	1.5

Source: JSA analysis of PLIDA data, 2025 (Skills Tracker project), JSA Data on Occupation Mobility (DOM) 2022-23

Table 4: Exposure varies across the most common occupations in remote and very remote areas

Gen AI exposure for most common occupations in remote areas; share of workers in remote areas (%)

Occupation	Augmentation Level	Automation Level	Percentage of total cohort workforce
Livestock Farmers	Medium	Low	4.2
Drillers, Miners and Shot Firers	Medium	Very Low	3.1
Sales Assistants (General)	Medium	Low	2.9
Mixed Crop and Livestock Farmers	Medium	Low	2.4
Metal Fitters and Machinists	Medium	Very Low	2.3
General Clerks	Medium	High	2.3
Truck Drivers	Medium	Very Low	2.1
Livestock Farm Workers	Low	Very Low	2.0
Registered Nurses	Medium	Low	1.9
Crop Farmers	Medium	Low	1.9

Source: JSA analysis of PLIDA data, 2025 (Skills Tracker project), JSA Data on Occupation Mobility (DOM) 2022-23

Exposure and mobility patterns among worker cohorts

To assess the impact of Gen AI on various cohorts through augmentation exposure, it is important to consider their occupational mobility and resilience. Analysis was undertaken using the Data on Occupation Mobility (DOM) 2022-23 dataset (extracted from ABS PLIDA data (Skills Tracker project).

Occupations were grouped into the following categories to capture their mobility and augmentation exposure level: 'Lower Augmentation Potential & Higher Mobility'; 'Lower Augmentation Potential & Lower Mobility'; 'Higher Augmentation Potential & Lower Mobility'; and 'Higher Augmentation Potential & Higher Mobility'. Table 5 below provides a summary of each of the categories.

Table 5: Occupation mobility and AI augmentation

Classification of occupation	Intuition
Higher Augmentation Potential & Higher Mobility	Higher augmentation scores and higher mobility indicate potential to use Gen AI in the current occupation, and with heightened potential to transition to other occupations with potential opportunities for Gen AI use.
Higher Augmentation Potential & Lower Mobility	Higher augmentation scores and lower mobility indicate potential to use Gen AI in the current occupation, but with fewer or weaker exit paths than the first quadrant.
Lower Augmentation Potential & Higher Mobility	Lower augmentation scores and higher mobility indicate relatively high occupational resilience, with more or stronger exit paths than the final quadrant. This quadrant may include occupations with potential opportunities for Gen AI use.
Lower Augmentation Potential & Lower Mobility	Lower augmentation scores and lower mobility indicate relatively high insulation from the effects of Gen AI.

Source: JSA analysis of PLIDA data, 2025 (Skills Tracker project), JSA Data on Occupation Mobility (DOM) 2022-23

Table 6 below outlines the shares of workers within each category by cohort. Almost 80% of workers across cohorts were in either the 'Lower Augmentation Potential & Higher Mobility' or 'Higher Augmentation Potential & Higher Mobility' categories. Furthermore, only a small share of workers across each cohort was in the 'Higher Augmentation Potential & Lower Mobility category'.

The analysis showed that over 60% of women were in occupations exposed to high augmentation ('Higher Augmentation Potential & Higher Mobility' and 'Higher Augmentation Potential & Lower Mobility'). More specifically, they had the highest share of workers in the 'Higher Augmentation Potential & Higher Mobility' category due to large numbers of women working as Registered Nurses, Primary School Teachers, Office Managers and Retail Managers. The relatively high share of workers in the 'Higher Augmentation Potential & Lower Mobility' category reflected relatively high numbers of women working as Information Officers, Keyboard Operators, General Practitioners and Resident Medical Officers.

Young workers were more likely to be in the 'Lower Augmentation Potential & Higher Mobility' category, given their high number of Sales Assistants (General), Kitchenhands, Checkout Operators and Office Cashiers and Waiters. 'Higher Augmentation Potential & Higher Mobility' was the second highest share, given the number of young people working in Receptionist jobs.

A relatively high share of First Nations workers were in the 'Lower Augmentation Potential & Higher Mobility' category, given the number of workers employed as Sales Assistants

(General), Aged and Disabled Carers, Truck Drivers and Child Carers. First Nations workers also had a reasonable percentage of people in the 'Higher Augmentation Potential & Higher Mobility' category, reflecting Receptionists, General Clerks, Education Aides and Welfare Support Workers.

Finally, in relation to remote and very remote workers, the high share of workers in the 'Lower Augmentation Potential & Higher Mobility' category can be explained by large numbers employed as Livestock Farmers, Drillers, Miners and Shot Firers and Metal Fitters and Machinists.

Table 6: Share of workers (%) by Cohort and Mobility/AI augmentation category

Cohort	Lower Augmentation Potential & Higher Mobility	Lower Augmentation Potential & Lower Mobility	Higher Augmentation Potential & Lower Mobility	Higher Augmentation Potential & Higher Mobility
Demographics				
Young (15-24)	67	6	4	23
Women	34	4	11	51
First Nations	50	9	7	35
CALD	40	8	9	43
Men	43	11	8	38
Geography				
Major Cities	36	7	10	47
Inner Regional	46	9	8	37
Outer Regional	48	10	9	33
Remote	46	11	10	33
Very Remote	49	10	8	34

Source: JSA analysis of PLIDA data, 2025 (Skills Tracker project), JSA Data on Occupation Mobility (DOM) 2022-23

3. Digital adaptability across the workforce

The previous section discussed how the labour market impacts – both potential opportunities and challenges – could vary for different groups of people. In addition to considering the occupation people are in, it is important to consider digital literacy across Australia, given it has the potential to also influence the impacts that might be associated with Gen AI adoption. In this regard, the 'digital divide' is central to this discussion.

Adoption experiences also appear to vary according to age, highlighting the importance of considering age in intersectional analysis. Fusion Digital Agency (2024) found that white-collar workers aged 58 years or more were the least likely to use AI, with only 75% identifying as AI users. In contrast, those aged 27-42 years were the most likely to adopt AI tools at work, with 90% of respondents in this age group indicating they used AI.

Cohort specific digital inclusion

The ‘digital divide’ refers to unequal access and use of digital technologies (Min, 2011). At the centre of discussions around the ‘digital divide’ is the Australian Digital Inclusion Index (ADII) 2023[1] which provides a detailed measure of digital inclusion in Australia, that identifies critical barriers to inclusion. The publication uses survey data to measure digital inclusion across Access, Affordability and Digital Ability.

The latest ADII report revealed that, while digital inclusion at the national level has improved, 9.4% of the population remained highly digitally excluded (Thomas, et al., 2023). There remains a considerable digital gap (7.5 points) between non-First Nations people and First Nations people. The gap widens to 21.6 points between remote First Nations people and non-First Nations people who reside in the same areas, with accessibility being a critical issue (see Case Study 4 for more discussion on this issue).

Moreover, while the Affordability Gap between capital cities and other areas of the country continued to narrow, the ‘Digital Ability’ gap increased from 7.0 points to 7.7 (between 2021 and 2023). Older Australians were more digitally excluded, particularly those aged over 75 where the disparity was considerable (41.6 points below the national average for Digital Ability and 18 points below the national average for Access).

Figure 3 provides an overview of cohorts overrepresented among the ‘digitally excluded and highly excluded’. People living in remote/very remote areas, First Nations people, and people with a disability, were well above the national average. Moreover, the extent of digital inclusion was strongly correlated with other indicators of advantage such as income, education and employment. The ADII showed that higher digital inclusion was associated with higher levels of education and income levels (Figure 4).

---

[1] The Australian Digital Inclusion Index is a relative measure of inclusion. Using a score of 0 to 100, it compares the degree to which individuals can be considered more or less digitally included than others based on three dimensions: Access, Affordability and Digital Ability.

Figure 3: People with disability, Very Remote Australians and First Nation People are more digitally excluded

Proportion of digital exclusion among different cohorts

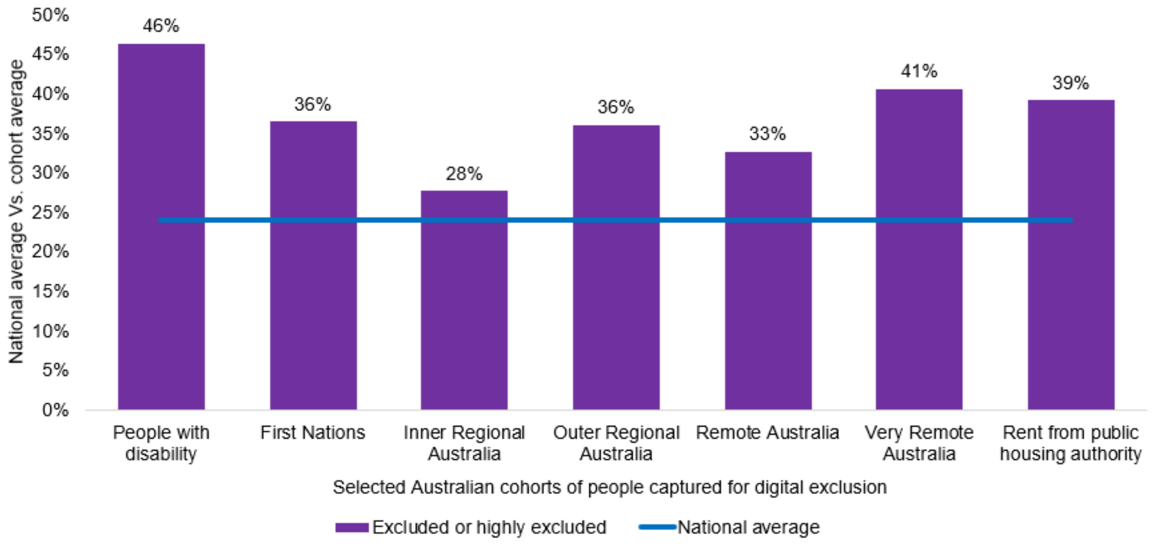

Source: Thomas et al (2023)

Figure 4: Australians with higher education and higher income are more digitally included

Digital inclusion scores of Australian adults by selected characteristics

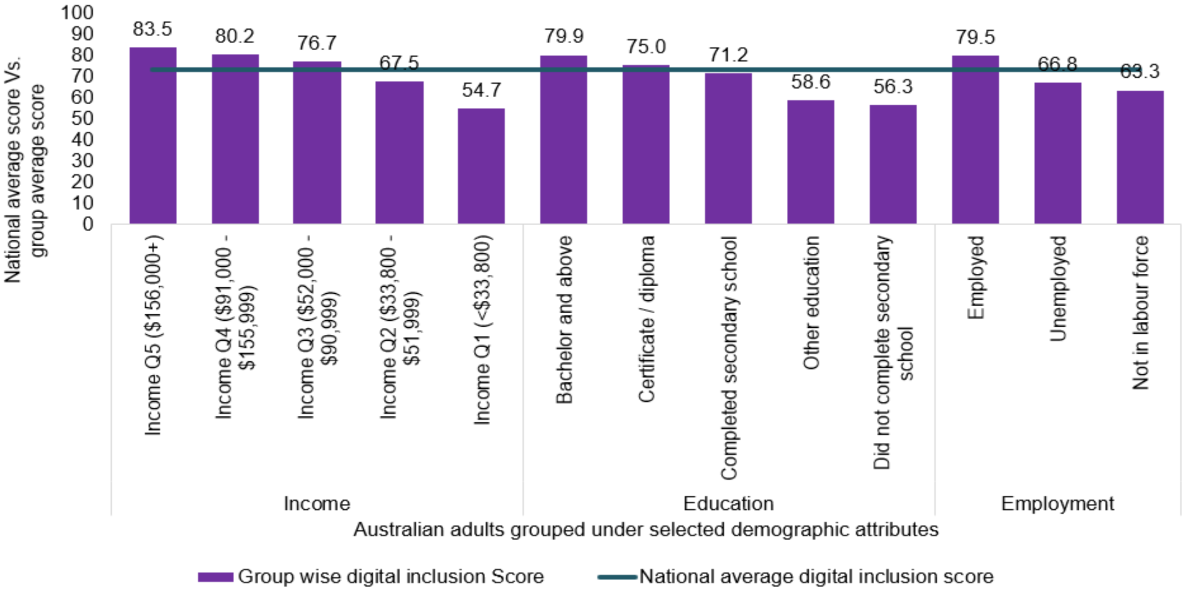

Source: Thomas et al (2023)

Digital exclusion in an increasingly digital society can have a major impact on a person's life across a wide range of domains including work, education, social interaction, commerce, and health. While each of these is important, work and education are the most important for our Study, in considering how different levels of digital inclusion could intersect with digital literacy, and how people can navigate the opportunities and challenges of the AI transition.

A foundational level of digital engagement is required across the vast majority of occupations in the Australian labour market. According to the Australian Skills Classification developed by the National Skills Commission (2022) (Figure 5) around 17% of occupations require basic digital skills at proficiency level 2 (commensurate with sending a short email reply) or level 3 (commensurate with entering information into a database).[2]

Figure 5: Over 80% of Australian occupations require intermediate digital proficiency skills

Distribution (%) of occupations by digital engagement proficiency level

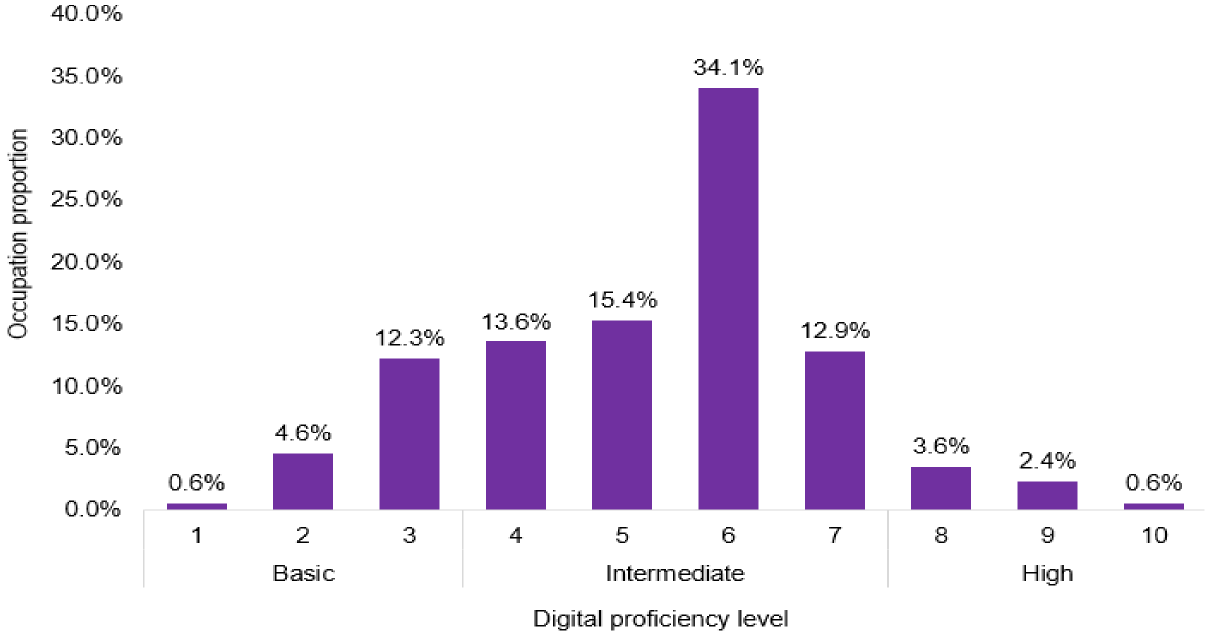

Source: National Skills Commission (2022), Australian Skills Classification

Digital exclusion not only limits the employment and earning opportunities available to people now but also impacts resilience and adaptability to future change as the transition to an increasingly digital economy continues, including the AI transition.

Digital exclusion also impacts the job search and application process. Research indicates that higher digital resources and skills facilitate the process of online job-seeking, increase the likelihood of being offered job interviews and reduce job-seeking burnout (De Marco, et al., 2023). This is significant given advertising on internet job boards is the most common method of recruitment in Australia, used by 63% of recruiting employers. In contrast, word of mouth recruitment – the most common offline method of recruitment – is used by only 30% of recruiting employers (JSA, 2023). The increase in AI-assisted recruitment within digital channels therefore makes digital inclusion and digital literacy even more important to consider across the workforce.

---

[2] The Australian Skills Classification (ASC) was the first attempt by the Australian Government to develop a common language to understand the skills and tasks within the Australian workplace. The product has not been updated since December 2023 and results should be interpreted with caution.

Box 2 Understanding the levels of digital divide

Digital divide research often focuses on the sequential relationship between the three levels of the digital divide which are material access, skills and use, and outcomes.[3] The sequential aspect is important as it suggest that one type of digital exclusion is dependent on another (Helsper, J.A.M. van , & Eynon, 2015; Van Deursen, Helsper, Eynon, & van Dijk, 2017).

When a person lacks digital skills, they are unable to use the Internet in a variety of ways which subsequently leads to being unable to achieve outcomes. (Van Deursen, Helsper, Eynon, & van Dijk, 2017)

As the Australian Network for Quality Digital Education observes:

Only if digital access is paired with digital literacy, and those skills put to productive use is digital inclusion achieved. (Loble & Stephens, 2024).

Fundamental to these three levels of the digital divide is the motivation of potential users and their ability to adopt, obtain, learn, use and leverage digital technologies (van Dijk, 2005). Figure 6 summarises the sequential and cumulative nature of the process which produces the digital divide. Importantly, the digital divide can be reproduced for each significant new digital innovation (including Gen AI) as experiences with previous technologies contribute to different starting lines for adoption of the next.

Figure 6: Framework for understanding the digital divide

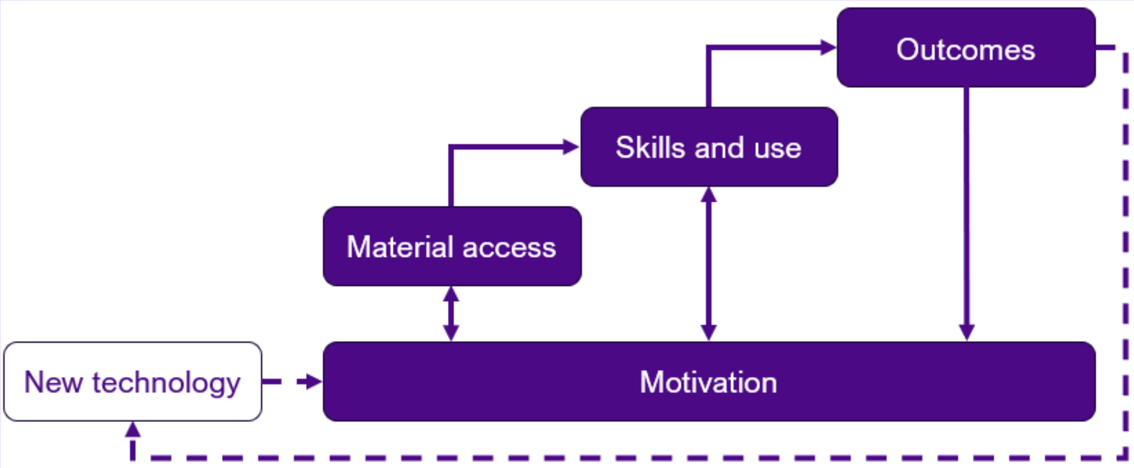

Source: Adapted from The Deepening Divide: Inequality in the Information Society (van Dijk, 2005)

The 2024 Australian Senate Report on Adopting Artificial Intelligence (The Senate, 2024) outlined a series of barriers regarding AI adoption in Australia. Amongst the more prevalent were insufficient AI-related skills.

The Productivity Commission (2024) also identified a lack of AI skills as being a fundamental factor to AI adoption in Australia.

Importantly, skills gaps – both among workers and managers – have been identified as one of the main barriers to AI adoption in workplaces across the globe (WEF, 2025).

Effective use of AI will require workers to at least become proficient in digital and foundational skills. The Australian Digital Inclusion Index, 2021 outlined a list of requisite digital ability skills (Table 7, Column 1). However, human (soft) skills are also extremely important in relation to AI adoption as evidenced by a global survey commissioned by Workday (Table 7, Column 2). See Analysis Paper E for more discussion of skills.

To this effect, the skills and knowledge that workers possess will be a key enabler of Gen AI adoption. The entire skills system will have a fundamental role to play in the uplift of skills across the entire population. Access to short form training will also be critical in the AI capability uplift.

Table 7: Skills required for AI adoption

Digital ability skills	Human skills
Basic operational: including downloading and opening files, connecting to the internet, and setting passwords	Adaptability and resilience: Ability to adjust to changes, Cultural sensitivity and adaptation, and Emotional intelligence and empathy
Advanced operational: including saving to the cloud, determining what is safe to download, customising devices and connections, and adjusting privacy settings	Analytical thinking and decision-making: Creative thinking and innovation; Critical thinking and problem-solving; Information synthesis and analysis Strategic planning
Information navigation: including searching and navigating, verifying trustworthy information, and managing third party data collection	Interpersonal connection and collaboration: Collaboration, Communication, Conflict resolution, Human networking and relationship building, and Negotiation and persuasion
Social: including deciding what to share, how, and who with, manage and monitor contacts, and communicate with others	Leadership and guidance: Ethical decision-making and moral judgment; Leadership and management; and Mentorship
Creative: including editing, producing, and posting content, as well as a broad understanding of the rules that may apply to these activities	
Automation: including connecting, operating, and managing smart devices and IoT technologies	

Source: Australian Digital Inclusion Index (2021); Workday (2025) Elevating Human Potential: The AI Skills Revolution

The AI divide

While the above discussion focused on the ‘digital divide’, this section looks more specifically at Gen AI exclusion.

The term ‘AI divide’ is broadly defined as a sub dimension of the digital divide and is based on the three levels of the digital divide identified by Carter, Liu and Cantrell (2020):

- Access to AI (the first level divide)

---

[3] Outcomes are conceptualised into four tangible outcome fields including Economic, Personal, Cultural and Social.

- Ability to use AI (the second-level divide), and
- Outcomes of AI engagement (the third level divide).

In addition, the authors highlighted the important role of individuals' perceptions, beliefs, and AI attitudes in relation to the AI divide (Carter, Liu, & Cantrell, 2020).

In relation to vulnerable cohorts, Wang, Boerman, Kroon, Möller, and de Vreese (2024) undertook research on the AI-shaped online news and entertainment environment focusing on the AI knowledge, skills and attitudes of users. The authors identified five user groups varying in AI knowledge, skills, and attitudes which included:

- the 'average users'
- the 'expert advocates'
- the 'expert sceptics'
- the 'unskilled sceptics', and
- the 'neutral unskilled'.

They found an emerging AI divide amongst users. In particular, characteristics such as age, gender and education were key determinants regarding AI marginalisation. For instance, the 'unskilled sceptics' and 'neutral unskilled' cohorts were older, had lower levels of education and privacy protection skills relative to the 'average user' cohort.

The authors also found that attitude played an important part in the AI divide. For example, the 'expert advocate' and 'expert sceptic' cohorts had high levels of AI knowledge, but the 'expert sceptic' group possessed significantly lower levels of AI skills. The main reason for this was that the willingness to utilise AI technology was far more positive for the 'expert advocate' group, which also led to them having a higher level of AI competency.

Moreover, the findings also highlighted the importance of privacy protection skills. The authors found that the 'expert advocates' cohort tended to have greater levels of privacy protection skills than other cohorts. This, therefore, highlights the fact that improving privacy protection skills amongst vulnerable cohorts can lead to increased AI engagement and competence.

In addition to this, studies cited in Bentley, Naughtin, McGrath, Irons, and Coope (2024), identified many factors that contributed to the AI divide for different cohorts. For instance, research on public sector AI adoption in the US showed that it benefited wealthy, male, educated workers with technological experience. In Australia, a lack of AI knowledge was cited as a key factor for a lack of AI adoption. In addition, fear and unwillingness to accept AI were identified in the UK, Germany, and China as inhibitors of adoption. Men in these 3 countries displayed a greater willingness to accept AI. In Australia, women, older people, lower salaried worker and people with less digital access are generally excluded from AI adoption due to less digital confidence (Bentley, Naughtin, McGrath, L. Irons, & S. Cooper, 2024).

The digital transformation has already benefited female workers across the globe, including Australia, by contributing to greater female empowerment (OECD, 2024). This has been made possible due to increased female accessibility to digital tools, the Internet, online platforms, digital financial services, and attainment of necessary digital knowledge and skills. This in turn has increased employment opportunities for women (OECD, 2024). AI adoption in the workplace has the potential to enhance women's employment participation through fairer talent acquisition, Gen-AI powered personalised training, and developing immersive learning experiences (Munshi & Wakefield, 2024)

Under the National Agreement on Closing the Gap, Target 17 aims to ensure Aboriginal and Torres Strait Islander people have equal levels of digital inclusion by 2026. A range of policy levers could potentially contribute to greater digital inclusion (First Nations Digital Inclusion Advisory Group, 2024).

4. How sections of the workforce perceive and experience technology

Varied perceptions of people who use Gen AI

Workers may be using AI in an unofficial capacity for different reasons. For instance, a survey conducted by Slack listed the most common reasons why employees hide their AI usage were: that they ‘feel that using AI is cheating’; they have a ‘fear of being seen as lazy’; and they have a ‘fear of being seen as less competent’. These findings align with survey evidence from Microsoft that found 19% of respondents reported feeling like they are ‘cheating’ when they use Gen AI to undertake their work tasks (Microsoft, 2024). Another survey showed individual workers were more likely to use AI tools if they were employed by larger businesses.[4]

This may depend significantly on the workplace context. In smaller enterprises, where reporting lines are closer and digital experimentation less formally sanctioned, workers often conceal their use of AI tools. In larger organisations, by contrast, the Fusion Digital Agency (2024), survey found that greater resources and clearer governance protocols translated into more openness. Around 83% of employees in large enterprises (over 5,000 staff) reported being transparent about their AI use with employers, compared with just 57% in small businesses. These differences show how organisational capability and governance settings condition individual behaviour.

Differences are also seen across different groups of people in the labour market. For example, recent research has pointed to women’s relationship to AI being shaped as much by perceptions and representation as by occupational exposure. Evidence points to lower reported confidence in digital skills and to women’s continued underrepresentation in STEM pathways and AI development roles (OECD, 2024; OECD, 2024). As Aldasoro, Armantier, Doerr, Gambacorta, & Oliviero (2024) note, this underrepresentation extends beyond technical positions into project leadership and governance, limiting women’s influence over how AI systems are designed and deployed.

Our roundtables and consultations have underscored the need for considering intersectionality more broadly. Women from culturally and linguistically diverse backgrounds described how they encountered both gendered assumptions and language-based bias in recruitment systems using AI. First Nations participants spoke about the compounding of cultural load with digital exclusion. People with disability highlighted that algorithmic systems often compound barriers they already face.

---

[4] Survey evidence from Fusion Digital Agency (2024) found that 81% of white-collar employees used AI tools in small businesses (1–100 employees), compared to 87% in medium-sized businesses (100–500 employees), 88% in large businesses (500–5,000 employees), and 93% for enterprises with over 5,000 employees.

These voices illustrate that AI does not land on neutral terrain: it carries the risk of amplifying existing lines of disadvantage unless inclusive design, along with appropriate and explicit safeguards, is built in.

Recruitment bias

An increasing body of research has shown how the use of AI in recruitment processes can lead to bias in talent selection process (Engler, 2019). Such tools are typically trained based on ‘successful applicant’ databases, which do not always adequately represent people with disability, and can subsequently lead to discrimination in recruitment (Center for Democracy & Technology, 2020).

As another example, recruitment using open-source foundational LLMs was found to affect labour market prospects of women and migrants, given the significant discrepancies between the skill levels and prospects assigned to men compared with women and migrants (UNESCO, IRCAI, 2024).

Recruitment applications could also ‘propagate built-in ableist biases and discriminat[ion]’ (OECD, 2023). For example, using resume data in automated hiring can amplify past biases and career interruptions, while affirmative recruitment measures risk flagging disability without addressing bias in recruitment. The implications of Gen AI depend on how well the technology handles diverse populations – whether datasets are representative, whether algorithms generalise based on the majority, and whether the results echo society’s biases.

In addition to this, Basic Rights Queensland in their submission to the Inquiry into the Digital Transformation of Workplaces, outlined a case study where female candidates with disability lost points in a video interview, due to AI assessing the ‘facial expressions’ of candidates.

Moreover, a report from the Centre for Democracy & Technology argued that employers must not only accommodate applicants with disabilities but also provide effective notice about how AI hiring tools work in the assessment process (Center for Democracy & Technology, 2020). Besides recruitment bias, AI tools are used at a disproportionally higher rate for surveying people with disability compared with their peers without disability, therefore raising concerns about privacy (Trades Union Congress, 2020).

Overseas, bias audits have become mandated in some areas before any automated recruitment tools can be used (see, for example, New York City Council (2021)).

JSA’s Recruitment Experiences and Outlook Survey (REOS) provides some insight into the use of Gen AI in recruitment processes in Australia. From 463 respondents, only 63 businesses (3%) used AI in the recruitment process. The most common use of AI was to generate job descriptions (44 respondents), although several also flagged the use of AI to screen candidates (20 respondents).

These experiences intersect with Australia’s broader policy stance on the safe and responsible use of AI. The Australian Government’s frameworks emphasise transparency, explainability, and fairness as core principles. Embedding these into recruitment applications is essential if they are to serve as enablers rather than amplifiers of bias. Without this grounding, AI in recruitment risks undermining trust not only in hiring processes but in labour market institutions more broadly.

Poor representation in the development of AI

Concerns about representation are central to debates on AI. A persistent issue is the **gender imbalance** in the AI workforce. Recent research highlights the low ratio of women across AI-related roles, shaping how technologies are designed and deployed (Aldasoro, Armantier, Doerr, Gambacorta, & Oliviero, 2024). This imbalance is compounded by race. The Australian Nursing and Midwifery Federation (ANMF) submission noted that health research has historically overrepresented Caucasian women, meaning AI systems trained on these datasets risk producing racially skewed outcomes.

Bias also arises from the **statistical foundations** of AI. Large language models trained on public platforms and social networks embed social distortions, producing harmful associations (WEF, 2023). Applying the same datasets to people with and without disability (WEF, 2023; Packin, 2021) erases structural barriers, while the **cognitive biases of developers and users**, together with the socio-economic context of adoption, further shape outcomes (Walkowiak, 2021).

The impacts on people with disability are especially stark. When underrepresented in training data, workers with disability face discriminatory outputs that limit access to employment and equitable treatment. The United Nations (2025) has stressed that inclusive design and oversight are essential to avoid exclusion disguised as neutrality.

At the same time, the literature points to **practical pathways** for more equitable outcomes. Ethical and responsible AI practices, dedicated research funding, use of disability-relevant datasets, certification of development standards, and better design are widely recognised (Accenture, 2021; OECD, 2023). AI intentionally designed for workers with disability can reduce employment gaps, while generative AI holds particular promise or risk for Aboriginal and Torres Strait Islander people with disability depending on whether cultural and social context is meaningfully embedded (OECD, 2023).

The development of AI applications often lacks involvement of First Nations people

The lack of Indigenous Cultural and Intellectual Property (ICIP) protections has meant AI-generated art continues to threaten the viability of Indigenous art and cultural knowledge (Worrell, 2024). Around 19,000 Aboriginal and Torres Strait Islander people are employed or receive income from the sale of visual arts (Productivity Commission, 2022). Gen AI tools could mass-produce inauthentic, Indigenous-style products, thereby diluting returns for artists. Failing to attribute Traditional Owners, providing false information, disregarding cultural protocols, and creating works that misappropriate ICIP without consultation can lead to an erosion over time of cultural capital as well as economic opportunities that derive from it.

To remedy this, Fitch et al. (2024) argue that ‘reinserting Indigenous sovereignty and self-determination practices within the AI space can alleviate this concern’. Carlson (2023) explains to this end that Indigenous lawyers are finding ways to integrate intellectual property with cultural rights while a global push for Indigenous Data Sovereignty continues to gain momentum to empower ‘economically’ in the AI age.

Working with Gen AI and being involved in how it is designed and applied, is creating opportunities for economic and technological empowerment on the ground. One example is Indigital, an Indigenous-owned profit-for-purpose company, which is also leveraging Gen AI augmentation. By combining Indigenous knowledge with the technology, Indigital is able to generate novel ways to preserve connection to Country, accelerate conservation efforts and educate young people. Another example is of conservation efforts prioritising Indigenous-led expertise in technological design to help preserve biodiversity across Kakadu National Park.

However, the current open Gen AI Tools and/or Vendor based Gen AI tools are mainly trained on ‘biased’ data sets that lean towards non-indigenous societies, thus creating greater risk of exacerbating existing social and economic inequalities (Ghosh, Venkit, Gautam, Wilson, & Caliskan, 2024). Misinterpretation of Indigenous voices, images, techniques, symbolic and linguistic knowledge is also a major concern (UNESCO, 2023).

When taking into consideration the ethics, culture, and equity dimension, First Nations people face AI adoption risks from three main sources which are: concerns over indigenous data and privacy; bias in the AI system leading to misinformed decisions; and underrepresentation of Indigenous engagement leading to cultural bias (Perera, et al., 2025). Moreover, the Australian Senate Select Committee on Adopting Artificial Intelligence (2024) noted that Gen AI was harming job opportunities of Australian First Nation artists in the creative industry and could ‘produce and perpetuate inauthentic and fake art, and [to] appropriate Aboriginal and Torres Strait Islanders’ art, design, stories and culture without reference to Traditional cultural protocols’.

Digital sovereignty

To mitigate the above risks, ‘digital sovereignty’ is one of the key policy initiatives to empower First Nations people to take control of their own data and privacy in the AI adoption process (Vinson, 2025). Digital sovereignty includes all ranges of digital capabilities and infrastructure that would remain free of influence from foreign forces. Its governance would also be in the hands of the host nation and institutions to safeguard national security, confidentiality, economic interests, and cultural values (Misra, Barik, & Kvalvik, 2025). Important areas of digital sovereignty include data sovereignty, infrastructure, standards and technology, ethics and regulations, cybersecurity and resilience, education and literacy, partnership and sustainability (Misra, Barik, & Kvalvik, 2025). Recognition and integration of indigenous knowledge in AI development and an Indigenous community-centred approach of AI development are also proposed as solutions (Misra, Barik, & Kvalvik, 2025; Ghosh, Venkit, Gautam, Wilson, & Caliskan, 2024).

To complement the above, ‘data sovereignty’ is a key measure in AI governance. In Australia, a lack of ‘data sovereignty’, has been identified as a key element undermining community control over AI data for the Indigenous community (House of Representatives, 2025). ‘Data sovereignty’ refers to the principle that governs data by the legal and regulatory framework of the jurisdiction in which data is collected (Patil, Mishra, Chockalingam, Misra, & Kvalvik, 2025). Carlson & Richards (2023) explains to this end that Indigenous lawyers are finding ways to integrate intellectual property with cultural rights while a global push for Indigenous Data Sovereignty continues to gain momentum to empower ‘economically’ in the AI age. The Productivity Commission in their study ‘AI raises the stakes for data policy’ stated that the Australian Government can implement better data design and enforcement of regulation to avoid harm to individuals from ensuring: greater clarity on the applicability and enforcement of existing laws in relation to privacy, cyber security, data protection; anti-discrimination, competition and consumer protection laws; and criminal activity (Productivity Commission, 2024).

Moreover, insights from this Study's consultation process, revealed that while Gen AI has transformative potential, its current trajectory risks deepening existing inequalities, particularly for Indigenous peoples. To mitigate potential risks, stakeholder consultation revealed that some of the following measures were needed:

- Embed Indigenous Data Sovereignty in AI governance
- Develop safeguards to protect Indigenous Cultural and Intellectual Property
- Ensure equitable access to digital tools and markets, and
- Promote critical digital literacy.

A key part of improving how Gen AI is designed and applied is to increase the participation and success of Indigenous people in the tech industry, which is the objective of not-for-profit organisation Indigitek. This organisation was established to normalise the presence of Aboriginal and Torres Strait Islander people in technology, foster culturally safe workplaces, and create accessible pathways into the digital economy.[5]

The organisation's approach to increasing Indigenous participation in the tech industry is underpinned by a strong network of partnerships with training providers and technology companies.

- With training providers, the organisation collaborates to design and deliver technology learning pathways that are tailored to the needs of Aboriginal and Torres Strait Islander learners.
- Tech companies are expected to participate in regular events, aligning with the organisation's strategic priorities, and offer tangible opportunities such as internships and employment placements.

The development of AI applications often lacks an understanding of disability

Gen AI is already having an impact on people with disability. One in every six Australians is living with a disability (AIHW, 2022). Just over half of Australians aged 15–64 with disability is in the labour force (53%), compared to 84% of those without disability – highlighting a notable labour market participation gap (AIHW, 2024). There are more than 1.3 billion people with disability, accounting for 17% of the global population. This makes people with disability the largest global minority group (WEF, 2025). Currently over 2.5 billion people worldwide require assistive products (WHO, 2024). With the global population ageing and an increase in noncommunicable diseases, it is projected that by 2050, approximately 3.5 billion people will need assistive technology (WHO, 2024).

Technological innovations such as wheelchairs, white canes, prostheses, hearing aids and personal computers have assisted people with disability to access employment opportunities and to be as productive as workers without disability (Grijseels, Zuiderent-Jerak, & J. Regeer, 2023; Walkowiak, 2021). AI powered solutions[6] offer new possibilities for people with disability in accessing labour market opportunities, such as, bias neutralisation in interview processes; digital support for physical and mental wellbeing; inclusiveness of training by accommodating different learning styles and preferences; reduced cost of supervision for disabled workers; identification of accessible routes within localities; and, better team design (Walkowiak, 2021;

---

[5] Case study: Indigitek

[6] These include eye tracking and voice recognition software, adaptive learning platforms automated vehicles, image recognition algorithms robots, speech to text algorithm, ADM for hiring, virtual job interview coaching system.

Zhuang & Goggin, 2024; OECD, 2023). In addition, studies identified a list of 45 AI-relevant solutions that facilitate better labour market inclusivity for workers with disability (OECD, 2023).

We heard in our consultations that natural language models can be used in a number of enabling technologies for people with communication difficulties or disabilities. Yet, AI-produced captions can frequently be of poor and unusable quality.

If Gen AI-enabled solutions skew to lower quality supports (for the purposes of reducing costs) this puts at risk the economic and social participation of people with disability. These risks and opportunities are relevant to people in various occupations and skill levels.

Co-design in the development of Gen AI use cases is often key to ensuring AI-driven tools suit the needs of people with disability. One example involved Deaf Mob and aimed to co-design a translating tool between English, Auslan, and Aboriginal sign language (Box 3).

Box 3 Deaf Mob – Gen AI use to bridge communication gaps for Deaf and hard-of-hearing Aboriginal people

The Breaking the Silence: AI for Deaf Mob project was launched to examine whether technology could bridge the divide for deaf Aboriginal and Torres Strait Islander people.

The project research focused on creating a culturally appropriate AI-driven translation tool. The goal was to co-design a framework, methodology, and prototype for translating between English, Auslan, and Aboriginal sign language. The project aimed to make the tool accessible and useful for Aboriginal and Torres Strait Islander communities, as well as mainstream organisations in Health and Justice.

The result of this co-designed process was a DeafMobDoRite prototype app, which was developed in Wiradjuri country. It was designed to improve communication for Deaf and Hard-of-Hearing Aboriginal Australians (Deaf mob), especially in health and justice contexts. The app features several innovative tools:

- AI Chatbot (AskAura): David can use AskAura to upload photos of documents, like a legal document, and ask questions in a conversational manner. The chatbot, trained for Deaf mob and legal contexts, simplifies complex legal terms, empowering users to understand their legal rights and navigate the justice system.
- Pre-recorded Avatar Messaging: Pre-recorded messages allow David to introduce himself and explain his communication needs in unfamiliar environments, reducing anxiety and ensuring his perspective is acknowledged.
- Live Translation: Real-time text-to-speech and speech-to-text functionalities facilitate David's engagement in conversations with professionals, translating across modalities as required.
- Symbol-Based Communication: For users with limited English or Auslan, the app supports personalised symbol-based interaction, making it inclusive for those with additional cognitive or physical impairments. David can use a symbol-based mode to express thoughts through personalized symbols, which are then translated to text and speech, making communication more efficient and accessible for users with different language abilities and physical or cognitive challenges.

Source: Inclusive design, disability, and Gen AI case study

Leveraging co-design and adaptability through implementation

Adaptability is not only a matter of skills uplift but of how technology is introduced and governed in workplaces. While digital literacy and access remain essential foundations, the lived experience of adoption is equally important.

Analysis Paper B introduced the Gen AI personas framework (leaders, change drivers, professionals, enabled workers, affected workers, informed citizens, and educators) as a way of making visible the roles people play in AI adoption. Leveraging co-design is one practical way of connecting this framework to implementation. By embedding personas into trial environments, workers can test and refine how new technologies fit with their real tasks, and feel part of shaping change rather than having it imposed.

The **Digital Transformation Agency's Copilot trial** provides one example of this in practice (Digital Transformation Agency, 2024). By involving staff directly, through digital champions and iterative feedback loops, co-design helped move from shadow use of AI to safe and transparent adoption. Workers had space to raise concerns and suggest refinements, and this process built trust that safeguards and worker agency could sit alongside each other.

This approach is particularly valuable in the non-market sector, where service quality and duty of care heighten the risks of poorly implemented systems. The experimental **HealthMate hypothetical scenario in healthcare** showed how co-design with Gen AI personas could support adaptability in a complex setting. Clinicians were asked to trial the personas in discharge planning and triage.

Leveraging co-design is not the only approach, but it shows how adaptability can be strengthened by involving those who have valuable experience and expertise to inform effective adoption, and who will use and be affected by the technology. In non-market systems such as health, education and social care, there will also be a role for other forms of experimentation and hybrid implementation. Regulatory sandboxes, phased roll-outs, and assurance frameworks can all complement co-design, with each step involving people in testing, evaluating, and adapting new systems. These approaches recognise that adaptability requires a diversity of steps, cultivated through iterative cycles of learning and adjustment, and grounded in the participation of those implementing change.

5. Conclusions

While Gen AI has significant and widespread applications throughout the economy, labour market outcomes depend on how technology is adopted, how productivity impacts are translated and how its gains are shared. A range of decisions (made by employers, workers, governments, and others in society) will determine the extent of any productivity gains and their implication for workers and workplaces.

Moreover, the transition to Gen AI will vary throughout the economy – in terms of both the opportunities and the challenges of the transition for workforce planning to navigate. These impacts will vary for different occupations, skill levels, and industries, and on different groups in society, given the demographic composition of those workforces. In addition, people's experiences of the technology will also be influenced by personal and environmental factors. As such, an intersectional lens will be vital.

It is critical that all workers are able to develop the skills and expertise in demand. These will range from digital skills to critical thinking and interpersonal skills. The adaptation of education and training – including digital skills and literacy – through schools, VET, higher education and other forms of training will be critical to how the economy responds.

Cohorts marginalised due to a lack of digital literacy are also likely to be marginalised due to a lack of Gen AI skills and knowledge. Understanding the broader context of digital literacy in Australia is fundamental to Gen AI adoption. Only a narrowing of the 'digital divide' will enable all Australians to fully engage with Gen AI technology, participate in the labour market and minimise the risk of being left behind.

The readiness of the skills system will be critical in helping to bridge the 'digital divide'. Policies that promote mobilising the entire skills system to prioritise digital/AI capability uplift and activating short form training in digital/AI will ensure that all Australians are upskilled and can keep up with the rapid pace of change.

References

Accenture. (2021). *AI for disability inclusion:Enabling change with advanced technology.* Retrieved from https://metroatlantaexchange.org/wp-content/uploads/2021/06/Accenture-AI-For-Disablility-Inclusion.pdf

AIHW. (2022). *People with disability in Australia 2022.* Australian Government.

AIHW. (2024). *People with disability in Australia 2024.* Australian Government.

Aldasoro, I., Armantier, O., Doerr, S., Gambacorta, L., & Oliviero, T. (2024). The gen AI gender gap. *Economics Letters, 241*. Retrieved from https://doi.org/10.1016/j.econlet.2024.111814.

Australian Digital Inclusion Index. (2021). *Case study: Taking a deep dive into digital ability*. Retrieved from https://digitalinclusionindex.org.au/case-study-taking-a-deep-dive-into-digital-ability/

Bentley, S. V., Naughtin, C. K., McGrath, M., L. Irons, J., & S. Cooper, P. (2024). The digital divide in action: how experiences of digital technology shape future relationships with artificial intelligence. *AI and Ethics*. Retrieved from https://doi.org/10.1007/s43681-024-00452-3

Carlson, B. (2023, September 8). *Indigenous knowledges informing 'machine learning' could prevent stolen art and other culturally unsafe AI practices*. Retrieved from The Conversation: https://theconversation.com/indigenous-knowledges-informing-machine-learning-could-prevent-stolen-art-and-other-culturally-unsafe-ai-practices-210625

Carter, L., Liu, D., & Cantrell, C. (2020). Exploring the Intersection of the Digital Divide and Artificial Intelligence: A Hermeneutic Literature Review. *AIS Transactions on Human-Computer Interaction, 12*(04), 253-275. doi:10.17705/1thci.00138

Center for Democracy & Technology. (2020). *Algorithm-driven Hiring Tools: Innovative Recruitment or Expedited Disability Discrimination?*

De Marco, S., Dumont, G., Helsper, J. J., Díaz-Guerra, A., Antino, M., Rodríguez-Muñoz, A., & Martínez-Cantos, J.-L. (2023). Jobless and Burnt Out: Digital Inequality and Online Access to the Labor Market. *Social Inclusion, 11*(04). doi:https://doi.org/10.17645/si.v11i4.7017

Digital Transformation Agency. (2024). *Evaluation of the whole-of-government trial of Microsoft 365 Copilot: full report.* Retrieved from https://www.digital.gov.au/initiatives/copilot-trial/microsoft-365-copilot-evaluation-report-full

Engler, A. (2019, October 31). For some employment algorithms, disability discrimination by default. Retrieved from https://www.brookings.edu/articles/for-some-employment-algorithms-disability-discrimination-by-default/

First Nations Digital Inclusion Advisory Group. (2024). *First Nations Digital Inclusion Measures.*

Fitch, E., McKenzie, C., Janke, T., & Shul, A. (2024, June 24). *The new frontier: Artificial Intelligence, copyright and Indigenous Culture*. Retrieved February 10, 2025, from Terri Janke and Company: https://www.terrijanke.com.au/post/the-new-frontier-artificial-intelligence-copyright-and-indigenous-culture

FSO. (2024). *Submission 238 to the Senate Select Committee on Adopting Artificial Intelligence.*

Fusion Digital Agency. (2024). *The Impact of AI on the Australian White Collar Workforce.* Retrieved from https://www.google.com/url?sa=t&rct=j&q=&esrc=s&source=web&cd=&ved=2ahUKEwjUl_OpiLSLAxVJcPUHHdGsK6oQFnoECBIQAQ&url=https%3A%2F%2Ffusiondigital.agency%2Fblog-digital-pr%2Fai-australian-white-collar-workforce%2F&usg=AOvVaw2MZZ89PC7i8Oz7Jiuf35ov&opi=89978449

Ghosh, S., Venkit, P. N., Gautam, S., Wilson, S., & Caliskan, A. (2024). Do Generative AI Models Output Harm While Representing Non-western Cultures: Evidence from a Community-Centered Approach. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, (pp. 476-489).

Gmyrek, P., Berg, J., & Bescond, D. (2023). *Generative AI and Jobs: A global analysis of potential effects on job quantity and quality.* ILO.

Grijseels, M., Zuiderent-Jerak, T., & J. Regeer, B. (2023). Technologies for inclusive employment: beyond the prosthetic fix–social transformation axis. *Disability & society, 38*(09), 1534-1557.

Helsper, E., J.A.M. van , A., & Eynon, R. (2015). *Tangible outcomes of Internet use: from digital skills to tangible outcomes project report.* University of Twente. Retrieved from https://ris.utwente.nl/ws/portalfiles/portal/5138116/Tangible-Outcomes-of-Internet-Use.pdf

House of Representatives. (2025). *The Future of Work : Inquiry into the Digital Transformation of Workplaces.* Canberra: Parliament of Australia.

JSA. (2023). *Recruitment Experiences and Outlook Survey – 2023 in Review.* Retrieved from https://www.jobsandskills.gov.au/publications/reos-2023-review

Loble, L., & Stephens, K. (2024). *Securing digital equity in Australian education.* University of Technology Sydney. doi:10.57956/rrpc-5708

Mandala. (2024). *Submission 125 to the Senate Select Committee on Adopting Artificial Intelligence. .*

Microsoft. (2024). *From hype to habit: Exploring the value of generative AI at work.*

Min, Y. (2011). The digital divide among internet users: an analysis of digital access, literacy, and participation. *Journal of Communication Research, 48*(01), 150-187.

Misra, S., Barik, K., & Kvalvik, P. (2025). Digital Sovereignty in the Era of Industry 5.0: Challenges and Opportunities. *Procedia Computer Science*, (pp. 108-117).

Munshi, P., & Wakefield, N. (2024, March). *How AI is being adopted to accelerate gender equity in the workplace*. Retrieved July 2025, from https://www.pwc.com/gx/en/about/inclusion/gender-equity/ai-accelerating-womens-inclusion-workplace.html

National Skills Commission. (2022). *Australian skills classification, release 2.0.* Canberra: National Skills Commission.

New York City Council. (2021). Local Law No. 144 (Automated Decision Tools). *(25), 5*. Administrative Code of the City of New Yor. Retrieved August 25, 2025, from https://www.nyc.gov/site/dca/about/automated-employment-decision-tools.page

OECD. (2023). *Using AI support People With Disability In The Labour Market: Opportunities And Challenges.* Paris: OECD.

OECD. (2023). *Using AI to support people with disability in the labour market.* Paris: Organisation for Economic Co-operation and Development.

OECD. (2024). Algorithm and eve: How ai will impact women at work. *OECD Artificial Intelligence Papers*. Paris. Retrieved from https://doi.org/10.1787/14dc6f89-en

OECD. (2024). *OECD Digital Economy Outlook 2024 (Volume 2) : Strengthening Connectivity, Innovation and Trust.* Paris: OECD Publishing. Retrieved from https://doi.org/10.1787/3adf705b-en.

OECD. (2024). *Using AI In The Workpalce: Opportunities, Risks and Policy Response.* OECD.

Packin, N. (2021). Disability discrimination using artificial intelligence systems and social scoring: Can we disable digital bias. *Journal of International and Comparative Law*.

Patil, A., Mishra, B., Chockalingam, S., Misra, S., & Kvalvik, P. (2025). Securing financial systems through data sovereignty: a systematic review of approaches and regulations. *International Journal of Information Security, 24*(04). Retrieved from https://doi.org/10.1007/s10207-025-01074-4

Perera, M., Vidanaarachchi, R., Chandrashekeran, S., Kennedy, M., Kennedy, B., & Halgamuge, S. (2025). Indigenous peoples and artificial intelligence: A systematic review and future directions. *Big Data & Society*, 1-22. doi:10.1177/20539517251349170

Productivity Commission. (2022). *Aboriginal and Torres Strait Islander visual arts and crafts.*

Productivity Commission. (2024). *Making the most of the AI opportunity : Research paper 2-The challenges of regulating AI.* Canberra.

The Senate. (2024). *Select Committee on Adopting Artificial Intelligence.* Canberra: Parliament House. Retrieved from https://www.aph.gov.au/Parliamentary_Business/Committees/Senate/Adopting_Artificial_Intelligence_AI/AdoptingAI/Report

Thomas, J., McCosker, A., Parkinson, S., Hegarty, K., Featherstone, D., Kennedy, J., . . . Ganley, L. (2023). *Measuring Australia's digital divide: Australian digital inclusion index 2023.* Melbourne: ARC Centre of Excellence for Automated Decision-Making and Society.

Trades Union Congress. (2020). *Congress 2020: General Council Report.* London: Congress House. Retrieved from https://www.tuc.org.uk/research-analysis/reports/general-council-report-2020

UNESCO. (2023). *Indigenous People-Centered Artificial Intelligence: Perspectives from Latin America and the Caribbean,.*

UNESCO, IRCAI. (2024). *Challenging systematic prejudices: an Investigation into Gender Bias in Large Language Models.* Paris: United Nations Educational, Scientific and Cultural Organization. Retrieved from https://unesdoc.unesco.org/ark:/48223/pf0000388971

United Nations. (2025, June 10). Artificial Intelligence for Inclusion: Strengthening Workforce Participation for Persons with Disabilities. Retrieved from https://webtv.un.org/en/asset/k16/k16ieyhu1s

Van Deursen, A., Helsper, E., Eynon, R., & van Dijk, J. (2017). The compoundness and sequentiality of digital inequality. *International Journal of Communication, 11*, 452-473.

van Dijk, J. (2005). The Deepening Divide: Inequality in the Information Society. *Sage Knowledge*. doi:https://doi.org/10.4135/9781452229812

Vinson, N. (2025, April 26). *https://www.technologycouncil.ca.* Retrieved July 2025

Walkowiak, E. (2021). Neurodiversity of the workforce and digital transformation: The case of inclusion of autistic workers at the workplace. *Technological Forecasting and Social Change*.

Wang, C., Boerman, S. C., Kroon, A. C., Möller, J., & de Vreese, C. H. (2024). The artificial intelligence: Who is the most vulnerable? *New Media & Society, 27*(07), 3867-3889. doi:https://doi.org/10.1177/14614448241232345

WEF. (2023, Nov 03). Emerging Technologies: Generative AI holds great potential for those with disabilities - but it needs policy to shape it. Retrieved from https://www.weforum.org/stories/2023/11/generative-ai-holds-potential-disabilities/

WEF. (2025, June 3). Closing the disability inclusion gap with business leadership. Retrieved July 2025, from https://www.weforum.org/impact/disability-inclusion/

WEF. (2025). *Future of Jobs Report 2025.* Geneva. Retrieved from https://reports.weforum.org/docs/WEF_Future_of_Jobs_Report

WHO. (2024, January 2). Assistive technology. Retrieved from https://www.who.int/news-room/fact-sheets/detail/assistive-technology

Workday. (2025). *Elevating human potential: The AI skills revolution.*

Worrell, T. (2024, February 6). Generative AI in the classroom risks further threatening Indigenous inclusion in schools. *The Conversation*. Retrieved from https://theconversation.com/generative-ai-in-the-classroom-risks-further-threatening-indigenous-inclusion-in-schools-222254

Zhuang, K. V., & Goggin, G. (2024). New possibilities or problems for disability and inclusion? The case of AI and ADMs across work. *Telematics and Informatics, 92*.

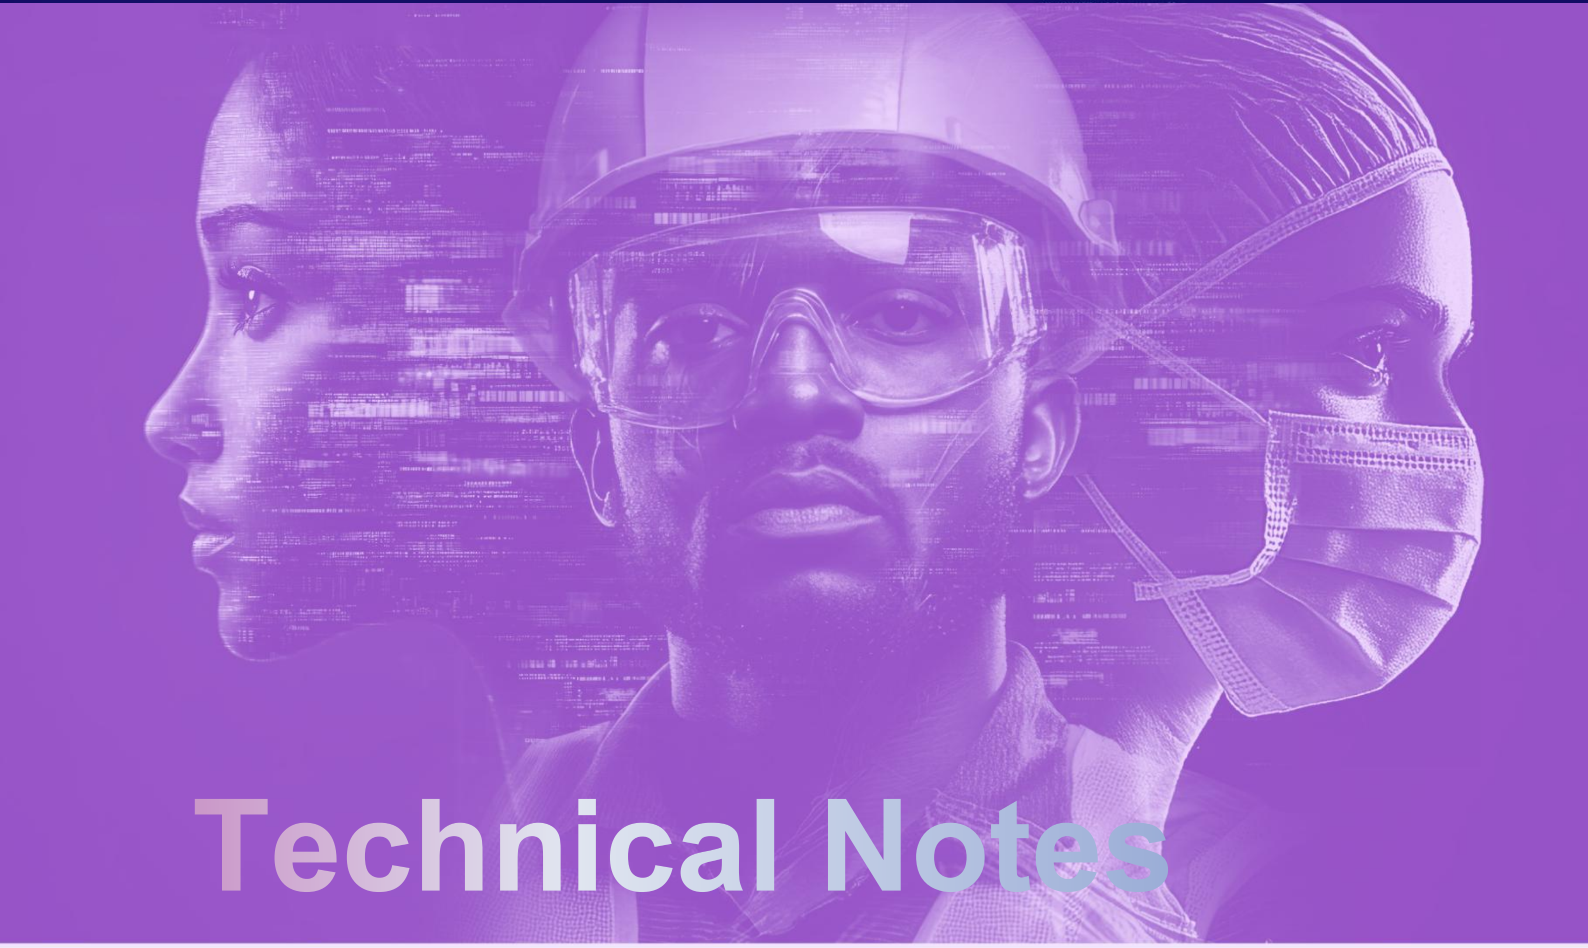

Technical Notes

Developing exposure measures of Gen AI

CGE illustrative scenarios

Hybridisation and specialisation potential

Exploring adaptation effects through online job advertisements

Occupational mobility and AI pathways

Entry-level analysis

Rate of skill change

Returns to experience and skill change

Qualitative research

Developing exposure measures of Gen AI: Automation and Augmentation

This paper outlines JSA's technical methodology for estimating the potential exposure of occupations to Generative AI (Gen AI). Exposure describes the feasible application of technology to existing work tasks, reflecting the state of technological progress in AI and related fields. Estimates of exposure give an indication of how technologies could be applied, including the extent to which tasks could be *either* augmented or automated.

The main purpose of the exposure scores was to assess how exposed different occupations and tasks within the Australian labour market are to the impacts of Gen AI. Exposure, in this context, refers to the potential for tasks to be either automated or augmented through Gen AI assistance.

At the start of the study, the most common measures used around the world to assess occupational exposure were drawn from 'Occupational Heterogeneity in Exposure to Generative AI' by Felten et al. (2021).[1] This therefore represented a natural starting point for our study in considering exposure.

However, after further consideration, this study instead drew upon the task-based framework introduced by the International Labour Organization (ILO) in 2023 (Gmyrek, Berg, & Bescond, 2023). While Felten's method provides a broad conceptual foundation, the ILO approach offers a more granular and contemporary analysis, particularly when adapted to the Australian context using ANZSCO task data and future taxonomies.

Here, we detail the key methodologies that we have tested, the datasets used, and sensitivity tests undertaken to ensure model robustness.

Data

Occupation data

The main data used in the development of exposure were ANZSCO (Australian and New Zealand Standard Classification of Occupations) v1.3 (ABS, 2013) occupations and their associated tasks.

At the start of the study the ANZSCO (Australian and New Zealand Standard Classification of Occupations) was the latest occupation standard. Although later versions of ANZSCO were present (i.e. those reflecting updates in 2022 and 2021) followed by the release of OSCA (Occupation Standard Classification for Australia) in November 2024, for our analysis we mainly use the v1.3 (2013) version. This is because labour market and skills data are still predominantly classified using the ANZSCO v1.3 classification system. In addition, tasks data for OSCA are not available at the unit code level (4-digit) and are instead at the detailed 6-digit level, which is the aggregation that is used for most of our analysis in this study. However, the methods could be adapted to OSCA in the future.

---

[1] For example, see Tech Council of Australia (2023) and Future Skills Organisation (2023).

Census (2021)

Census 2021 (TableBuilder) was used to translate occupational exposures to various cohorts (for example gender, industry and regions), weighted by the share of each category of cohorts as available in Census TableBuilder (2021). Please see **Data Release Bespoke Tables on the JSA website, Tables 5A - 5G** for detailed list of cohorts for which exposures were produced using this data**.**

Method

This study primarily derives exposures by drawing from the ILO method (Gmyrek, Berg, & Bescond, 2023), extends it to include direct assessments for augmentation and enhancing it for the Australian and New Zealand Standard Classification of Occupations (ANZSCO).

Calculating the exposure scores

The method uses data from each task in an ANZSCO unit code, and GPT-4 is used to obtain automation (drawing from the ILO method) and augmentation exposure scores.

Augmentation prompt

> *For the following task associated with the occupation {occupation} (ANZSCO Code: {anzsco_code}), provide a score of potential augmentation with Generative AI.*
>
> *The score of potential augmentation should capture to what extent Generative AI has the potential to impact the way the task is conducted such that the interaction increases either labour productivity or the quality of the output itself without reducing the amount of labour required.*
>
> *The score should range from 0 to 1.*
>
> *Do not provide any other commentary.*
>
> ***Task:** {Task name}*

Automation prompt

> *For the following task associated with the occupation {occupation_name} (ANZSCO Code: {anzsco_code}), provide a score of potential automation with GPT Technology, and a brief justification for the score, given that the job is located in Australia.*
>
> *The score should range from 0 to 1.*
>
> *Do not provide any other commentary*
>
> ***Task:** {Task name}*

- Estimated scores show both the potential to augment (augmentability) and automate (automatability) each work task within an occupation, with two scores per task ranging from 0 to 1.[2]
- The exposure scores of a particular occupation reflect the average potential automatability and augmentability for all work tasks within that occupation.

---

[2] The scores also reflect a non-zero potential of no augmentation or automation of each task, as all exposures are below 1.

- To assess the spread in task-level scores within an occupation, we then calculate the standard deviation of scores across tasks. That is, within the same occupation, some tasks could be highly automatable or augmentable and some the opposite.[3]

Exposure scores therefore allow us to compare how different tasks and occupations might interact with Gen AI across the labour market, rather than estimating the percentage of all tasks potentially automated or time saved per task.

For simplicity, we have not differentiated between 'core' and 'other' tasks, as this would require significantly more information than is available in any existing datasets.[4] Nor have we weighted the tasks in order of complexity, priority or importance.[5] [6] Moreover, exposure scores do not reflect the practical context of the task (including the inherent value of human involvement) which may prevent the use of AI.

Finally, these exposures are measures of the potential for automation and augmentation, and they assume that current Gen AI technologies are implemented in full to their peak potential. Actual automation and augmentation will be lower and more variable where the technology is only partially adopted and evolving. Table 1 shows an example of automation and augmentation scores for Child Care Centre Managers. The accompanying bespoke tables published on the JSA website have scores and AI-generated justifications for each of these scores.

Table 1. Gen AI exposure scores for Child Care Centre Managers (ANZSCO 1341)

Tasks for Child Care Centre Managers (ANZSCO 1341)	Automation	Augmentation
Directing and supervising child carers in providing care and supervision for young children	Low (0.2)	Medium (0.6)
Ensuring the centre is a safe area for children, staff and visitors	Low (0.2)	Medium (0.6)
Providing care for children in before-school, after-school, day, and vacation care centres	Low (0.2)	Medium (0.4)
Complying with relevant government requirements and standards	Low (0.3)	High (0.7)
Developing and implementing programs to enhance the physical, social, emotional, and intellectual development of young children	Low (0.3)	High (0.7)
Liaising with parents	Low (0.3)	High (0.7)

---

[3] A small number of occupations have only one task descriptor, therefore the standard deviation is zero for those.

[4] For instance, the job of a surgeon is made up of several tasks including undertaking surgery. If the surgery itself is the only task that is augmented by Gen AI, the occupation could show a low average score – even though the most important task is affected.

[5] Measuring potential exposure at the task level, rather than human ability or skill levels (as the Felten methodology does) avoids applying strong assumptions about Gen AI performance or the complexity of an occupation. This means we avoid exaggerating exposure across occupations with abilities that are more correlated to language modelling.

[6] We have undertaken analysis by applying time taken for each task as a weight or to rank tasks within an occupation. We found the weighted scores were not more statistically significant than, nor did they vary greatly overall from, the unweighted scores. At this stage, this analysis has been treated as experimental and a check of the robustness of the ILO methodology, as we assessed the time weightings potentially assumed more (or false) accuracy than warranted. Even so, evaluating the complexity and priority of tasks could improve the exposure framework, giving better insight into how well Gen AI could handle tasks and impact work.

Recruiting staff and coordinating professional development	Low (0.3)	High (0.7)
Maintaining records and accounts for the centre	High (0.7)	High (0.8)
Occupation scores (average of tasks)	**Low (0.31)**	**Medium (0.65)**

Source: JSA analysis, ABS (2021) (Tablebuilder).

Guardrails

The use of GPT technologies can come with significant variability in outputs, as the algorithm itself is stochastic, rather than deterministic in nature. To ensure consistent outputs that address the intent of the prompt, the following techniques were used:

1. Temperature settings: The temperature settings were kept very low (i.e. 0.3) to ensure outputs are as deterministic as possible.
2. Function calling: We used OpenAI's function calling features, where we specify that the score column would be a numeric value between 0 and 1, and the justification column would be a justification of the score with a maximum token length of 150 characters.

Noting Gen AI's cutting-edge nature, established data for ‘ground truthing’ analysis is rare and sparse. However, a number of sources were cross referenced to ensure the task scores were representative of the current understanding of Gen AI exposures. A key quantitative method used to assess the differences was in the use of online sentiment data (that is, an analysis of online opinions) as a comparison point.

Comparison of exposure methods

We undertook a systematic comparison of results produced by Felten et al. (2023) and the ILO (2023), including either using a direct 'occupational' concordance from the ILO (2023), or using the method developed by the ILO (2023) and applying it to ANZSCO tasks, the latter being the one that we used. Box 1 briefly describes why we decided to change from using the more established measure of exposure from Felten, to the ILO method.

Box 1: Short comparison of Felten and the ILO methods

Felten – an agnostic measure of exposure

Several recent studies of AI exposure use the method proposed by Felten et al. (2021) for the US labour market. This is based on mapping of AI applications (such as Language Modelling and Image Generation) to 52 human abilities using a crowd-sourced online survey that indicates how each AI application and human abilities relate.

Felten et al. (2023) focuses on Image Generation and Language Modelling for Gen AI specific impact. The key results are that Gen AI exposure rates are highest for high-skilled, white-collar workers such as Professionals and Clerical and Administrative Workers, and lower for unskilled work.

- The method does not distinguish between augmentation and adaptation, meaning the generated scores do not clarify how the technology might be implemented. This raises key questions as to what ‘exposure’ means.

- Occupations with abilities correlated to the ‘Language Modelling’ AI task have the highest exposure scores. This can exaggerate exposure of some occupations relative to others, and could be augmentation-biased – for instance, Judges and Barristers have higher scores than Legal and Court Clerks.
- The database used by Felten et al. (2023) may not capture the abilities of current technology. They select Language Modelling and Image Generation as Gen AI tasks from a list of 10 tasks from 2010 based on the Electronic Frontier Foundation (EFF). There have been several AI tasks that have been identified since then, such as ‘Learning to Learn’ and ‘Safety and Security’.

Accounting for automation as well as augmentation

The latest method used by the ILO addresses these limitations to some degree. The ILO study focused on exposure of occupations to Gen AI via a tasks- and occupation-based analysis (Gmyrek, Berg, & Bescond, 2023).

Measuring potential exposure at the task level (rather than human ability or skill levels as Felten does) means the ILO method does not make assumptions about Gen AI performance or the complexity of an occupation.

The ILO method used GPT-4 to understand which tasks could be augmented or automated for each occupation – by first generating tasks for each occupation and then evaluating the tasks' exposure.

Our chosen methodology involves replicating the ILO approach. However, instead of using AI generated tasks as the ILO method does and mapping the International Standard Classification of Occupations (ISCO) to ANZSCO, we use the ANZSCO tasks data that the Australian Bureau of Statistics has identified for each occupation, at a point in time.

We therefore adapt the ILO methodology to the Australian labour market. Exposure is then calculated for each task within the ANZSCO occupation (example below). We also include *an additional direct and ordinal* measure of augmentation, as opposed to the residual binary measure the ILO applies. This means that our method estimates two exposures for each task. The ILO’s method, in comparison, infers relative augmenting potential for an occupation from task-level automation exposure spreads alone.

As noted above, the scores are largely to aid comparison between tasks and occupations. Our interpretation of the scores are as follows: for a given task – say, *ensuring the centre is a safe area for children, staff and visitors* – a low automation score suggests that very little (if any) of that task could be automated. This suggests that a human worker would have to undertake most of that task in most situations. The medium augmentation score suggests that the worker completing the task would benefit from working *with* Gen AI tools (for instance, in undertaking risk assessments or analysing safety data). The scores cannot be used to infer the quality with which Gen AI might acquit a task, or how easily the potential augmentation or automation could be realised.

Assumptions

The method and analysis here also replicated four of the ILO's original assumptions. First, the framework is based on assuming full technological potential (100% adoption). Second, tasks are performed in high-income countries (how tasks are acquitted varies based on infrastructure, institutional and other settings). Third, GPT-4 predictions likely skew towards technological optimism, missing practical limitations. Fourth, the framework focuses on current automation and augmentation potential without speculating on future job creation (adaptation potential).

Limitations and exclusions

Key limitations reflect those of the underpinning dataset, along with the method used.

First, the tasks used come from the ANZSCO v1.3 classification, which was last updated in 2013. Noting that it is likely tasks for occupations have changed over the past 12 years, the exposures may reflect a more dated view of tasks.

Second, it is assumed that all tasks equally contribute to the overall score. In reality, there is a difference between primary (or distinguishing tasks) for an occupation and secondary (or associated) tasks of an occupation that are incidental aspects of a job. For example, hair styling for hairdressers, who will also likely undertake some clerical or administration work. This method has not attempted to draw these tasks out to compare exposure for core or primary tasks, given the major challenges in constructing such a dataset. Instead, we undertook targeted sensitivity analysis by mapping the ANZSCO tasks to the latest release of the Australian Skills Classification, which contained data on time spent per ASC specialist task (%). When weighting the tasks with time spent per task, there were no considerable differences in the exposure scores in the approaches.

Despite these inherent assumptions and limitations, the large database of tasks and use of standard deviations mean the ILO's framework is reasonably robust. Moreover, the Australian results in this study align strongly with what we have heard from our consultations – that a transformation in how work is done is more likely than widespread automation of work and transitional strategies should capture this.

Validation

Online sentiment data

The website 'Will Robots Take My Job' (willrobotstakemyjob.com) was identified early on, as well as through our steering group consultations, as a unique source of self-reported workers' sentiment on whether the job could be replaced by robots (which we assumed as automation as a whole). The table below shows high scoring occupations for each method. It shows that online sentiment and the ILO-ANZSCO method has a higher degree of concordance for top occupations with online sentiment than Felten results, particularly for higher skilled occupations. From a Pearson's correlation standpoint, the correlation score between online sentiment and ILO-ANZSCO was 0.25, compared to -0.125 for Felten based score.

Table 2: Top 3 highest automation scores for each method compared

Occupation	'Automation' Sentiment	ILO-ANZSCO	Felten
Keyboard Operators	High (91)	High (81)	Medium (61)
Telemarketers	High (90)	High (81)	High (100)
Filing and Registry Clerks	High (68)	High (76)	Medium (55)
Tourism and Travel Advisers	High (68)	High (73)	High (75)
Information Officers	High (82)	Medium (66)	High (76)
Bank Workers	High (88)	Medium (54)	Medium (60)
Judicial and Legal Professionals	Low (28)	Low (34)	High (92)
University Lecturers and Tutors	Low (44)	Low (43)	High (90)
Management and Organisation Analysts	Low (35)	Medium (53)	High (90)

Source: JSA analysis, Gmyrek, Berg, & Bescond (2023), ABS (2013), Felton (2021).

Consultations

The scores were provided to our steering group members as well as the Ministerial Advisory Board as part of the advisory mechanisms for the study, and directly to the ILO for feedback. The ILO noted how our approach aligned with their approach, in particular the use of tasks from classification systems (rather than AI generated tasks).

The ILO also noted their latest release, which expanded the automated approach by first sampling human based responses of potential automatability, and interleaving this into a Retrieval Augmented Generation (RAG) system to provide context for the automated score generation (See (Gmyrek, Berg, Kaminski et al., 2025)). A future study may incorporate this nuance where local data from workers is collected.

CGE Illustrative Scenarios

This technical note documents the methodology applied in the computable general equilibrium (CGE) modelling for this Study as presented in the **Overarching Paper** and **Analysis Paper B - Adoption**. It amalgamates methods from the three scenarios, aligning them to the broader analytical structure of the study.

The CGE illustrative scenarios are intended as a transparent demonstration of an approach, not a prediction of future outcomes. Its purpose is to show how stylised adoption of Gen AI could affect the composition of employment growth over time. The methodology builds on the Victoria University Employment Forecasting (VUEF) model, adapted with a task-based production framework following Acemoglu and Restrepo (2019).

This note sets out the data foundations, the extended production function used to model augmentation and automation, the design of adoption scenarios, the simulation method, and the limitations inherent in this modelling. It also explains how results should be interpreted: as relative shifts in employment growth trajectories, not forecasts of job creation or loss.

Data and model foundation

The CGE analysis used the Victoria University Employment Forecasting (VUEF) model, a dynamic CGE framework developed by the Centre of Policy Studies. It provides a highly disaggregated representation of the Australian economy, covering 124 industries, 97 occupations, and 67 education cohorts. Labour supply is integrated with demographic and education pathways, while demand is generated from production structures at the industry level.

At its base, the model uses a nested constant elasticity of substitution (CES) production structure. Standard CGE applications nest intermediate inputs with value-added components, the latter being a CES composite of labour and capital. For the purposes of this study, the labour nest was expanded to reflect task composition and exposure to AI.

This required embedding the task-based production approach of Acemoglu and Restrepo (2019). In their framework, output is an aggregation of tasks, each of which may be performed by labour, augmented by AI, or automated. We extended this formulation into the VUEF labour nest, linking it to JSA's exposure indices and adoption scenarios.

Incorporating Gen AI into the CGE framework

The Acemoglu and Restrepo task-based production function can be written as:

$$Y = A\int_0^1 [(1 - \varphi_i - \theta_i)L_i + \varphi_i (1+\lambda)L_i + \theta_i K_i]^{\rho} di$$

where output Y depends on total factor productivity A, tasks indexed by i, labour L_i, capital K_i, the augmentable share of tasks φ_i, the automatable share θ_i , the productivity uplift from augmentation λ, and the elasticity parameter ρ.

This structure captures three distinct possibilities for each task: remaining manual, being augmented, or being automated. Augmentation raises the marginal product of labour, while automation substitutes labour with a capital input.

In VUEF, we operationalised this framework through a two-level CES nest:

- Task-level nest (occupation-specific):

$$\mathrm{Q_i} = \left[\alpha_\mathrm{i}\left(\left(\mathrm{L_i^M}\right)_\mathrm{T}\right)^{\sigma_\mathrm{T}} + \beta_\mathrm{i}\left(\mathrm{L_i^A}\right)^{\sigma_\mathrm{T}} + \gamma_\mathrm{i}\left(\mathrm{M_{AI}}\right)^{\sigma_\mathrm{T}}\right]^{(\frac{1}{\sigma_\mathrm{T}}}$$

- Industry-level nest:

$$\mathrm{Y_j} = \left[\delta_\mathrm{j}\,\mathrm{K_j^{\sigma_I}} + (1-\delta_\mathrm{j})\left(\sum_\mathrm{i}\mathrm{Q_i}\right)^{\sigma_\mathrm{I}}\right]^{(\frac{1}{\sigma_\mathrm{I}})}$$

- Economy-wide nest:

$$\mathrm{Y} = \prod_\mathrm{j}\mathrm{Y_j^{\omega_j}}$$

Cost minimisation ensures firms choose the least-cost mix of manual, augmented, and automated tasks subject to:

Objective: $\min_{\{\mathrm{L_i^M},\mathrm{L_i^A},\mathrm{M_{AI}^i}\}}\mathrm{wL_i^M} + \mathrm{w^A\,L_i^A} + \mathrm{c^M M_{AI}^i}$

Constraints: $\mathrm{Q_i} = \alpha_\mathrm{i}\,\mathrm{L_i^M} + \beta_\mathrm{i}\,\mathrm{L_i^A} + \gamma_\mathrm{i}\,\mathrm{M_{AI}^i}$

where w is the wage for manual labour, w^A is the effective wage of augmented labour (reflecting productivity gains), and c^M is the cost of automation (assumed negligible).

Scenario design

To model how Gen AI enters the production system, we imposed stylised adoption curves based on exposure shares. These curves follow an S-shaped trajectory: adoption is slow initially, accelerates in the 2030s, and approaches saturation by 2050.

The three scenarios differ in sequencing:

Scenario 1: Simultaneous adoption. Augmentation and automation proceed together, reaching 50 per cent adoption by 2033. This represents an even-paced rollout across sectors and technologies.

Scenario 2: Non-market lag. Market industries adopt earlier, while health, education, and public administration lag by around five years. This captures possible sectoral constraints on adoption, particularly in labour-intensive services.

Scenario 3: Augmentation first. Augmentation is front-loaded, reaching 50 per cent adoption by 2030, while automation lags to 2035. This reflects a possible sequencing where augmentation tools such as co-pilots and assistants diffuse faster than full task automation.

In all cases, adoption reaches 100 per cent of potential by 2050. These scenarios do not forecast adoption but illustrate different plausible pathways consistent with international evidence and consultation feedback.

Method

The modelling followed a structured process.

First, exposure indices (ANZSCO 4-digit level) for augmentation and automation were mapped to occupations at the ANZSCO 3-digit level. These indices, derived from task-level analysis, provided the base weights for the CES task nest.

The exposure indices were translated to the CES task nest as task shares per type of technology using Frechet Scores (Box 1)

Second, adoption paths were imposed through S-curves. In Scenario 2, adoption differed by sector; in Scenario 3, by mode of technology. These paths determine how the weights α, β, γ evolve over time.

Third, the production function was calibrated to reproduce the JSA baseline employment projections in the absence of AI. Calibration aligns initial task weights, substitution elasticities, and productivity terms so that the model replicates known employment outcomes.

Fourth, simulation shocks were applied from 2025 to 2050, moving task shares from manual to augmented and automated according to each scenario's adoption curve.

Fifth, results were decomposed into their drivers: task substitution (manual → augmented → automated), industry-level reallocation, and demographic impacts. Decomposition differs from calibration. Calibration ensures the baseline model is correct; decomposition explains which channels drive deviations when shocks are applied.

Because total factor productivity (TFP) is exogenous, results reflect the reallocation of labour across tasks and industries, not endogenous innovation or spillover effects. For this reason, results should be interpreted as changes in employment growth relative to the baseline projection, not as forecasts of output levels.

Box 1 Task-level use of Gen AI: from exposure scores to task shares

Each occupation is assigned two exposure scores: one for automation and one for augmentation (Analysis Paper A – Exposure). These indicate the extent to which tasks in that occupation could, on average, be automated by AI or supported by Gen AI tools. Not all tasks are equally exposed. Some are far more suitable for automation or augmentation than others.

To capture this variation in the CGE, a Fréchet distribution is used. This distribution represents the likelihood that Gen AI is the most efficient way to perform each task. It reflects the reality that while many tasks remain largely manual, a minority may be extremely amenable to Gen AI. The “heavy tail” of the distribution allows for a few tasks to be highly suited to automation or augmentation, while most are less so.

By *fitting* a Fréchet curve to the automation and augmentation scores for occupations, we estimate the share of tasks in an occupation that are manual, augmented or automated. In cases where tasks could plausibly fall into both automation and augmentation categories, an overlap adjustment ensures they are not double-counted.

The *fitted distribution* provides an upper bound on the share of tasks in each occupation that could eventually be undertaken by Gen AI. These maximum task shares represent a long-run ceiling, assuming AI is fully available and adopted wherever it is more productive than human labour.

The task shares are then applied in the CGE model. As with other technological change, adoption is assumed to follow an S-curve: slow at first, faster as diffusion accelerates, and tapering as saturation is reached. Over time, the scale parameters of the Fréchet distribution shift along this adoption curve, gradually moving tasks from manual to augmented or automated modes.

This approach produces dynamic, occupation-specific estimates of how task content may change as Gen AI diffuses. These evolving task shares are the key input into the CGE modelling to estimate employment outcomes by occupation and industry.

Figure 1 shows examples of Fréchet cumulative distribution functions. Different shape and scale parameters alter how many tasks are highly exposed to Gen AI. The more slowly the curve rises, the larger the share of tasks that can eventually be automated or augmented.

Figure 1 - maximum task shares by AI technologies in each occupation

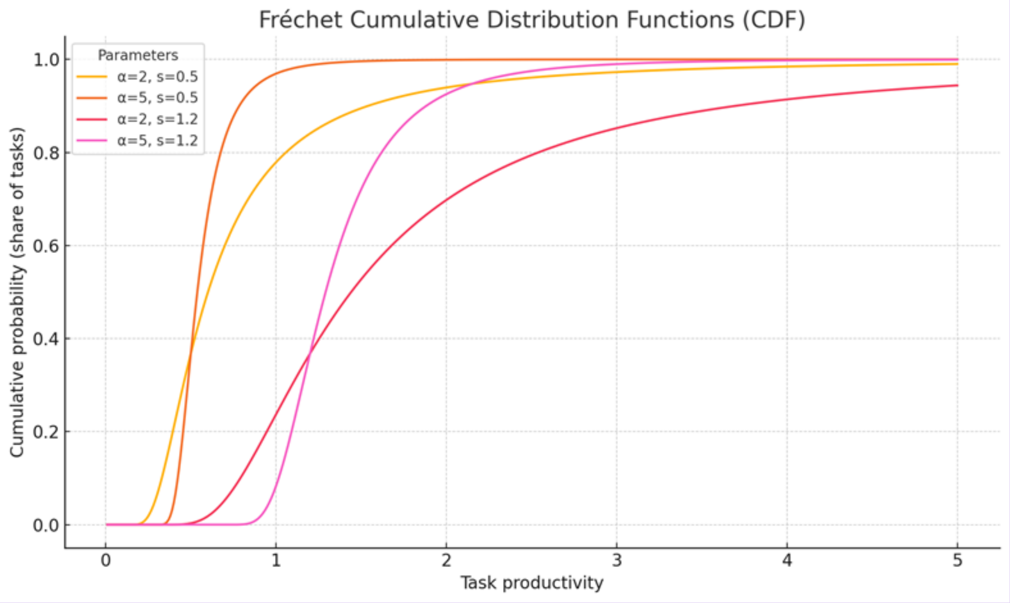

Assumptions

In undertaking CGE modelling of this nature, a range of assumptions are required. These included:

AI is assumed to be costless. Automation is treated as having negligible capital cost, and augmentation requires no additional investment. This isolates productivity effects but overstates potential gains.

Adoption is modelled as occurring in Australia alone. Other countries are not assumed to adopt AI, meaning trade competition effects are excluded.

The labour market retains frictions. Workers cannot instantly change occupations, and wages adjust gradually. This allows for periods of below-baseline employment growth during adjustment.

The skill stock is held constant. The distribution of education and qualifications does not change, meaning the model cannot capture how reskilling or migration could mitigate adjustment costs.

Finally, productivity is exogenous. While task-level productivity shifts are incorporated, broader total factor productivity effects are not generated endogenously.

Limitations

CGE models are powerful tools for analysing structural change, but they have significant limitations.

The most important limitation is that productivity gains are uncertain. Treating AI as costless exaggerates potential efficiency improvements. In practice, firms face substantial costs for infrastructure, data governance, and workforce capability.

A second limitation is that the model does not capture interaction with other technologies. AI adoption will likely be complementary to robotics, cloud services, and data systems, but these complementarities are excluded here.

A third limitation is the assumption of static skills. Because the model holds education and qualification distributions constant, it cannot represent how targeted skills policies might accelerate gains or accelerate recovery.

A fourth limitation is the absence of adaptation. In reality, firms and workers adapt, creating new tasks and occupations. These reinstatement effects, which the literature emphasises as critical, are not generated in the CGE framework.

Finally, while the model disaggregates results by gender and age, it cannot capture other important dimensions of distribution, such as different regional effects, socio-economic background, or firm size.

For these reasons, results should be read as a stylised illustration of how adoption sequencing alters employment growth trajectories, not as predictions of aggregate employment levels or sectoral outcomes.

Exclusions

Several dimensions were deliberately excluded from scope. Firm-level adoption strategies are not represented. International adoption and trade effects are omitted. Sub-state regional differences are not modelled. Institutional feedback loops, such as how AI might alter funding models in education or health, are not considered.

Results framework

Across all scenarios, results exhibit a two-phase pattern: a period of adjustment followed by a period of recovery.

In the adjustment phase, employment growth dips below the baseline projection as AI adoption alters the mix of tasks and workers are reallocated across industries. The depth and timing of this dip differ across scenarios.

- Scenario 1 shows the sharpest adjustment, with simultaneous augmentation and automation creating more immediate disruption.
- Scenario 2 spreads the adjustment over time by delaying adoption in the non-market sector, cushioning short-run impacts.
- Scenario 3 produces a staggered adjustment: augmentation first displaces some labour, followed by a smaller automation dip later.

In the recovery phase, employment growth rises above the baseline. Productivity gains feed through into higher output, and displaced labour is reabsorbed into expanding occupations. Again, the scenarios differ:

- Scenario 1 shows a stronger long-run rebound.
- Scenario 2 has a smoother recovery but smaller overall gain.
- Scenario 3 a balanced trajectory with early cushioning and delayed adjustment.

The key point from these illustrative scenarios is that sequencing matters. How and when AI is adopted changes the timing and magnitude of employment deviations. This shows the importance of transition management, particularly for sectors and cohorts more exposed to automation or augmentation.

Applications and extensions

The modelling provides a transparent framework for considering adoption sequencing. It demonstrates how different adoption pathways shape the trajectory of employment growth, highlighting the role of timing and sectoral context.

Future extensions could address current limitations.

Endogenous skill dynamics

In the current model, the distribution of skills and qualifications across the workforce is held constant, even as tasks shift between manual, augmented, and automated modes. In reality, education, training, and migration will shape the future stock of skills, altering the capacity of individuals to transition into new roles. Embedding a mechanism where training investments or qualification flows feed back into the occupational distribution would allow the model to represent how active skill policies shorten the adjustment phase and accelerate recovery.

Capital costs of adoption

At present, automation is treated as costless and augmentation requires no additional investment. This isolates productivity effects but is unrealistic. Firms face costs for hardware, data infrastructure, governance systems, and workforce capability. Including these costs would alter both the adoption trajectory and the scale of productivity gains, particularly in non-market sectors where budgets are constrained.

Realistic adoption patterns

The current scenarios assume smooth, S-shaped adoption curves imposed exogenously. In practice, adoption is likely to be uneven, subject to regulatory interventions, firm-level capacities, and international diffusion. Introducing stochastic adoption shocks or allowing heterogeneity across firm types would generate more plausible and policy-relevant adoption trajectories.

Linking CGE with microsimulation

CGE provides results at industry, occupation, and demographic cohort levels but cannot capture household- or individual-level distributional impacts. A microsimulation layer could translate CGE outputs into changes in household income, inequality, or welfare, offering richer insights into the social impacts of AI adoption.

Cohort analyses

The cohort analyses by age and gender in the scenarios demonstrate how employment trajectories might differ across groups. However, these should be read with caution. They are highly stylised and do not capture how individuals actually move between occupations, nor do they account for the hours that different groups more commonly work (such as full-time and part-time work). As such, while they illustrate broad structural pressures, they cannot be treated as forecasts of specific cohort outcomes.

Changing exposures

The exposure indices used in this modelling are static, but in practice they will evolve as technology develops. The indices were calculated using Fréchet scores, which provide a probabilistic measure of exposure but are **fitted** to the actual task-level scores. In reality, there is greater task-level granularity than the Fréchet distributions capture.

ANZSCO mismatch

Exposure scores are calculated at the ANZSCO four-digit level, which provides detail for specific occupations. The CGE model, however, operates at the ANZSCO three-digit level, aggregating across multiple four-digit occupations. This aggregation masks variation where different four-digit occupations within the same three-digit group have very different exposure profiles. As a result, the modelling cannot capture fine distinctions in how AI may affect related occupations.

Validation

Results were validated against three sources. First, exposure indices from JSA's task-level analysis ensured occupational outcomes were consistent with exposure patterns. Second, adoption dynamics from Analysis Paper B ensured scenarios were aligned with observed evidence on adoption to date. Third, results were compared with the international literature, particularly Acemoglu and Restrepo (2019), which also finds a two-phase pattern of adjustment followed by recovery for general-purpose technologies.

Hybridisation & specialisation potential

This technical note sets out the methodology used to assess how occupations may adapt under the influence of Gen AI through processes of hybridisation and specialisation. It complements the adaptation analysis presented in **Analysis Paper C - Adaptation**, extending the exposure framework to explore how automation and augmentation could reshape the similarity of occupations to one another.

The analysis provides a benchmark for understanding where existing roles may become more alike, enabling hybridisation, and where new specialised variants may emerge through augmentation. It is not a prediction of future occupations. Rather, it is a diagnostic approach, helping policymakers and practitioners to understand where opportunities or pressures for adaptation may be strongest within the existing occupational taxonomy.

Data

The analysis uses the ABS ANZSCO v1.3 task dataset, which provides task-level descriptions for each four-digit occupation. These tasks are linked to automation and augmentation exposures derived by JSA in our Study.

Each task description is encoded into a semantic vector representation using natural language processing models. This allows similarity between occupations to be measured by comparing their task vectors. These similarity scores form the foundation for assessing hybridisation and specialisation potential.

Method

The analysis proceeds in three structured steps, producing baseline similarity scores and differential measures of hybridisation and specialisation potential. This approach follows the framework documented in Analysis Paper C.

Computing hybridisation and specialisation potential scores

As discussed in the analysis papers, the theoretical framework for job adaptation from Gen AI at the task level was applied as follows, with the mathematical expressions included in Box 1.

Step 1: Compute baseline similarity

We estimate how similar occupations are by comparing the words used to describe their tasks. We use a model called ‘sentence-transformers/all-MiniLM-L6-v2’ - a popular approach for finding similarities in sentences - to convert each task description into a 384-dimensional vector, with equal weighting. Then, we calculate the cosine similarity between these vectors to get a baseline similarity score between each task for each pair of occupations.[1]

Step 2: Derive candidates for job hybridisation

We identify jobs that could become more similar as AI reduces the intensity of certain tasks. This means one person could handle multiple tasks that were previously divided between AI and non-AI roles. As with Step 1, we use a model called ‘sentence-transformers/all-MiniLM-L6-v2’ to convert each task into a vector.

---

[1] See Sentence Similarity Models – Hugging Face for further details.

We then take a weighted average of these vectors, where the weights are based on how much the tasks can be automated. This helps us calculate how similar jobs are without the AI-intensive tasks.

The difference between this new similarity score and the original score shows how much job roles could merge due to AI. Table 1 lists the job pairs with the highest increase in similarity.

Table 1: Some occupation pairs would become more similar if automation-exposed tasks are automated

Occupation pairs with greatest increase in similarity – higher automation exposed tasks weighted lower

Occupation 1	Occupation 2	Non-AI intensive Similarity score (*similarity score increase)*
Gallery, Library and Museum Technicians	Visual Merchandisers	0.75 *(+0.10)*
Inspectors and Regulatory Officers	Other Clerical and Office Support Workers	0.66 *(+0.07)*
Electronic Engineering Draftspersons, Technicians	Electronics Trades Workers	0.76 *(+0.07)*
Human Resource Professionals	Policy and Planning Managers	0.63 *(+0.06)*
Other Miscellaneous Labourers	Ticket Salespersons	0.67 *(+0.06)*
Importers, Exporters and Wholesalers	Telemarketers	0.72 *(+0.06)*
Electronics Trades Workers	Mechanical Engineering Draftspersons, Technicians	0.76 *(+0.06)*
Conference and Event Organisers	Supply, Distribution and Procurement Managers	0.71 *(+0.06)*
Gallery, Museum and Tour Guides	Security Officers and Guards	0.60 *(+0.06)*
Gallery, Library and Museum Technicians	Gallery, Museum and Tour Guides	0.66 *(+0.05)*

Note: Excludes occupation pairs with a similarity score below 0.60.

Source: JSA analysis; ABS ANZSCO v1.3.

Step 3: Derive candidates for job specialisation

We investigate candidates for job specialisations by estimating the similarity of occupation pairs with a revised weighting. Tasks with higher automation exposure are weighted higher to reflect potential changes in job content due to augmentation. Close clusters of jobs under this scenario compared to the baseline may indicate new specialised jobs that combine similar tasks from previous roles.

As with Step 1 and Step 2, we again use the ‘sentence-transformers/all-MiniLM-L6-v2’ model to encode all tasks into an embedding structure with equal weighting. Then, we take a weighted average of the task embeddings based on the augmentation scores. This allows us to calculate the AI-intensive job-pair similarity. The difference between the AI-intensive similarity score and the baseline similarity indicates the potential increase in specialisation due to Gen AI adaptation

This approach provides a conceptual view of 'first order' adaptation. That is, where automation and augmentation may cluster tasks and which can help identify where specialised or hybrid roles could emerge within existing occupations and roles. However, it cannot account for the creation of entirely new tasks, (for example to support the roles of Gen AI Personas).

Box 1 Estimating hybridisation and specialisation potential

Step 1. Baseline similarity

The starting point is the baseline similarity between occupations, calculated according to current tasks and conditions. For each occupation j, the baseline vector is constructed by averaging the embeddings of all its tasks:

$$v_j^{base} = \left(\frac{1}{T_j}\right) \Sigma_{\{t \in j\}} v_t$$

where T_j is the number of tasks in occupation j and v_t is the embedding of task t.

The cosine similarity between occupations j and k provides the baseline similarity score:

$$S_{jk}^{base} = cos(v_j^{base}, v_k^{base})$$

This establishes a benchmark for how similar occupations are under the current distribution of tasks.

Step 2. Hybridisation potential

Hybridisation potential measures the extent to which occupations become more similar once automation-exposed tasks are down-weighted. The rationale is that if automation reduces the importance of certain tasks, the remaining task sets may converge, opening possibilities for hybrid roles.

For each occupation, tasks are reweighted by (1 - A_t), where A_t is the automation exposure of task t. The non-AI-intensive vector is then:

$$v_j^{nonAI} = \left(\frac{1}{T_j}\right) \Sigma_{t \in j} (1 - A_t) v_t$$

Pairwise similarity is recalculated as:

$$S_{jk}^{nonAI} = cos(v_j^{nonAI}, v_k^{nonAI})$$

The hybridisation differential is then:

$$\Delta S_{jk}^{hybrid} = S_{jk}^{nonAI} - S_{jk}^{base}$$

A *positive differential* indicates that occupations are drawn closer together once automation-exposed tasks are reduced in weight.

Step 3. Specialisation potential

Specialisation potential captures the extent to which occupations become more similar once augmentation-exposed tasks are emphasised. The rationale is that augmentation may cluster tasks into new combinations, suggesting opportunities for specialisation.

For each occupation, tasks are reweighted by their augmentation exposure U_t:

$$v_j^{AI} = \left(\frac{1}{T_j}\right) \Sigma_{t \in j}(U_t) v_t$$

Pairwise similarity is recalculated as:

$$S_{jk}^{AI} = cos(v_j^{AI}, v_k^{AI})$$

The specialisation differential is then:

$$\Delta S_{jk}^{hybrid} = S_{jk}^{AI} - S_{jk}^{base}$$

Positive values suggest that *augmentation increases similarity*, pointing to potential for specialisation. *Negative values* suggest *augmentation reinforces distinctiveness*.

Source: JSA Analysis

Assumptions

Several assumptions shape the results.

First, the analysis assumes that the ANZSCO task taxonomy accurately reflects the content of occupations. In practice, tasks vary across workplaces, industries, and regions. Some tasks are omitted, and others may differ in emphasis. The taxonomy provides a consistent baseline but cannot capture all real-world variation, or provide a comprehensive or highly representative set of tasks.

Second, automation and augmentation exposures are treated as static. These scores reflect potential as of 2024-25 but will evolve as AI technologies develop. The analysis cannot account for future changes in exposure or for differences in the pace of adoption across sectors. This is a consistent limitation of methodologies in the Study that rely on the exposure scores.

Third, similarity scores are based on task language. Embedding models capture semantic similarity in descriptions, but tasks that are described using similar words or phrases are not always functionally equivalent. The method assumes language maps closely to task content and therefore that similarity in language reflects similarity in the nature of the tasks.

Fourth, adaptation is assumed to flow from task-level changes. Organisational, institutional, and regulatory influences are not modelled, even though they strongly shape how roles evolve in practice.

Finally, the method assumes that hybridisation and specialisation are captured through changes in similarity. This provides a structured benchmark, but it does not predict the creation of entirely new tasks or occupations.

Limitations

The framework is diagnostic, as with a number of the methodologies applied in this Study. Results should be read as benchmarks - to understand the potential for change - and not forecasts.

- **Size bias in hybridisation measures**. Occupations with few tasks can display volatile similarity scores. Their relative weights shift significantly when automation or augmentation exposures are applied. Large occupations, by contrast, have more stable profiles.

- **Evenness but not relatedness**. The method identifies whether occupations become more or less similar, but it does not measure the relatedness of tasks. Two occupations may appear closer under the model even if their tasks are functionally distinct.
- **Noise in diversity of task overlaps**. A small number of shared tasks can disproportionately affect similarity scores. This is analogous to noise in diversity metrics, where one-off transitions or overlaps carry the same weight as substantive flows.
- **Inheritance of sensitivities**. Because the framework compares baseline and differential scores, any instability in the baseline is carried through to hybridisation and specialisation measures. Occupations near thresholds can shift category with small changes in weighting.
- **Dependence on task wording**. Embedding models are sensitive to the language of task descriptions. Two tasks that are written differently but performed similarly may appear less related. Conversely, tasks that share common language may appear closer than they are functionally.
- **Static exposures**. Automation and augmentation scores are treated as fixed. In practice, exposures will change as AI technologies evolve and diffuse.

Parameter sensitivities & robustness checks

Parameter sensitivities and robustness can be tested in several ways as an extension of this analysis.

- **Multi-year pooling**. Constructing similarity measures across multiple snapshots of the ANZSCO dataset allows us to assess whether results are stable across time rather than reflecting one year’s composition.
- **Threshold analysis**. Sensitivity to similarity changes can be tested. Excluding weak overlaps ensures that results are not driven by trivial commonalities.
- **Weighting schemes**. Alternative weights on automation and augmentation exposures can be applied to test whether results depend heavily on parameter choices.
- **Bootstrapping**. Resampling tasks within occupations can produce confidence intervals around baseline and new similarity scores. This helps explain where hybridisation or specialisation signals are robust and where they are more susceptible to variability.

For the Study, our checks focused on threshold analysis and weighting schemes, given tasks are relatively stable over time.

Exclusions

The framework does not attempt to predict the creation of entirely new tasks. It is grounded in the existing task taxonomy in ANZSCO, which is static by design.

It does not model labour demand or wages. Similarity shows where occupations may converge or cluster, but it does not indicate whether demand for these roles will rise or fall. Individual-level adaptation such as retraining, part-time to full-time shifts, or occupational mobility across sectors, is not represented. These dynamics are critical in practice but outside scope.

Regulatory barriers are also excluded. Professional licensing, training requirements, and industrial relations all shape adaptation but are not captured here.

Results

The results highlight where hybridisation and specialisation signals are strongest, both in quantitative terms (per the examples of occupations in Table 2) and through illustrative examples of how occupations may cluster differently when automation and augmentation are emphasised (as can be seen in Figures 1 to 3).

Table 2. Hybridisation and specialisation potential of selected occupations

Occupation	Hybridisation potential	Specialisation potential
Telemarketers	-2.13	2.50
Textile and Footwear Production Machine Operators	-6.25	11.14
Ticket Salespersons	-9.13	-17.53
Timber and Wood Process Workers	-2.36	11.29
Toolmakers and Engineering Patternmakers	-5.43	2.59
Tourism and Travel Advisers	-1.04	0.28
Upholsterers	-6.69	9.06

Note: Excludes occupation pairs with a similarity score below 0.60.

Source: JSA analysis of ANZSCO v1.3 task dataset.

Interpretation of the data for selected occupations in Table 2

Hybridisation

Most occupations record negative hybridisation scores. This means that when automation-exposed tasks are down-weighted, occupations tend to become less similar. Hybridisation potential is therefore limited. Only a handful of occupations, such as Vehicle Painters and Vehicle Body Builders and Trimmers, show positive signals.

Specialisation

Specialisation scores are more frequently positive, particularly in trades and manufacturing. Textile and Footwear Production Machine Operators, Timber and Wood Process Workers, and Upholsterers show high clustering under augmentation.

Some occupations show negative specialisation potential. Ticket Salespersons, Veterinarians and Travel Attendants become more distinct under augmentation weighting, reinforcing their uniqueness.

Graph-based analysis.

Network analysis shows similar patterns. For example, Conference and Event Organisers become more similar to Technical Sales Representatives and Supply, Distribution and Procurement Managers under augmentation (Figure 1). Secondary pathways also appear (Figure 2), and clusters of specialisations emerge, particularly in engineering, health, and clerical work (Figure 3).

Figure 1: Jobs closer to Conference and Event Organisers (green) and previously closer

Figure 2: Jobs closer to Conference and Event Organisers with proximity to Technical Sales Representatives and Supply, Distribution and Procurement Managers.

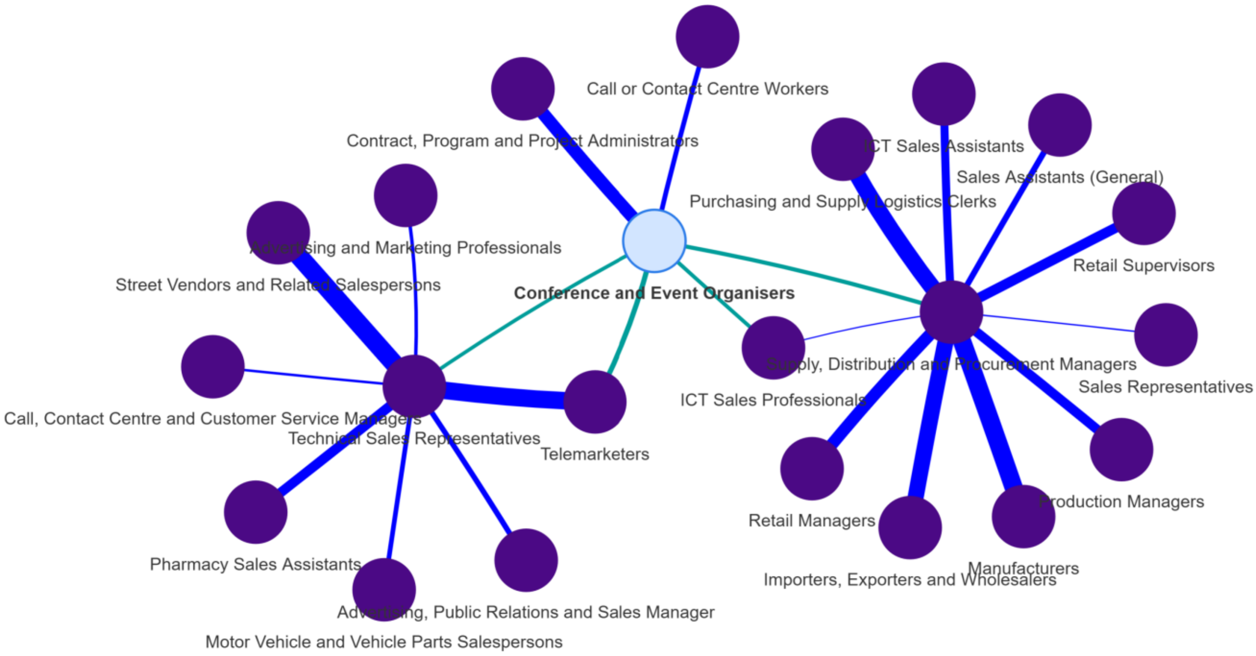

Figure 1. Potential specialisation clusters as a result of Gen AI augmentation

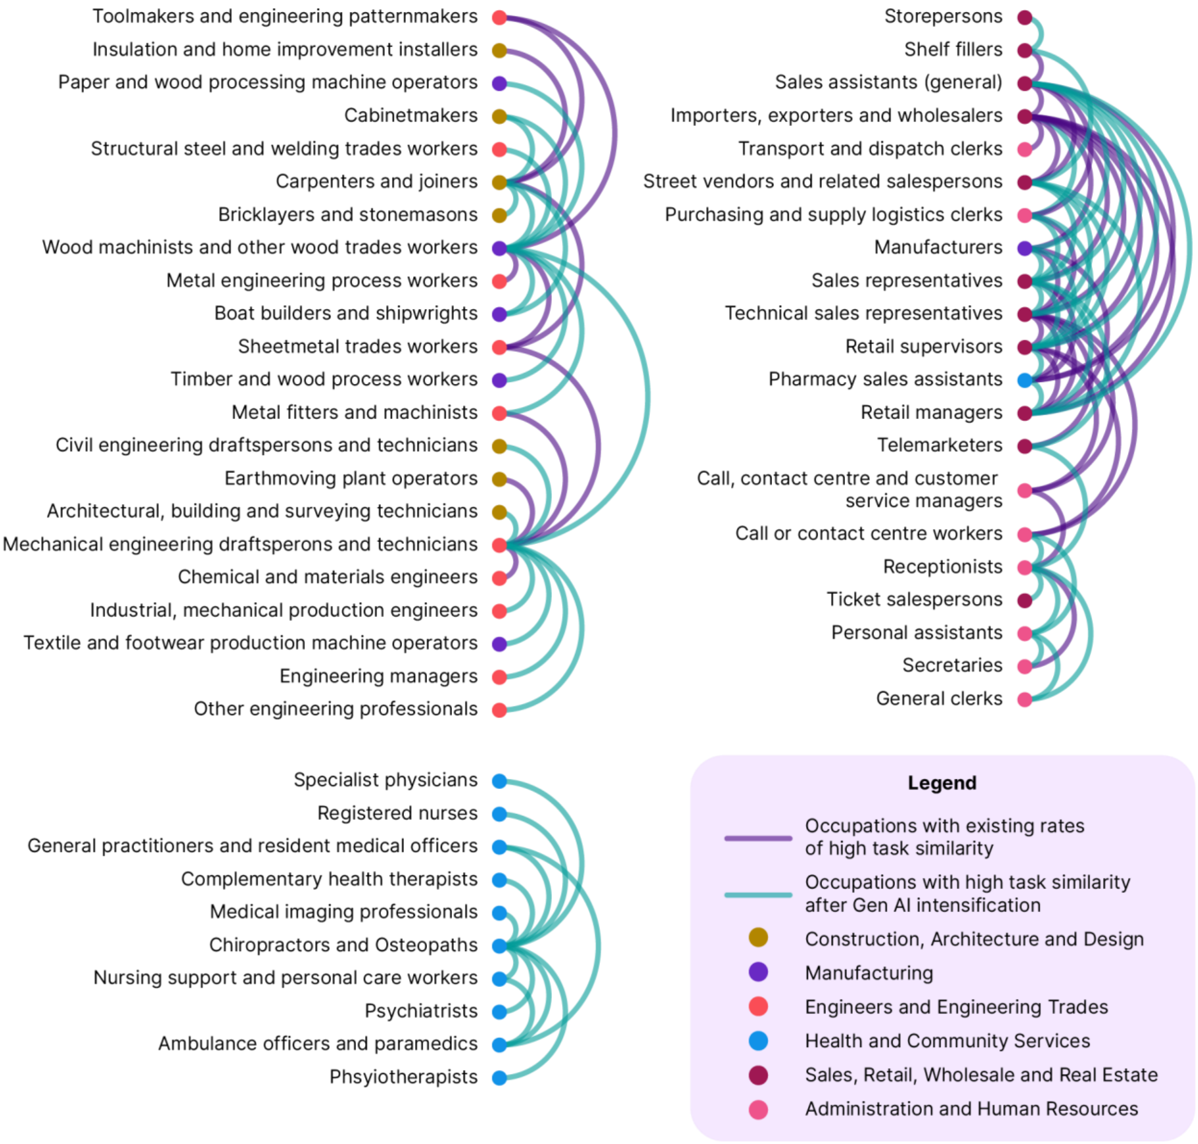

Source: JSA analysis; ABS ANZSCO v1.3.

Applications and extensions

The framework provides a structured means of assessing occupational adaptation and can be extended in several directions to improve its relevance over time.

Dynamic exposures would update automation and augmentation scores as technologies evolve, ensuring results reflect the changing task landscape rather than a static snapshot.

Integration with observed mobility flows (DOM) would link conceptual similarity-based pathways with actual observed labour market movements between occupations, offering insight into both potential and realised transitions.

Network analysis would capture secondary pathways of similarity, showing how changes in one domain can cascade through related occupations and revealing indirect vulnerabilities or opportunities.

Applying **sectoral lenses** would highlight domain-specific dynamics, such as the role of regulation and licensing in health, education, or the trades, and indicate where transferable skills or new specialisations emerge.

Together, these extensions enhance the framework's **policy utility**, providing a stronger evidence base for where training investment, regulatory attention, or support for emerging roles is most needed.

Validation

Validation was carried out in three ways, as part of the full validation for the Study.

First, by comparing results with labour market evidence. The occupations showing the strongest specialisation potential -clerical, administrative, and sales roles -align with skill changes observed in job advertisement data.

Second, by cross-checking with Analysis Paper C. The task-based similarity clusters identified here match adaptation signals documented in Paper C, especially in clerical and digital roles.

Third, by testing sensitivity. Results were re-estimated with alternative weighting schemes and similarity thresholds. While marginal occupations shift categories, the overall patterns of weak hybridisation and stronger specialisation remain consistent.

Exploring adaptation effects through online job advertisements

This technical paper relates to analysis published in Analysis Paper C: Adaptation (Paper C) in *Our Gen AI Transition: Implications for Work and Skills Analysis Papers* that explored job adaptation. It extends on the brief overview of the approach included in Box 1 Exploring adaptation effects through online job advertisements in Paper C.

Data

We study changes in the skill content of jobs using online job advertisement data from Lightcast, comparing the years 2012 and 2024. This data is scraped by Lightcast from Australian job advertisements posted on publicly available online job boards.

Online job advertisement data is well-suited to measuring change over time in the skill requirements of jobs. This is due to the detailed content, granularity, timeliness of online job advertisement data (Tsvetkova, D'Amico, Lembcke, Knutsson, & Vermeulen, 2024). Limitations of online job advertisement data are discussed in detail in the Limitations section below.

Method

Adaptation and change in the skill content of jobs

Change in the skill content of jobs is a common form of adaptation as technological advances are harnessed into new production processes. Our approach to exploring such changes through online job advertisements involves two steps.

First, we capture the prevalence of specific skills in Australian online job advertisements by occupation in 2012 and 2024 (the earliest complete year of comparable data and the latest complete year), using the Lightcast skill taxonomy. For this analysis, we use the most granular level of the Lightcast skill taxonomy (which includes over 33,000 unique skills). The start and end year for the analysis were chosen with the intention of capturing medium term changes in the skill composition of jobs (i.e. across a period of 10 to 15 years).

Second, we compare the composition of skills listed in online job advertisements by occupation between the two periods. This enables us to identify:

- which skills have substantially risen (or fallen) in demand
- which skills are listed in 2024 which were not listed at all in 2012 and vice-versa.

While this approach does not directly measure changes in production processes and the reorganisation of work, by examining which specific skills are rising and falling in demand we can draw reasonable inferences about changes within a given occupation.

Adaptation and change in the skill complexity of occupations

We extend on our analysis of changes in the skill content of jobs by examining their implications for the skill level (or skill complexity) of occupations. Our approach for exploring changes in skill complexity involves three steps.

First, we calculate a skill level value for each unique skill in the Lightcast skill taxonomy. This skill level value is the weighted average of the ANZSCO Skill Level of the occupations in which that skill is listed in job advertisements. For example, if a particular skill is listed only in Skill Level 1 occupations, then its skill level value would be 1. Whereas if 80% of occupations listing a particular skill are Skill Level 2 occupations and the other 20% are Skill Level 3 occupations, the skill level for that skill would be 2.2.[1]

Second, we determine the weighted average skill level value of each in-scope occupation for 2012 and 2024 based on the skill level values calculated in the step above. That is, skills that were listed in a higher proportion of online job advertisements for the occupation in the relevant year are weighted higher.

Third, we isolate newly appearing skills (i.e. skills that were not present in online job advertisements for an occupation in 2012 but were present in 2024) to analyse whether these skills have tended to increase the skill complexity of the occupation.

Assumptions

Each unique skill is assigned a single skill level value

The approach for calculating the skill level value of each skill does not account for potential systematic differences between groups of occupations in relation to the proficiency levels required for the same skill. The potential impact of this assumption is likely to be greatest in relation to common skills (i.e. skills that are applicable across a wide range of jobs). For example, the skill 'Communication' was present across occupations from every ANZSCO Skill Level in 2024, with the weighted average skill level of Communication being 2.3. In practice, the proficiency of communication required may differ between occupation groups.

This assumption is required in the absence of sufficiently granular, regularly updated skills-related information for occupations in the Australian labour market and the proficiency level related to these different skills.

Limitations

Online job advertisements are not representative of all job openings

Online job advertisements data over- or under-represent certain occupation groups based on their recruitment methods. For instance, a higher proportion of businesses use recruitment websites and job boards when recruiting for Managers (64%) and Professionals (64%) compared with Technicians and Trades Workers (56%), Labourers (48%) and Sales Workers (44%) (National Skills Commission, 2022).

Given the focus of this analysis on within-occupation changes in skill composition, not having comprehensively representative data on the distribution of job openings across occupation groups is not a critical limitation. However, it is possible that imperfect representativeness of online job advertisement within occupations could have an effect. For example, businesses recruiting for a particular occupation using online job boards may differ from those recruiting using word of mouth with respect to the extent of technology integration more broadly. If this is

---

[1] Some occupation unit groups contain occupations across multiple ANZSCO Skill Levels. Where this is the case, the mid-point between the Skill Levels is chosen. For example, the unit group Medical Technicians includes occupations at Skill Levels 2 and 3. Therefore, the Skill Level of this occupation for the purposes of the calculation is 2.5.

the case, it is possible this approach could potentially overstate technology-related dynamism in the skill requirements for the occupation.

Some skills may be implicit in job advertisements
Online job advertisements are not primarily intended as a comprehensive stocktake of the skills required for a job. In some instances, skills that are essential for the job may not be explicitly stated, e.g. due to possession of the skill being assumed for those who possess the required qualification, experience or related skills. Changes in relation to previously or newly assumed or implicit skills will not be captured by this approach.

Exclusions

We exclude from our analysis any occupations with fewer than 100 online job advertisements in either 2012 or 2024. This is done to improve the robustness of the results, ensuring each in-scope occupation has a minimum viable sample of online job advertisements contributing to the analysis. Applying this threshold leaves us with 245 out of 358 occupation unit groups (4-digit ANZSCO occupations) as in-scope for analysis.

Results

At the unique skill level, skill level values range across the full spectrum from 1 to 5. As Table 1 reveals, the most listed skills at the higher end of skill level values typically relate to leadership and management whereas the most listed skills at the lower end typically relate to the performance of routine manual work.

Table 1. Common skills with high skill level values centre around leadership and management

Top 5 skills listed in online job advertisements in select skill level value ranges, 2024

Top skills (1 to 1.99)	Top skills (2 to 2.99)	Top skills (3 to 3.99)	Top skills (4 to 5)
Management	Contract Management	Invoicing	Showrooms
Leadership	Student Services	Gardening	Coffee Making
Planning	Go-to-Market Strategy	Listening Skills	Machine Operation
Project Management	Growth Mindedness	Sales Process	Physical Stamina
Mentorship	Vendor Management	Pruning	Loading And Unloading

Note: A high skill level value (closer to 1) indicates that the skills listed in online job advertisements for this occupation are heavily concentrated among highly skilled occupations per ANZSCO skill levels. A low skill level value (closer to 5) indicates the opposite.

Source: JSA analysis of Lightcast job advertisement data.

Table 2 presents the occupations with the highest and lowest skill level values based on the skill content of online job advertisements for the occupation. The highest skill level values are predominantly found among Professional occupations, with the lowest skill level values concentrated among Machinery Operators and Drivers, and Labourers.

Table 2. The skill content of Professional occupations skews most heavily toward higher skill levels

Occupations with the highest and lowest skill levels, 2024

Panel A. Occupations with the highest skill level values	Skill Level Value	Panel B. Occupations with the lowest skill level values	Skill Level Value
Multimedia Specialists and Web Developers	1.41	Forklift Drivers	3.36
Software and Applications Programmers	1.44	Packers	3.12
Psychiatrists	1.48	Laundry Workers	3.12
Other Medical Practitioners	1.50	Commercial Cleaners	3.08
Specialist Physicians	1.54	Storepersons	3.08
Computer Network Professionals	1.55	Shelf Fillers	3.07
Secondary School Teachers	1.55	Truck Drivers	3.05
Surgeons	1.55	Freight and Furniture Handlers	3.02
Database and Systems Administrators, and ICT Security Specialists	1.55	Concreters	3.02
General Practitioners and Resident Medical Officers	1.59	Domestic Cleaners	2.99

Note: A high skill level value (closer to 1) indicates that the skills listed in online job advertisements for this occupation are heavily concentrated among highly skilled occupations per ANZSCO skill levels. A low skill level value (closer to 5) indicates the opposite.

Source: JSA analysis of Lightcast job advertisement data.

Comparing the average skill level value of newly appearing skills with that of the occupation overall reveals that new skills may be higher or lower skill depending on the occupation. Figure 1 shows the professional occupations with greatest differential between new skills and the occupation overall. For example, the new skills for Technical Sales Representatives tend to have a higher skill level value than the average for skills across this occupation. The inverse is seen for Psychiatrists.

Figure 1. New skills may have higher or lower skill values than the occupation overall

Professional occupations with the largest differentials between newly appearing skills and all skills

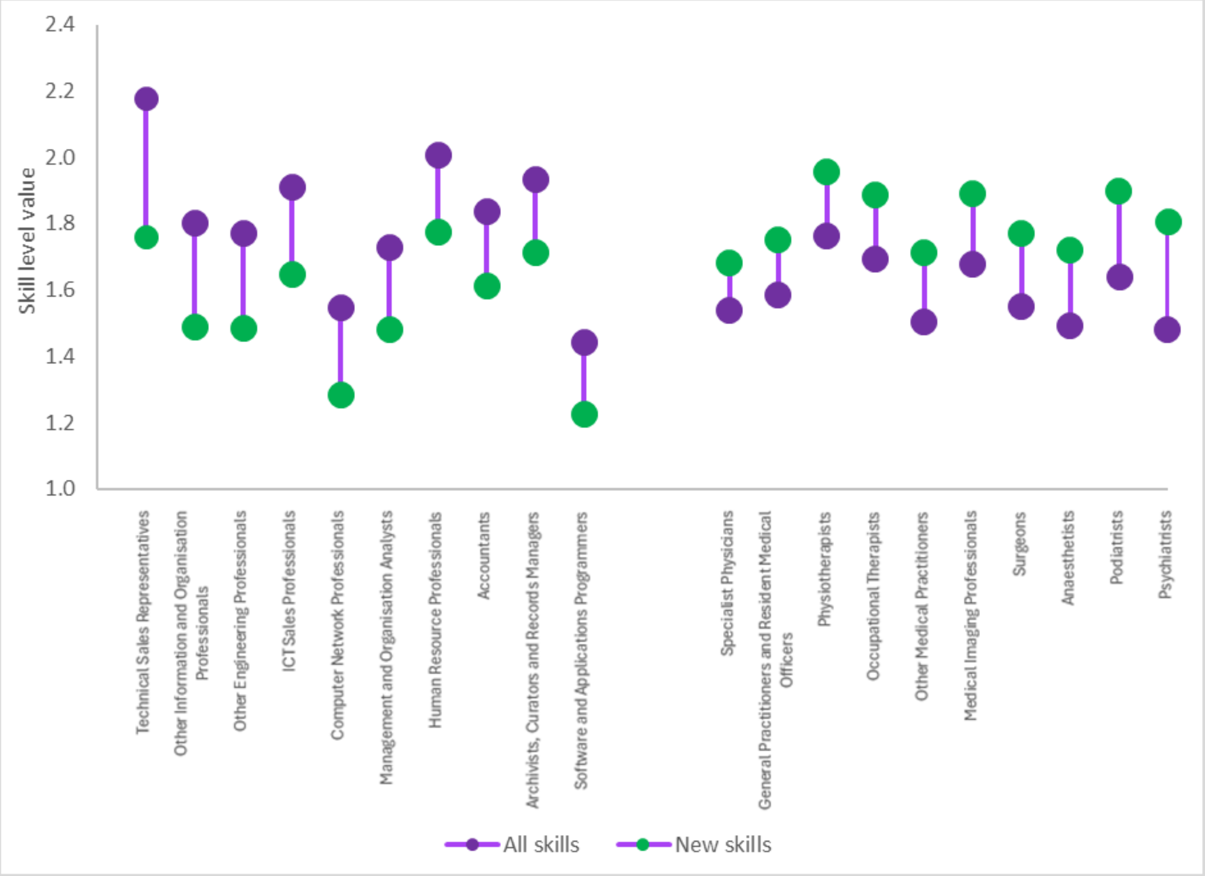

Note: A high skill level value (closer to 1) indicates that the skills listed in online job advertisements for this occupation are heavily concentrated among highly skilled occupations per ANZSCO skill levels. A low skill level value (closer to 5) indicates the opposite.

Source: JSA analysis of Lightcast job advertisement data.

Results relating to changes in the skill content and complexity of occupations produced through this approach are descriptive in nature. Any causal inferences should be made in relation to the specific context of the relevant occupation/s, after consideration of the specific growing (or declining) skills and other sources of evidence as appropriate. The analysis published in Analysis Paper C: Adaptation provides examples of this approach in relation to select occupations including Arts and Media Professionals, Business, Human Resources and Marketing Professionals, Secretaries, and Welfare Support Workers. Discussion of potential explanations for changes in the skill content of jobs beyond these examples is beyond the scope of this technical paper.

Applications and extensions

This approach is applied in Analysis Paper C: Adaptation to explore real-world examples of within-occupation adaptation effects. This includes how these effects can manifest in multiple ways, such as emerging trends, job specialisations and job hybridisation.

The approach in this paper could be applied to exploring the medium-term effects of Gen AI adoption on the skill content of occupations. As firms increasingly progress from tentative adoption, to integration into workflows, to maturity (where AI use is scaled and governed strategically), the adaptation effects of generative AI will become more pronounced.

Observing changes in the skill content of jobs as the Gen AI transition progresses will support labour market and skills system actors to align workforce planning and skills strategies with medium-term investments in technology. Observing changes in the skill complexity of occupations could also provide insight into the skill bias of Gen AI-driven technological change.

The approach in this paper is not specific to the exploration of technological change. It may also provide useful insights in relation to adaptation stemming from other drivers of change, such as demographic shifts, changes in consumer demand, or regulatory change.

Further work extending on this approach could also usefully isolate certain types of skill. For example, excluding common skills and certification skills would enable exploration of how adaptation effects have changed the complexity of specialised skills in an occupation.

Validation

As part of our validation process, we cross-check the average skill level value of occupations in each ANZSCO Skill Level. This cross-check indicates that the approach in this paper produces a set of results which, on average, demonstrates the same distribution of occupations by skill level (i.e. Skill Level 1 occupations are higher skilled than Skill Level 2 occupations and so forth) (Table 3).

Skill level values under our approach tend to converge nearer to the middle of the possible range. In part, this is due to many online job advertisements listing common skills (skills that are prevalent across many different occupations) alongside specialised skills and certification skills. As a result, job advertisements typically do not exclusively list skills that are concentrated in occupations at a single skill level.

Table 3. Average skill level values for occupations broadly align with expectations

Average skill level value of occupations by ANZSCO Skill Level

Skill Level (ANZSCO)	Average Skill Level Value (Online job advertisements)
Skill Level 1	1.80
Skill Level 2	2.15
Skill Level 3	2.57
Skill Level 4	2.63
Skill Level 5	2.80

Source: JSA analysis of Lightcast job advertisement data.

Occupational Mobility and AI Pathways

This technical note documents the methodology used in this Study to measure occupational mobility and combine that mobility with forward-looking measures of exposure to Gen AI. It underpins Analysis Paper D and Analysis Paper F, and forms part of the Study's technical documentation.

The Data on Occupational Mobility (DOM) dataset provides a historical view of the recently observed transitions between occupations in 2022–23. By contrast, automation and augmentation exposures are forward-looking, reflecting how tasks may change as Gen AI diffuses.

The methodology summarised in this paper combines these two perspectives to produce pathway-based measures. These measures show not only how exposed an occupation is, but how that exposure plays out given the destinations workers are most likely to move into.

The approach rests on three complementary indicators of occupational mobility: churn, entropy and diversity.

These indicators are standardised and combined into a Mobility Index. The index is then paired with forward-looking measures of automation and augmentation to construct Automation Pathway Scores and Augmentation Pathway Scores.

The framework classifies occupations by both their exposure to Gen AI and their capacity to adjust.

Data and model foundation

The DOM dataset is derived from the ABS Person-Level Integrated Data Asset (PLIDA) as part of the Skills Tracker project. It records transitions made by individuals between occupations during 2022–23, based on occupation information recorded in personal income tax data. Each transition includes an origin occupation, a destination occupation and the number of workers moving between the two.

For analysis, these data are reshaped into square origin–destination matrices. Rows represent origin occupations, columns represent destinations, and the diagonal captures people remaining in the same occupation. This structure ensures that inflows, outflows and transition probabilities can be measured consistently for all occupations.

Constructing occupational mobility measures

Occupational mobility is measured using churn, entropy and diversity. Each indicator captures a different aspect of occupational dynamism.

- **Churn** reflects turnover intensity: the share of workers entering or leaving an occupation relative to its employment stock. High-churn occupations, such as sales assistants, see large numbers of workers moving in and out each year. Low-churn occupations, such as primary school teachers, are characterised by stability.
- **Entropy** reflects the balance of outflows. It measures whether workers leaving an occupation disperse evenly across multiple destinations or concentrate into one. An occupation whose outflows are spread evenly records high entropy, while one that directs nearly all flows to a single destination records low entropy. High entropy is a sign of greater flexibility.

- **Diversity** reflects the breadth of options. It counts the number of distinct destinations workers move into. An occupation with flows into five destinations has a diversity score of five, while one that directs all flows to a single destination has a score of one.

Churn, entropy and diversity are standardised into z-scores and averaged to create a composite Mobility Index. Positive values indicate higher than average mobility, negative values lower than average mobility. The Mobility Index provides a single measure of occupational resilience to the impacts of labour market disruption and other change. Box 1 provides mathematical write-ups for each measure.

The international literature supports the use of these measures. Turnover measures comparable to churn are common in labour economics and are used to capture job and worker flows. Entropy-style dispersion measures have been applied in studies of occupational reallocation, including by the ILO. Destination counts, equivalent to diversity, appear in comparative work by labour economists affiliated with the IZA. Taken together, these indicators provide a balanced picture of volume, balance and breadth in occupational transitions, though such measures might vary depending on the data.

Box 1 Mathematical expressions – Churn, Entropy, Diversity, Pathways

Below are the mathematical specifications for **churn, entropy, diversity**, and the composite **Mobility Index**.

A. **Churn** is the turnover intensity relative to employment stock:

$Churn\_j = \frac{\left(E_j^{in} + E_j^{out}\right)}{2N_j}$ where:

- E_j^{in} = total inflows to occupation j
- E_j^{out} = total outflows from occupation j
- N_j = employment stock of occupation j

Churn is defined for $N_j > 0$. Small occupations with *limited* employment stocks can display *inflated* churn values due to denominator effects.

B. **Entropy** measures the evenness of outflows from occupation j across destinations:

$Entropy\_j = -\Sigma\_\{i \neq j\}\ p\_ji * \log(p\_ji)$

where:

- $p\_ji = f\_ji / E\_out\_j$, the probability of a transition from occupation j to i
- f_ji = number of transitions from occupation j to occupation i
- $E\_out\_j = \Sigma\_\{i \neq j\}\ f\_ji$

Entropy is set to zero if $E\_out\_j = 0$. The measure captures *balance*, not the *relatedness* of destinations.

C. **Diversity** counts the number of distinct destinations with positive flows:

$Diversity\_j = \Sigma\_\{i \neq j\}\ 1\{ f\_ji > 0 \}$

where:

- $1\{\cdot\}$ is an indicator function equal to 1 if the condition is *true*, 0 otherwise

Optional thresholds can be applied (e.g., $f\_ji \geq \tau$) to reduce the influence of *very small flows*.

D. To standardise the composite **Mobility Index**, each measure is converted to a z-score across occupations:

Z score across occupations (or z_j,m*)* $= (M\_j,m - \mathrm{mean}(M\_m)) / \mathrm{sd}(M\_m)$

- for $m \in$ {Churn, Entropy, Diversity}.

E. The **Mobility Inde**x is then calculated as:

$Mobility\text{-}Index\_j = (z\_j,Churn + z\_j,Entropy + z\_j,Diversity) / 3$

Alternative weightings could be used, but *unweighted averaging* ensures transparency and simplicity, with an equal treatment applied to each partial indicator.

Pathway exposure scores for workers combine own (original occupation) exposure with the exposure of destinations:

$P\hat{}X\_j = \alpha X\_j + (1 - \alpha)\ \Sigma\_\{i \neq j\}\ p\_ji\ X\_i$

where:

- $P\hat{}X\_j$ = pathway score for occupation j, for either automation or augmentation exposure
- X_j = own exposure score for occupation j
- p_ji = transition probability from j to i
- **$\alpha \in [0,1]$ is the weighting parameter, with $\alpha = 0.6$ in this study**

The values of **α** in real terms could differ between occupations and over time, depending on, for example, if the skills stock of a worker might influence the probability of flow from the existing occupation to another one.

This aims to reflect how exposure to the technology that a worker in transition might face is shaped both by the potential exposure of the original occupation and the occupations that workers in those occupations would most commonly move into.

Automation and Augmentation Pathway Scores

Traditional exposure measures assign automation or augmentation scores to occupations in isolation. However, resilience depends not only on the job a worker currently holds, but also on the jobs they are likely to move into. To capture this, Automation Pathway Scores and Augmentation Pathway Scores combine an occupation's own exposure with the exposures of its common destinations.

Pathway scores are a weighted combination of own exposure and destination exposures. A weight of 0.6 is applied to own exposure, with the remainder distributed across destinations according to observed transition probabilities.

This design means that if an occupation is itself relatively stable but commonly leads to highly automatable roles, its Automation Pathway Score will be higher than its own exposure alone suggests. If an occupation is exposed but often leads to resilient destination occupations, its pathway score will be lower.

These measures link forward-looking task exposure with historical mobility. They reflect both the characteristics of the job itself and the characteristics of the jobs workers tend to move into.

Quadrant classification framework

To interpret the implications of the analysis, occupations are positioned on a two-dimensional plot: Mobility Index on the vertical axis and pathway score on the horizontal axis. Thresholds, based on the mean or median of each distribution, are used to divide the space into four quadrants.

- High automation pathway scores and low mobility indicate potential vulnerability: workers face higher substitution potential from technology and fewer or weaker exit paths.
- High automation pathway scores and high mobility also indicate potential vulnerability, but with greater potential ability to adjust than the first quadrant of occupations.
- Low automation pathway scores and high mobility indicate relatively high occupational resilience.
- Low automation pathway scores and low mobility indicate apparent stability, but with limited scope for adaptation if exposure to automation changes.

When augmentation pathway scores are used, the interpretation changes. Occupations with high augmentation scores and high mobility are well placed to benefit from Gen AI. Those with low augmentation and low mobility are the least likely to gain from new tools.

Methodological assumptions

The approach rests on several assumptions.

First, DOM data are historical and focused on a single year of transitions. They capture occupational mobility in 2022–23 and are sensitive to the labour market context. If measured during the COVID years or under other labour market conditions, churn and entropy would have looked different.

Second, AI exposures are forward-looking. They describe potential changes in tasks over time, but are treated as fixed in this analysis.

Third, the weighting parameter influences results. A weight of 0.6 favours own exposure, but other values would shift scores and reclassify some occupations.

Fourth, classification depends on the choice of threshold. Mean and median cut-offs give slightly different quadrant assignments, particularly for marginal cases.

Finally, aggregation would make a difference. Exposures are calculated at the four-digit ANZSCO level, while the DOM data operate at the three-digit level. This aggregation may mask variation within occupational groups.

Limitations

The DOM methodology is diagnostic rather than predictive, and the measures it relies on each carry technical limitations.

- **Churn**, defined as inflows and outflows relative to employment stock, is sensitive to occupation size. Smaller occupations can appear highly dynamic because even modest flows inflate the ratio, while large occupations can appear stable despite substantial movement.
- **Entropy** measures the balance of outflows across destinations, but it does not account for the relatedness or quality of those destinations. Two occupations that are very different in skill content are treated the same as two that are highly similar. Entropy is also less reliable for occupations with very few observed flows.
- **Diversity** simply counts the number of distinct destinations. It treats all destinations equally regardless of size, which means that a single worker moving into a new role can shift the measure in the same way as thousands moving into a dominant exit. This makes diversity particularly sensitive to noise.

Because the Mobility Index is an average of these three indicators, it inherits their sensitivities. Small changes in flows can shift results, especially when measures are based on a single year of data. DOM results are also sensitive to labour market conditions. Flows observed in a year like 2022–23 will differ markedly from flows observed during periods of disruption such as COVID.

These limitations point to the need for robustness testing. Results could be checked across multiple years, bootstrapped to assess confidence intervals, or re-estimated excluding very low-flow transitions. Sensitivity analysis using alternative weightings in the composite index would also help identify where results are stable and where they are most fragile.

Results framework

The DOM analysis highlights relative positioning across occupations.

For example:

- Sales Assistants record high churn and high mobility but relatively low augmentation potential. They are resilient in terms of exits, but not as strongly positioned for productivity gains.
- General Clerks record moderate churn, lower mobility and higher automation exposure, placing them in a more vulnerable position with fewer exits.
- Registered Nurses record low churn and low mobility, but moderate augmentation potential, suggesting stability but also more limited flexibility.

Differences are also seen for various groups of people in the labour market. For example:

- Women are concentrated in augmentable but less mobile occupations.
- Young workers are concentrated in more mobile but less augmentable roles.

- First Nations workers and those in remote areas are represented in both more resilient and more vulnerable quadrants.

However, as noted earlier, results from this type of analysis are highly sensitive to methodology. Adjusting the weighting parameter changes quadrant placement. Choosing mean or median thresholds alters classifications for some occupations. Results are also sensitive to labour market conditions, as churn and entropy reflect the flows occurring under those conditions.

Applications and extensions

The methodology provides a structured way to link forward-looking AI exposures with historical occupational mobility.

Extensions are possible. Exposures could be updated dynamically as technology evolves and task content changes. Transition probabilities could be adjusted to reflect training, credentialing or migration policy interventions. Analyses could be conducted at the four-digit ANZSCO level to reduce aggregation bias. Cohort tags could be embedded into transition matrices.

The approach can also be applied at the industry level. Using a Data on Industry Mobility (DIM) dataset, similar measures of churn, entropy and diversity could be constructed.

Automation and augmentation pathway scores could then be combined with industry-level mobility to assess resilience and exposure across industries as well as occupations.

Validation

Validation was undertaken in stages.

First, comparisons with established labour market patterns confirm the expected behaviour of the measures. For example, entry roles in sales occupations record higher churn and diversity. More stable occupations, such as teachers and nurses, record lower churn and entropy. Pathway scores align with known task structures.

Second, sensitivity tests have been conducted. Varying the weighting parameter and changing the thresholds between mean and median alter the scores and categorisation for marginal cases but preserve the broad structure of the analysis and results.

Finally, given this approach relies on an innovative combination of a range of underlying methods, JSA is continuing to validate the approach, to inform its potential future use in other analyses. Until then, the method and the results should be regarded as preliminary.

Entry-Level Analysis

The entry-level analysis is used to analyse whether demands of new entrants to the labour market are changing. We are particularly interested in exploring whether Gen AI has affected entry-level hiring or changes in demand for junior workers, since Gen AI is best at repetitive knowledge-based tasks, the same tasks many new entrants to the market would be undertaking in their work. Our qualitative analysis also indicates that there has been a shift across various industries in the approach to junior hiring and training as Gen AI has become more mainstream.[1]

This analysis is the first of its kind to be applied to the Australian context, though a similar technique has been used in analyses overseas (Casselman, 2018; Almeida, 2025). To date there is no standard definition or measurement of an entry-level role across either Australia or comparable countries (see **Box 1**).

Data and methodology

To consider whether Gen AI has affected entry-level hiring, we review online job advertisements for positions targeted to new entrants and look for changes in demand over time. We use the Lightcast job advertisement dataset of Australian jobs between 2018-2025, including data up to June 2025. For a comprehensive discussion on the benefits and limitations of this dataset, please refer to **Analysis Paper D**.

We define entry-level workers as any worker who is seeking a first job in a career pathway node. This refers to any new worker in a role, including recent graduates, those with no or limited (2 or fewer years of work experience), and anyone who will need to undergo further training on the job. This also includes those who have recently graduated from secondary or tertiary education, and new or returning workers to the workforce trying to break into new fields, as well as those who have some experience in one field but are moving into the first node of a different career pathway.

To identify a job ad as being entry-level, we searched the text of the job advertisement and flagged any ads containing any entry-level related keywords, including but not limited to words like 'junior', 'assistant', 'graduate', 'apprentice', 'trainee', 'no experience needed', '2 years of experience, 'training provided', and 'career starter'. Job ads were grouped by occupation and industry and the proportion of each measured over time to produce the entry-level related graphs and charts in this study.

---

[1] Case study: Gen AI and entry-level roles in legal, health care, and creative industries

Box 1: Official definitions of entry-level workers

Internationally, entry-level roles are typically not monitored as part of official labour market statistical collections. The ABS discusses 'new entrants' to the labour market in their Participation, Job Search and Mobility Survey (Australian Bureau of Statistics, 2025) as those who have 'entered or re-entered employment without working during the last 12 months (i.e. they did not change jobs during the year).' The measure includes those who started their first job and were never previously employed and those are returning to the workforce after an absence of more than 12 months.

The ABS also uses the Occupation Standard Classification for Australia (OSCA) to group jobs by skill level. The skill levels describe the education, prior experience, and on-the-job training required to competently perform the tasks required of a particular occupation and can be matched to typical degree/certificate requirements (Australian Bureau of Statistics, 2024).

Many organisations also discuss 'graduates' and use this term to refer to recent tertiary education graduates, but even this is not consistent. Nor does it include those who haven't graduated and are entering the labour force for the first time.

In the US, the Bureau of Labor Statistics (BLS) defines entry-level roles as 'the starting level for workers who are new to an occupation'. The BLS also notes that entry-level work for different occupations may require different education, training, or experience at entry, but does not measure any entry-level statistics specifically (U.S. Bureau of Labor Statistics, 2025). In the O*NET occupation taxonomy, another American occupation taxonomy, roles are differentiated by 'job zones' which group jobs with similar training, experience, and educational requirements. Entry-level roles are usually considered as those with the lowest needs and in job zone 1 (O*NET Resource Center, 2025).

Both the EU and the UK also have no official definition of an entry-level role. Instead, both use skill level to define a job’s minimum starting requirements (based on the relevant local occupation classifications, the ISCO-08 and the UK Standard Occupation Classification [SOC] respectively), where skill level defines the training and experience needed to do the job (International Standard Classification of Occupations (ISCO), 2024) (Office for National Statistics (ONS), 2023). The UK does have an entry-level category of education (Entry 1-3) where anything below level 1 is considered base education, but this does not refer to jobs, rather the educational attainment of a worker (Gov.UK, 2025).

Singapore has no definition of an entry-level role but does have a SkillsFuture Skills Framework where career pathways are published, and the first node in each path is considered entry-level for the specific occupation/sector (SkillsFuture Singapore, 2018).

Results and interpretations

The share of entry-level job ads overall decreased during the COVID-19 pandemic and increased after that, reflecting a tight post-pandemic labour market. After 2022, the proportion of entry-level job ads decreased slightly but remained stable through to the end of 2024. Looking at individual industries and occupations, the share of entry-level job ads shows the same dip mid-pandemic (2021) and the subsequent increase in entry-level roles as the economy recovered, though this recovery was more pronounced for certain industries and occupations. Similarly, after 2022 the overall trend has been that the share of entry-level roles has decreased slightly for most occupations and industries. For further discussion, please refer to **Analysis Paper D**.

The results show that entry-level job ads, and by extension, the demand for entry-level workers has not declined since Gen AI has become more widespread (between 2022 and mid-2025). This is a significant finding, given the capability of the technology to compete the same kinds of tasks often allocated to entry-level workers, and the anecdotal evidence we heard in our qualitative research and elsewhere. The caveat to this result, however, is that the effects of Gen AI may not yet be evident in job ad trends or other official statistics. Nor can the extraordinary labour market disruption of the COVID-19 period be understated, with some larger effects potentially drowning out any emerging signals related to the AI transition.

The clear takeaway from this analysis is that so far there is little evidence of Gen AI affecting entry-level demand. However, JSA should continue monitoring these statistics to readily identify if and when changes start to occur. Junior workers finding adequate employment is a cornerstone to a healthy and dynamic labour market, with appropriate career and skilling pathways. Therefore being able to identify early whether the situation for entry-level workers is changing will allow the government and other actors in the jobs ad skills system adequate time to respond with appropriate policy action and strategies.

Exclusions/Amendments

The analysis identifies entry-level roles for almost every occupation. One of the limitations of this current method and analysis is that all occupations are included as possible entry-level roles. For example, C-suite roles are included in our analysis as they also advertise for those looking to first move into the C-suite from other specialties without any prior executive experience. These roles are often geared towards small business looking for an administrator, or not-for-profit (NFP) organisations who are more willing to train candidates on the job provided they have an adequate background in either the particular industry the small business operates in or other relevant experience.

Whether an experienced professional moving into a new field like this (i.e. starting another career path) is considered an entry-level worker to the same extent that a newly graduated student or a worker with no work experience whatsoever is questionable. Future iterations of this work may choose to distinguish between the two and create a two-tier entry-level analysis, one of which focuses on entry-level workers with no experience at all, and one which considers those with some professional background who are choosing to move into a new career pathway.

Validation

At this stage, no comparable analyses on entry-level roles are available in the Australian labour market. The ABS does report new entrants to the labour market in their Participation, Job Search and Mobility Survey, however, this is limited to those who started their current job in the last year but did not have a job at any time over the twelve months prior. It is not limited to junior roles (Australian Bureau of Statistics, 2025).

Other employment statistics which look at recent graduates (a subset of all entry-level workers) demonstrate trends similar to our analysis. For example, unemployment for VET graduates spiked with a shock in 2019 and a subsequent decline until 2022, whereafter it has remained constant with a very slight increase through 2024 (National Centre for Vocational Education Research (NCVER), 2024). The ABS also reported a slight decrease in employment for school leavers between 2023 and 2024 (Australian Bureau of Statistics, 2024), whereas the Graduate Outcomes Survey reported similar levels of employment after graduating from university between 2022-2023 (Quality Indicators for Learning and Teaching, 2025).

Entry-level demand and graduate outcomes overseas

Overseas evidence around entry-level roles is mixed. Further, there are no official labour market statistics released on entry-level workers (See **Box 1**).

In the US, Revelio Labs, a job advertisement data company, claimed that demand for entry-level workers in January 2025 is around 35% less than that of January 2023 (Simon, 2025). In contrast, Indeed, a job advertising engine claimed that the share of entry-level roles, particularly in tech, has remained stable since the post-COVID recovery with only a very slight *increase* in demand (from 17% to 18% of job ads) between 2022-2025. In the EU, the LinkedIn EMEA Labour Market Outlook also shows that hiring for new entrants and junior positions has remained stable during the 2022-2025 period, while hiring has only declined for more senior roles (LinkedIn, 2025).

On the other hand, the 2025 Handshake Internship Index, an American annual survey of university students and recruitment monitoring index, noted a decline of advertised internships by 15% between 2023 and 2025 (Handshake, 2025). The Federal Reserve of New York shows that unemployment for recent graduates spiked in 2020, declined until 2022 and has been increasing very slightly since (The Federal Reserve Bank of New York, 2025). In Singapore, employment for university graduates six months after graduation had slightly decreased between 2023 and 2024 (Ministry of Education, Singapore, 2025), whereas Japan recorded its second highest share ever of recent graduate employment in April 2025 (Jiji , 2025).

There seems to be some consensus that the composition of work has changed for graduates and those undergoing study, even as employment outcomes might vary across sectors and over time (Sigelman, Mamertino, de Zeeuw, Levanon, & Guilford, 2025).

Applications/Extensions

It is still too early to tell whether Gen AI is having an effect on the entry-level workforce, but with regular monitoring, we may be able to identify when (if) it has a material effect on entry-level jobs. In the event that demand changes, early warning will ensure that policymakers can better support vulnerable groups who are looking for work for the first time and to ensure that new entrants to the labour market are able to secure employment.

For future research, the entry-level analysis could also distinguish between levels of work experience (see discussion in the **Exclusions/Amendments** section above) and split the analysis by those who are entering the labour market for the first time entirely and those with some experience who are moving into a different career path for the first time.

How entry-level work has changed over time can also be considered on the "supply" side by considering job profiles data and the stock of skills, experiences and qualifications entry-level workers may have across time and industries.

Rate of skill change

This technical paper relates to analysis published in Analysis Paper E: Skills in *Our Gen AI Transition: Implications for Work and Skills Analysis Papers* focused on skills changes in occupations over time. It extends on the brief overview of the approach included in Box 1 Rate of skill change.

Data

We study the rate of skill change by occupation between 2012 and 2024 using online job advertisement data from Lightcast. This data is scraped by Lightcast from Australian job advertisements posted on publicly available online job boards.

Online job advertisement data is well-suited to measuring the rate of skill change. As Deming and Noray (2020a) observe, 'vacancies directly measure employer demand for specific skills, and vacancy data are sufficiently detailed to measure changing skill demands within occupations over time.' The limitations of online job advertisement data are discussed in detail in the Limitations section below.

Method

We construct a measure of the rate of skill change between 2012 and 2024, replicating the approach of Deming and Noray (2020a) in their earlier work on the US labour market. The start and end years for our analysis were chosen with the intention of capturing changes over the medium term (i.e. across a period of 10 to 15 years). To improve the relevance of the results to the Australian labour market, we map online job advertisements to the Australian and New Zealand Standard Classification of Occupations (ANZSCO) rather the US Standard Occupational Classification system used in Deming and Noray's original analysis.

For both years (that is, 2012 – the earliest complete year of comparable data - and 2024 – the latest complete year), we collate all the skills listed in online job advertisements for each occupation unit group (4-digit ANZSCO occupations). We then calculate the share of job advertisements for the occupation in which each skill appears in the relevant year.

Next, we calculate the absolute value of the difference in shares for each skill between the two years, and then sum these values by occupation to obtain an overall measure of change applying the formula below.

$$Skillchange_o = \sum_{s=1}^{S} \left\{ Abs\left[\left(\frac{Skill_o^s}{JobAds_o} \right)_{2024} - \left(\frac{Skill_o^s}{JobAds_o} \right)_{2012} \right] \right\}$$

As a final step, to account for differences between 2012 and 2024 in the frequency of job advertisements and skills per job advertisement, we multiply the equation above by the ratio of total skills in 2012 to total skills in 2024, for each occupation. This step accounts for compositional changes in the Lightcast data and prevents confusing changes in the frequency of job advertisements with changes in the average skill requirements of any given job advertisement.

Assumptions

That the Lightcast skills taxonomy is a reasonable framework for describing skills in the Australian labour market.

By its nature, this exercise is dependent on the use of a specific taxonomy for classifying skills. This taxonomy is then assumed to be a reasonable framework for providing useful insight into skills in the labour market. While a range of skills taxonomies and frameworks exist in Australia, there is no single nationally recognised or agreed framework or taxonomy that consistently describes the range of skills employers and individuals need across the Australian labour market and education system (Jobs and Skills Australia, 2024).

In the absence of such a nationally agreed taxonomy, our approach uses the skills taxonomy provided by Lightcast. While potentially limited in some respects (e.g. the extent to which it reflects the Australian context), the Lightcast skills taxonomy has advantages as a detailed taxonomy (with over 33,000 unique skills) that is regularly reviewed and updated.

Limitations

Online job advertisements are not representative of all job openings.

Online job advertisements data over- or under-represent certain occupation groups based on their recruitment methods. For instance, a higher proportion of businesses use recruitment websites and job boards when recruiting for Managers (64%) and Professionals (64%) compared with Technicians and Trades Workers (56%), Labourers (48%) and Sales Workers (44%) (National Skills Commission, 2022).

Given the focus of this analysis on within-occupation changes in skill composition, not having, comprehensively representative data on the distribution of job openings across occupation groups is not a critical limitation. However, online job advertisements, resumes and worker profiles are also the key inputs for the development of the Lightcast skill taxonomy. As a result, the Lightcast skill taxonomy may be relatively more detailed and comprehensive in relation to skills more commonly listed for Managers and Professionals and potentially understate the extent of skill change in occupation groups that are underrepresented in online job advertisements and worker profiles.

Some skills may be implicit in job advertisements.

Online job advertisements are not primarily intended as a comprehensive stocktake of the skills required for a job. In some instances, changes in skill requirements for an occupation may not be observable in skills listed in job advertisements. For example, in occupations with mandatory qualifications (or other certifications required under legislative, licensing and regulatory arrangements), skill change may occur through changes in education and training requirements rather than changes to the unique skills demanded in job advertisements (Deming & Noray, 2020a).

Exclusions

We exclude from our analysis any occupations with fewer than 100 online job advertisements in either 2012 or 2024. This is done to improve the robustness of the results, ensuring each in-scope occupation has a minimum viable sample of online job advertisements contributing to the analysis. Applying this threshold leaves us with 245 out of 358 occupation unit groups (4-digit ANZSCO occupations) as in-scope for analysis.

Results

Professionals and managerial occupations tend to exhibit the fastest rate of skill change, with Labourers exhibiting the slowest rate of skill change (Figure 1).

Figure 1. Higher skilled occupation groups experience faster skill change

Average rate of skill change of occupations by occupation major group, 2012-2024

Source: JSA analysis of Lightcast job advertisement data.

Technology-intensive occupations also appear to demonstrate a faster rate of job skill change (Figure 2). This can be seen by correlating the rate of skill change with the percentage of job advertisements – by occupation – requesting at least one skill in the Information Technology category per the Lightcast skills taxonomy.

Figure 2: Technology-intensive occupations tend to exhibit faster skill change

Rate of skill change and percentage of job advertisements requesting Information Technology (IT) skills by occupation

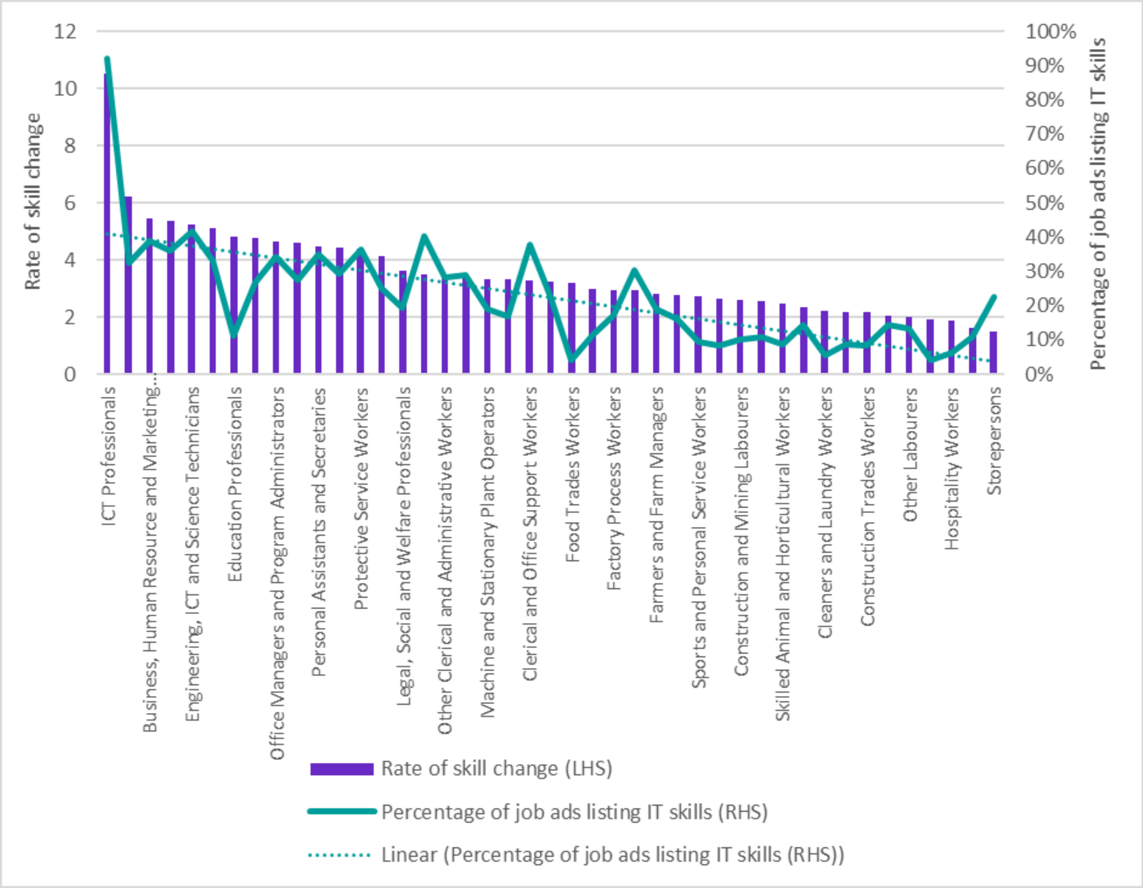

Source: JSA analysis of Lightcast job advertisement data.

The rate of skill change for the full list of in-scope occupation unit groups is available at Table 15 of the JSA AI Study Bespoke Tables.

Applications and extensions

This approach is applied in Analysis Paper E: Skills to explore whether technology-intensive occupations exhibit a faster rate of skill change. The pace of job skill change has important labour market and skills system implications, including in relation to the dynamism of education and training systems and adaptive capacity of firms and workers.

Future work building on this analysis could include:

- explanatory analysis seeking to understand how technology impacts the pace of change
- consideration of a wider range of variables beyond technology-intensity that could influence the pace of job skill change, and
- restricting the analysis to certain categories of skills to isolate or compare the rate at which different types of skills change within occupations of interest.

Validation

As part of our validation process, we compare our results against those produced by Deming and Noray (2020a). Notwithstanding that direct comparison is challenging given the different occupation classifications and time periods used, the results of our analysis are broadly in line with expectations based on Deming and Noray's earlier work. For example, ICT, marketing, engineering and sales occupations are represented at the high end of rates of skill change under both sets of results (Table 1).

Table 1: Similar occupations exhibit the fastest rate of skill change in both sets of results

Occupations with the highest rates of skill change, comparison of Deming and Noray analysis and JSA analysis

Deming and Noray (2020)	Jobs and Skills Australia (2025)
Computer Occupations	ICT Professionals
Advertising, Marketing and Sales Managers	Arts and Media Professionals
Sales Representatives, Services	Business, Human Resource and Marketing Professionals
Operations Specialties Managers	Design, Engineering, Science and Transport Professionals
Life, Physical, and Social Science Technicians	Engineering, ICT and Science Technicians
Electrical and Electronic Equipment Mechanics, Installers, and Repairers	Specialist Managers
Engineers	Education Professionals
Financial Specialists	Sales Representatives and Agents
Business Operations Specialists	Office Managers and Program Administrators
Supervisors of Installation, Maintenance, and Repair Workers	Sales Support Workers

Source: Deming and Noray (2020b), JSA analysis of Lightcast job advertisement data.

Returns to experience and skill change

This technical note provides the regression framework and evidence underlying **Box 8 of the Overarching Paper**, which examines how the rate of occupational skill change shapes the lifecycle profile of returns to experience.

We adapt the Deming and Noray earnings–experience model to the Australian context using HILDA microdata and link it to a measure of occupational skill change derived from online job ads, calculated in the companion Technical Note on the Rate of Skill Change.

The purpose of this paper is to explain the approach used to produce the analysis presented in the Box 8 figure and Analysis Paper E.

Data

The analysis combines the Household, Income and Labour Dynamics in Australia (HILDA) survey with an independent, occupation-level measure of the rate of skill change. HILDA waves from 2012 to 2023 are used to construct log hourly wages, labour market experience and tenure, and to identify occupations using ANZSCO codes.

The rate of skill change is computed from Lightcast online job ads between 2012 and 2024, following the approach set out in the Technical Note on the Rate of Skill Change. This replicates Deming and Noray's skill-turnover metric and maps vacancy skills to ANZSCO 4-digit occupations with a compositional adjustment for the number of skills per advertisement.

Occupations are assigned to quintiles of skill change for the primary specification, with decile groupings used as a robustness check. The occupation-level index is then merged onto HILDA data using the corresponding ANZSCO code to assign each individual-year observation to a skill-change group.

Method

We estimate wage–experience profiles with person fixed effects and year fixed effects, allowing the slope and curve of experience to vary by the rate of skill change.

In our model, w_it denotes log hourly wages for individual i in year t.

Experience (Exp_it) and its square enter the model, interacting with indicators for the occupational skill-change group g.

Model: $\ln \boldsymbol{w_{it}} = \boldsymbol{\alpha_i} + \boldsymbol{\lambda_t} + \boldsymbol{\beta}_1 \cdot \boldsymbol{Exp_{it}} + \boldsymbol{\beta}_2 \cdot \boldsymbol{Exp_{it}}^2 + \boldsymbol{\Sigma_g}\ [\boldsymbol{\gamma}_1\boldsymbol{g} \cdot \boldsymbol{Exp_{it}} + \boldsymbol{\gamma}_2\boldsymbol{g} \cdot \boldsymbol{Exp_{it}}^2] \cdot \mathbf{1}\{\boldsymbol{SkillChange} \in \boldsymbol{g}\} + \boldsymbol{X}'_{\boldsymbol{it}}\boldsymbol{\theta} + \boldsymbol{\varepsilon_{it}}$

In the model, α_i are person fixed effects and λ_t are year fixed effects.

The vector X_{it} includes time-varying controls used in sensitivity checks.

Group-specific wage–experience profiles are obtained by evaluating the fitted quadratic terms at each g and solving for the experience level at which the derivative equals zero to identify the peak of returns.

Assumptions

Our method involved the following assumptions:

- The occupation-level rate of skill change derived from job ads text is taken as a valid proxy for the pace at which skill requirements evolve in practice.
- Experience is assumed to follow a smooth quadratic relationship with wages within each skill-change group.
- Occupation coding in HILDA is assumed sufficiently accurate to support merging with the job-ad index.
- The group assignment is treated as occupation-specific in a given year, acknowledging that some workers transition between occupations over the panel.

Limitations

Our method had the following limitations:

- Online job ads under-represent some occupations and may omit skills that are implicit or embedded in qualification requirements.
- Firm-level wage-setting, bargaining arrangements and institutional features are not directly modelled.
- Classification and timing misalignment when merging the vacancy-based index to HILDA data may introduce some measurement error, though sensitivity checks indicate stable qualitative patterns.
- Results summarise occupation-average dynamics and do not isolate individual reskilling events.

Exclusions

Occupations with fewer than 100 online job ads in either 2012 or 2024 are excluded from the skill-change computation to ensure minimum coverage.

Additionally, HILDA observations not able to be mapped to in-scope occupations are excluded.

Extreme outliers in wages and tenure are also trimmed from the distribution to improve the stability of the results from the model.

Comparison with the Deming and Noray approach

The starting point of both studies is the same conceptual framework: wages over the career reflect a race between on-the-job learning and the obsolescence of older skills.

Deming and Noray (2020a) modelled this as a function where earnings rise with experience at a diminishing rate, but where the slope and curvature of the curve depend on the pace at which tasks evolve. When task change is slow, experience compounds more fully; when tasks change quickly, the curve flattens.

Our application adopts this framework for Australia, linking HILDA microdata to an index of skill change built from job ads data.

Methodology

Overall, the specification is consistent between our approach and Deming and Noray:

$$\ln w_{it} = \alpha_i + \lambda_t + \beta_1 \cdot Exp_{it} + \beta_2 \cdot {Exp_{it}}^2 + \Sigma_g \left[\gamma_1 g \cdot Exp_{it} + \gamma_2 g \cdot {Exp_{it}}^2\right] \cdot 1\{SC \in g\} + \varepsilon_{it}$$

- **Alignment**: both treat earnings as a quadratic in experience; both allow the slope and curve to vary with the rate of skill change; both derive a skill-change index from job ads by summing absolute changes across time.
- **Differences**: we use HILDA panel with person and year fixed effects, Deming and Noray used cross-sectional data; we group data into quintiles/deciles, they used continuous interactions; we use ANZSCO with compositional adjustment, they used SOC and Burning Glass.

Data

Both approaches rely on job ads data to measure skill turnover within occupations. Deming and Noray's data cover 2007–2019 U.S. job ads, while our data cover 2012–2024 Australian job ads from Lightcast. Our returns are estimated using HILDA microdata, which provides longitudinal structure absent in their cross-sectional analysis.

Results

In both studies, high skill-change contexts show later peaks in experience returns, while low-change contexts peak earlier.

Our quintile results emphasise a potential 'sweet spot' in the high group, where experience is rewarded longer without being eroded by very high change.

Deming and Noray emphasised the flattening of returns in the most dynamic fields, consistent with faster obsolescence.

Results

The estimated wage–experience profiles differ systematically across skill-change groups. In occupations with the lowest rates of skill change, returns to experience rise steeply in early career and reach an earlier peak before flattening.

In occupations with higher rates of skill change, returns accrue over a longer horizon and peak later, consistent with continuous learning against a moving skills frontier. At the very top of the distribution, peaks still arrive later but can be modestly lower than those for the high group, suggesting that extreme skill turnover may erode mature returns unless reskilling keeps pace.

The figure reproduced in Box 8 of the overarching report for the Study summarises these findings by reporting, for each quintile of skill change, the years of tenure at the wage peak, the level of peak returns, and the years of experience at peak returns.

Decile-based robustness checks confirm the gradient observed with quintiles and indicate some flattening in the top decile, while residual diagnostics from the HILDA model show centred residuals with no systematic curvature, supporting the adequacy of the specification.

Table 1. Regression estimates and tests by skill-change group (person FE; year FE; clustered at person level)

Skill-change group	β_1 (Exp)	β_2 (Exp^2)	$\gamma_1 g$ (Exp×group)	$\gamma_2 g$ (Exp^2×group)	Joint test (p-value)
Q1: Very low	pos., sig.	neg., sig.	ref.	ref.	<0.01

Q2: Low	pos., sig.	neg., sig.	sig.	sig.	<0.01
Q3: Medium	pos., sig.	neg., sig.	sig.	sig.	<0.01
Q4: High	pos., sig.	neg., sig.	sig.	sig.	<0.01
Q5: Very high	pos., sig.	neg., sig.	sig.	sig.	<0.01

Notes: All specifications include person fixed effects and year fixed effects. Coefficients are significant at the 1% level unless otherwise indicated.

Why we presented the data in the Overarching Paper using bars, not parabolas

The regressions yield fitted parabolas for each group. These are useful technically but harder to compare for a broader audience.

For Box 8 we summarised each parabola into three scalars – years of tenure at peak wage, magnitude of the peak, and years of experience at peak returns – and presented them as bars (Table 2 for interpretation).

This makes the message clearer for a broader audience, while the full parabolas remain available on request.

Quintiles, deciles, and the choice of grouping

We tested both quintile and decile groupings. Quintiles provide stability and clear monotone patterns. Deciles confirm the gradient but show volatility at the tails. The top decile peaks late but with somewhat lower returns, reflecting obsolescence and sampling noise. Quintiles were used in Box 8 for clarity, with deciles included in annex materials.

Table 2. Box 8 summary metrics by quintile (experience–wage profile)

Skill-change quintile	Years of tenure at peak wage	Peak returns (relative)	Years of experience at peak returns
Q1: Very low	≈ earlier peak	higher peak level	≈ mid-teens
Q2: Low	earlier	high	mid- to late-teens
Q3: Medium	intermediate	intermediate	late-teens
Q4: High	later	high but later	around two decades
Q5: Very high	latest	slightly lower peak	around two decades

Notes: Relative levels and timing summarise fitted profiles underlying the Box 8 bars. Exact values are drawn from the HILDA estimation used to produce the figure.

Alternative approach that could be used: continuous specification

An alternative model treats the rate of skill change as a continuous regressor, augments the specification with detailed occupation identifiers, and estimates how the slope and curvature of experience vary with the continuous index rather than with grouped indicators. We include a comparison of the original and alternative approaches below as validation of our original methodology.

The alternative approach reduces discretisation and captures within-group heterogeneity. The estimated profiles are consistent with the grouped model: as the skill-change index rises, wage–experience curves peak later and the curvature flattens at longer tenures, while at the extreme upper tail peak levels can be slightly lower.

The continuous specification therefore shows a similar pattern to the narrative in Box 8 and provides a complementary diagnostic for future extensions and comparisons.

A brief comparison of the continuous and grouped approaches

The grouped specification (as used in Box 8) assigns occupations to quintiles or deciles based on their rate of skill change and estimates separate wage–experience profiles for each group.

The continuous specification instead uses the raw skill-change index as a regressor, interacted with experience and its square, while controlling for detailed occupation identifiers.

Both approaches yield consistent substantive findings but differ in their level of granularity. The grouped model is easier to communicate and visualise, while the continuous model better captures heterogeneity within groups and avoids potential arbitrariness in group cut-points.

Regression outputs from the continuous specification show that the interaction terms between the continuous skill-change index and experience are statistically significant and align with the expected signs: positive for the linear term and negative for the squared term at higher values. This implies that as skill change increases, returns to experience rise for longer but eventually flatten and can decline. These results validate and extend the grouped model findings, providing confidence that the Box 8 patterns are not artefacts of the grouping choice.

Qualitative Research

In addition to quantitative research, this Study used qualitative research to explore the views, experiences, and perceptions of participants regarding their use and adoption of Gen AI within the Australian labour market and skills system.[1] The qualitative research was designed to capture nuanced perspectives on how individuals and firms interact with Gen AI technologies, examine the perceived impacts of Gen AI developments, and identify opportunities and challenges for successful integration into the labour market.

The primary objective of the qualitative research was to explore participants' views and experiences regarding their use and adoption of Gen AI within the Australian labour market context. The qualitative research was structured around three core investigative areas that guided both data collection and analysis:

- Understanding interaction patterns: The qualitative research aimed to examine how people interact with Gen AI in labour market contexts, including their lived experiences, evolving perceptions, and attitudes toward using Gen AI for work-related activities, and study and training purposes.
- Exploring perceived impacts: The qualitative research sought to investigate participants' perceptions of how Gen AI developments affect the labour market and skills system, with particular attention to both benefits and challenges experienced by different groups including employees, employers, and students. This exploration included understanding different impacts across various demographic groups and different industries and firms.
- Identifying opportunities and challenges: The qualitative research aimed to uncover opportunities and practical challenges related to Gen AI adoption, including examining support mechanisms for transitioning to a Gen AI-integrated labour market, clarifying roles and responsibilities in facilitating this transition, and identifying cohort-specific support needs with careful consideration of intersectionality.

This qualitative technical report provides a detailed examination of the qualitative research methods, including data collection approaches, analytical frameworks, and quality assurance measures used throughout the Study.

Qualitative data collection

The qualitative research used purposive sampling as the primary recruitment strategy, which is widely recognised as appropriate for qualitative research seeking to gather rich, detailed information from participants with specific characteristics or experiences relevant to the research objectives. Purposive sampling enables researchers to deliberately select participants who can provide insights into the phenomena under investigation, ensuring that the sample includes individuals with relevant knowledge and experience.

For focus groups and some in-depth interviews, the recruitment process was guided by a detailed screening instrument developed by ORIMA Research in collaboration with experienced

---

[1] The Gen AI team would like to acknowledge the contribution to this qualitative research (including technical report) from commissioned research undertaken by ORIMA Research. As part of the Study, JSA commissioned ORIMA Research to conduct qualitative research, particularly in-depth interviews and focus groups with individuals who are part of the Australian labour market and skills system, exploring their experiences and perceptions regarding the use of Gen AI in these contexts.

qualitative researchers in the Gen AI Study team. This screener ensured that participants met specific criteria related to their awareness of, experience with, or impact from Gen AI technologies within labour market contexts. Local specialist qualitative research recruiters were engaged to facilitate participant identification and recruitment, leveraging their expertise in accessing diverse participant pools across different geographic regions and demographic groups. For consultations and some interviews, the Gen AI Study team used snowball sampling and direct contact with potential research participants.

The sampling approach recognised that participants' identities were multifaceted, with overlapping characteristics that influenced their labour market experiences. For example, the Study included participants who simultaneously belonged to multiple categories, such as long-term unemployed individuals considering entrepreneurship in creative industries, or culturally and linguistically diverse participants over 55 years of age working in health care sectors. This intersectional approach to sampling ensured that the research captured the complexity of participants' lived experiences within the labour market system.

The research involved over 150 participants including consultations, roundtables, in-depth interviews, and focus groups.

ORIMA Research's conducted qualitative research included 59 participants across two rounds of data collection (**Error! Reference source not found.**). The decision to conduct data collection in two distinct rounds allowed for iterative refinement of research questions and analytical focus based on emerging insights from initial data collection:

- Round 1 – conducted from 10-12 December 2024, with a total of 36 participants across three face-to-face focus groups, one online focus group and three online in-depth interviews. Participants were grouped by cohort and had at least some awareness of Gen AI prior to the research.
- Round 2 – conducted from 6-13 February 2025, with a total of 23 participants across two online focus groups and five in-depth interviews. Participants were grouped by industry and had used Gen AI and/or had their work impacted by it.

Data collection methods

The qualitative research used multiple data collection methods to ensure comprehensive exploration of participant experiences and perspectives. This multi-method approach enhanced the validity and reliability of findings through triangulation of data sources.

In-depth interviews

Individual in-depth interviews were conducted with selected participants to gather detailed, personal narratives about their experiences with Gen AI technologies. These interviews used a semi-structured format that allowed for both systematic coverage of key topics and flexibility to explore unexpected themes that emerged during conversations. This method was designed to draw rich, contextual information about participants' experiences while maintaining consistency across different interview sessions.

Interview sessions were conducted both online and in-person, depending on participant preferences and logistical considerations.

Focus groups

Focus group discussions were used to capture group dynamics and explore how participants' views were influenced by interaction with others sharing similar experiences. The focus groups were particularly valuable for understanding collective perspectives on Gen AI adoption and for identifying areas of consensus and divergence among participants with similar backgrounds or experiences.

Focus groups were organised by cohort in Round 1 and by industry in Round 2, ensuring that participants shared relevant commonalities that facilitated meaningful discussion.

The focus group discussions were moderated using structured discussion guides that ensured consistent coverage of research topics while allowing flexibility to explore emerging themes. Moderators used techniques to encourage balanced participation, manage group dynamics, and maintain focus on research objectives throughout the sessions.

Case study development

The qualitative research incorporated targeted case study development, involving participants specifically selected to address critical evidence gaps and test hypotheses that emerged during early stakeholder consultations. These case studies used a hybrid design approach that combined multiple data collection methods and analytical perspectives to provide a comprehensive understanding of how Gen AI is and could be reshaping work, skills, and institutional arrangements.

Each case study was designed to explore specific aspects of Gen AI adoption within particular contexts, examining both current impacts and potential future developments. The case studies used multiple data sources, including participant interviews, document analysis, and literature review where appropriate. This multi-source approach enabled triangulation of findings and provided rich, contextual understanding of Gen AI adoption and integration processes.

Additional data sources

The Study incorporated several additional data collection methods to enhance the comprehensiveness of findings:

- 'Have Your Say on Gen AI' Consultation Hub: This platform facilitated targeted input from technical experts, clinical professionals, industry representatives, and individual stakeholders. The consultation hub received 12 responses that provided practical insights, optimisation suggestions, and identified concerns about adoption pace and exclusion risks.
- Government submission reviews: The research team analysed over 200 formal submissions from relevant government processes, including the Select Committee on Adopting AI, the Inquiry into the Digital Transformation of Workplaces, and the Inquiry into the use of generative artificial intelligence in the Australian education system. This analysis focused on submissions covering ground relevant to the study objectives and noted the rapidly evolving evidence base in this domain.
- Stakeholder consultations: JSA conducted roundtable consultations with various labour market actors and subject matter experts, providing additional perspectives on institutional and policy dimensions of Gen AI adoption.

Data collection procedures

Ethical considerations and consent processes

Prior to data collection (particularly in-depth interviews and focus groups), the Study team sought ethics review and approval from the Department of Employment and Workplace Relations' internal Ethics Review Panel. The approval from the Panel received on 29 November 2024.

Recording and transcriptions

All interview and focus group sessions were audio-recorded with participant consent to ensure accurate capture of discussion content.

Transcription processes used both automated and manual approaches to ensure accuracy and efficiency. Initial transcriptions used automated transcription software (e.g., Microsoft Copilot), followed by manual review and correction to ensure accuracy of content, speaker identification, and contextual nuances.

Data management

The Study implemented comprehensive data management procedures to ensure participant confidentiality and data security throughout the research process. All digital files were stored on secure, encrypted systems with appropriate access controls limited to authorised research team members.

Analytical framework and approach

Abductive thematic analysis

The study used an abductive thematic analysis approach, which represents an analytical framework that integrates elements of both inductive and deductive reasoning. This approach was particularly appropriate for the Study's objectives because it enabled systematic examination of data while remaining open to unexpected patterns and insights that emerged from participant experiences. This abductive approach integrates the elements of both inductive and deductive reasoning (Saunders, Lewis, & Thornhill, 2019). The approach involves examining data iteratively to identify themes and patterns, test theoretical assumptions, and refine insights through continuous engagement with the emerging themes (Saunders & Lewis, 2018).

In the context of this Study, abductive analysis enabled the research team to identify themes that not only categorised participant experiences but also suggested underlying patterns and relationships that could explain the complexities of Gen AI adoption within labour market contexts.

The abductive approach was iterative and reflexive, involving movement between empirical data and theoretical frameworks as understanding developed throughout the analysis process. This iterative process enabled the research team to refine interpretations as new information emerged and to accommodate complexities and contradictions within the data.

Qualitative coding and theme development process

The analytical process followed established best practices for rigorous thematic analysis while maintaining the flexibility required for abductive reasoning. The analysis was conducted manually by the same team of specialist consultants who had conducted the data collection, ensuring continuity of understanding and appreciation of contextual factors that influenced participant responses.

- Initial coding phase: The analysis began with systematic coding of all transcripts and field notes. This process involved careful reading and re-reading of all materials to develop familiarity with the data and identify initial patterns of interest. Codes were developed both inductively from the data content and deductively through consideration of existing theoretical frameworks relevant to technology adoption and labour market dynamics.
- Code organisation and refinement: Following initial coding, the research team engaged in collaborative sessions to review, organise, and refine the coding framework. This process involved examining relationships between codes, identifying overlapping concepts, and developing hierarchical structures that reflected the complexity of participant experiences.
- Theme construction: The transition from codes to themes involved identifying overarching patterns and explanatory frameworks that could account for multiple coded segments across the dataset. Themes were developed to capture both semantic content (explicit meanings within the data) and latent content (underlying assumptions and conceptual frameworks). This process required careful consideration of how emerging themes related to existing theoretical knowledge while remaining grounded in participant experiences.

Use of NVivo and Gen AI tools

The research team used NVivo and Gen AI tools, particularly Microsoft Copilot, as supportive resources during the data analysis process. For instance, NVivo was used to facilitate systematic coding of qualitative data such as transcripts and submissions. Copilot was used to assist with tasks such as summarising public submissions, extracting key points from transcripts, and condensing researcher notes into more manageable forms. Copilot was also applied to help identify potential themes and patterns across the qualitative data, providing an initial layer of thematic organisation.

Importantly, the role of Gen AI was supportive rather than central, with human researchers maintaining full oversight of the process. All outputs generated by Copilot were carefully reviewed, validated, and refined by the research team to ensure accuracy, depth, and contextual sensitivity in the final analysis.

Quality assurance in analysis

The Study implemented multiple strategies to ensure analytical rigour and enhance the trustworthiness of findings:

- Collaborative analysis sessions: Regular research team meetings facilitated the collective review of coding decisions and theme development. These meetings enabled the identification of alternative interpretations and ensured that analytical decisions were well-supported by evidence.
- Validation using technology: A validation process also used Gen AI tool to review session notes, summaries, and transcripts to identify potential gaps in emerging themes. This technological validation complemented human interpretation while ensuring comprehensive coverage of the dataset.
- Peer and participant review: Where feasible, preliminary findings were shared with participants (e.g., case studies) and subject matter experts to verify that interpretations accurately reflected their experiences and perspectives. This validation process helped ensure that analytical interpretations remained grounded in participant realities rather than researcher assumptions.

Presentation of qualitative findings

Qualitative research provides valuable insights into the breadth and depth of participant views on a topic. However, it does not quantify the size and prevalence of these views, making it impossible to measure the exact number of participants holding a particular view on individual issues.

Where relevant, this Study uses specific terms to qualitatively indicate and approximate the size of the qualitative research participants who held particular views:

Most – refers to findings that relate to more than three quarters of the research participants.

Many – refers to findings that relate to more than half but less than three quarters of the research participants.

Some – refers to findings that relate to around a quarter but less than half of the research participants.

A few – refers to findings that relate to less than a quarter of research participants.

Throughout our Study, the term 'participant' refers to individuals who took part in focus groups, interviews, and consultations (including roundtables). When used without specifying a particular cohort, it refers to all participants collectively, as findings were often consistent across groups or demographics.

The most common qualitative findings are reported, except in certain situations where only a few participants raised particular issues. These findings are nonetheless considered important and potentially have wide-ranging implications and/or applications.

Integration of multiple data sources

The analytical approach systematically integrated findings from various data collection methods and sources. This integration process involved:

- Cross-method triangulation: Findings from in-depth interviews, focus groups, submissions, and consultation processes were systematically compared to identify areas of concordance and thematic variation. This comparison enabled identification of themes that were consistent across different data collection contexts as well as insights that were specific to particular methods or participant groups.
- Stakeholder perspective integration: The analysis incorporated perspectives from multiple stakeholder groups, including individual participants, organisational representatives, and subject matter experts. This multi-stakeholder approach enabled development of themes that reflected the complexity of Gen AI adoption from different vantage points within the labour market system.

Participant Characteristics and Representation

Demographic Composition

The Study included participants from diverse demographic backgrounds to ensure representation of different perspectives and experiences within the Australian labour market. The demographic profile encompassed variation across age groups, cultural and linguistic backgrounds, employment status, industry sectors, and skill levels.

Industry and occupational representation

The Study prioritised representation from specific industries based on areas of particular interest and relevance to current Gen AI adoption patterns. Participants included representatives from sectors such as health care, education, legal services, and creative industries.

While this sectoral focus enabled in-depth exploration of Gen AI impacts within specific professional contexts, it also represents a limitation in terms of broader generalisability to lower-skilled occupations and other industry sectors. The research design acknowledged this limitation while maximising insights within the prioritised areas of investigation.

Experience and awareness levels

Participants were selected based on having at least some awareness of Gen AI technologies, with Round 2 interviews and focus groups specifically targeting individuals who had direct experience using Gen AI or whose work had been directly impacted by these technologies. This sampling approach ensured that participants could provide informed perspectives on their experiences while acknowledging that the sample was not representative of the broader population's awareness levels.

Limitations and methodological considerations

Sample representativeness and size

The Study design included several important limitations that must be considered when interpreting findings. The purposive sampling approach, while appropriate for the research objectives, resulted in a sample that was intentionally skewed toward participants with existing Gen AI awareness and experience. This selective approach means that the findings may not represent the experiences of the broader Australian population, particularly those with limited technology awareness or access.

The prioritisation of specific industries and higher-skilled occupations further limits the generalisability of findings to other sectors and skill levels within the labour market. While this focus enabled detailed exploration within targeted areas, it also means that important perspectives from other segments of the workforce may be underrepresented.

Future research with larger samples from individual audience groups could potentially expand on these insights and uncover additional themes that may not have emerged in the current Study. The current Study provides a strong foundation for understanding Gen AI adoption experiences while acknowledging areas where additional research would strengthen the evidence base.

Integration with broader research

Connection to quantitative components

The qualitative research component was designed to complement and enhance quantitative analyses conducted as part of the broader Study. The qualitative findings provide contextual depth and experiential understanding that enriches interpretation of quantitative patterns and relationships identified through other analytical approaches.

The integration of qualitative and quantitative approaches enables a more comprehensive understanding of Gen AI adoption processes, combining statistical patterns with rich descriptions of lived experiences and meaning-making processes. This mixed-methods approach strengthens the overall robustness and applicability of study findings.

Table 1: ORIMA's conducted in-depth interviews and focus groups

	Online (National)	Melbourne (VIC)	Total
Online focus group (OFG) and face-to-face focus group (FG); Online in-depth interview (OIDI)			
Round 1			
Small business operators[2]	1 x OFG n=8	-	1 x OFG n=8
Young persons aged 18-24 years	1 x OFG n=9	-	1 x OFG n=9
Mature-aged persons aged 55 years and over	-	1 x FG n=9	1 x FG n=9
Long-term unemployed persons[3]	3 x OIDI n=3	-	3 x OIDI n=3
Culturally and Linguistically Diverse (CALD) persons[4]	1 x OFG n=7	-	1 x OFG n=7
Round 2			
Creative industry professionals	1 x OFG n=9	-	1 x OFG n=9
Health care industry professionals	1 x OFG n=9	-	1 x OFG n=9
Legal industry professionals (Skill Levels 1-5)	5 x IDI n=5	-	5 x IDI n=5
TOTAL			5 x OFG, 1 x FG, 8 x OIDI n=59 participants

---

[2] Defined as those who managed or owned a business with 2-15 employees.

[3] Defined as those who had received job seeker income support for at least two years.

[4] Defined as those who spoke a language other than English at home and were born overseas.

References

ABS. (2013). Australian and New Zealand Standard Classification of Occupations, Version 1.3. *Australian and New Zealand Standard Classification of Occupations, Version 1.3*.

ABS. (2021). Census of Population and Housing, 2021 (TableBuilder).

Almeida, L. (2025, June 30). *Number of new UK entry-level jobs has dived since ChatGPT launch – research*. Retrieved September 10, 2025, from The Guardian: https://www.theguardian.com/business/2025/jun/30/uk-entry-level-jobs-chatgpt-launch-adzuna

Australian Bureau of Statistics. (2024, May). *Education and Work, Australia*. Retrieved from Australian Bureau of Statistics: https://www.abs.gov.au/statistics/people/education/education-and-work-australia/latest-release

Australian Bureau of Statistics. (2024). *OSCA - Occupation Standard Classification for Australia.* Retrieved from Australian Bureau of Statistics: https://www.abs.gov.au/statistics/classifications/osca-occupation-standard-classification-australia/2024-version-1-0

Australian Bureau of Statistics. (2025, February). *Job Mobility*. Retrieved from Australian Bureau of Statistics: https://www.abs.gov.au/statistics/labour/jobs/job-mobility/latest-release

Bernard, B. (2025, July 25). *Experience requirements have tighteneed amid the tech hiring freeze*. Retrieved from Indeed Hiring Lab: https://www.hiringlab.org/2025/07/30/experience-requirements-have-tightened-amid-the-tech-hiring-freeze/

Casselman, B. (2018, January 2018). *As Labor Pool Shrinks, Prison Time Is Less of a Hiring Hurdle*. Retrieved September 10, 2025, from The New York Times: https://www.nytimes.com/2018/01/13/business/economy/labor-market-inmates.html

Daron Acemoglu, P. R. (2019). Automation and New Tasks: How Technology Displaces and Reinstates Labor. *Journal of Economic Perspectives*.

Deming, D. J., & Noray, K. (2020a). Earnings dynamics, changing job skills, and STEM careers. *The Quarterly Journal of Economics*, 1965-2005.

Deming, D. J., & Noray, K. (2020b). *Online appendix for "Earnings dynamics, changing job skills, and STEM careers".* Retrieved from https://static1.squarespace.com/static/60832ecef615231cedd30911/t/60871ffa59c9730432cc387d/1619468283035/DemingNoray_STEM_Appendix.pdf

Felten, E., Raj, M., & Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal, 42*(12), 2195-2217. Retrieved from https://onlinelibrary.wiley.com/doi/full/10.1002/smj.3286

Felten, E., Raj, M., & Seamans, R. (2023). *Occupational Heterogeneity in Exposure to Generative AI.*

Gmyrek, Berg, Kaminski et al. (2025). *Generative AI and Jobs: A Refined Global Index of Occupational Exposure.* International Labour Organisation.

Gmyrek, P., Berg, J., & Bescond, D. (2023). *Generative AI and jobs: A global analysis of potential effects on job quantity and quality.* ILO Working Paper 96: International Labour Organization.

Gov.UK. (2025). *What qualification levels mean*. Retrieved from Gov.UK: https://www.gov.uk/what-different-qualification-levels-mean/list-of-qualification-levels

Handshake. (2025). *Internships Index 2025.* Retrieved September 10, 2025, from https://joinhandshake.com/themes/handshake/dist/assets/downloads/network-trends/Handshake-Internships-Index-2025.pdf

International Standard Classification of Occupations (ISCO). (2024). Retrieved from International Labour Organization (ILO): https://ilostat.ilo.org/methods/concepts-and-definitions/classification-occupation/

Jiji . (2025, May 23). *98 Pct of New Grads in Japan Land Jobs*. Retrieved from Jiji (English): https://sp.m.jiji.com/english/show/40261

Jobs and Skills Australia. (2024). *National Skills Taxonomy Discussion Paper.* Retrieved from https://www.jobsandskills.gov.au/consultations/national-skills-taxonomy-discussion-paper

LinkedIn. (2025). *EMEA Labour Market Outlook - July 2025.* Retrieved September 9, 2025, from https://economicgraph.linkedin.com/content/dam/me/business/en-us/talent-solutions/resources/pdfs/emea-labour-market-outlook-july-2025.pdf

Ministry of Education, Singapore. (2025, February 20). *Graduate Employment Survey*. Retrieved from Ministry of Education, Singapore: https://www.moe.gov.sg/-/media/files/post-secondary/ges-2024/joint-web-publication-ges-2024-5-aus.pdf

National Centre for Vocational Education Research (NCVER). (2024, December 18). *New student outcomes data out now*. Retrieved from National Centre for Vocational Education Research (NCVER): https://www.ncver.edu.au/news-and-events/media-releases/new-student-outcomes-data-out-now

National Skills Commission. (2022). *Recruitment Insights Report – June 2022.* Retrieved from https://www.jobsandskills.gov.au/sites/default/files/2022-07/Recruitment%20Insights%20Report%20%20-%20June%202022_0.pdf

O*NET Resource Center. (2025). *Job Zone Reference*. Retrieved from O*NET Resource Center: https://www.onetcenter.org/dictionary/30.0/excel/job_zone_reference.html

Office for National Statistics (ONS). (2023, January 19). *SOC 2020*. Retrieved from Office for National Statistics: https://www.ons.gov.uk/methodology/classificationsandstandards/standardoccupationalclassificationsoc/soc2020

Quality Indicators for Learning and Teaching. (2025). *Graduate Outcomes Survey*. Retrieved September 10, 2025, from Quality Indicators for Learning and Teaching (QILT): https://qilt.edu.au/surveys/Data-Visualisation/gos

Saunders, M. N., & Lewis, P. (2018). *Doing research in business and management* (2 ed.). Harlow: Pearson.

Saunders, M. N., Lewis, P., & Thornhill, A. (2019). *Research methods for business students* (8 ed.). New York: Pearson.

Sigelman, M., Mamertino, M., de Zeeuw, M., Levanon, G., & Guilford, G. (2025). *No Country for Young Grads.* The Burning Glass Institute.

Simon, L. (2025, August 4). *Is AI responsible for the rise in entry-level unemployment?* Retrieved September 9, 2025, from Revelio Labs: https://www.reveliolabs.com/news/macro/is-ai-responsible-for-the-rise-in-entry-level-unemployment/

SkillsFuture Singapore. (2018, May). *Skills Framework for Human Resource: A Guide to Occupations and Skills.* Retrieved September 10, 2025, from SkillsFuture Singapore: https://www.skillsfuture.gov.sg/docs/default-source/skills-framework/skills-framework-for-hr.pdf

The Federal Reserve Bank of New York. (2025, August 1). *The Labor Market for Recent College Graduates - Unemployment Rates for Recent College Graduates versus Other Groups*. Retrieved from Federal Reserve Bank of New York: https://www.newyorkfed.org/research/college-labor-market#--:explore:unemployment

Tsvetkova, A., D'Amico, E., Lembcke, A., Knutsson, P., & Vermeulen, W. (2024). *How well do online job postings match national sources in large English speaking countries? Benchmarking Lightcast data against statistical sources across regions, sectors and occupations.* Paris: OECD Publishing.

U.S. Bureau of Labor Statistics. (2025, August 28). *Occupational Outlook Handbook*. Retrieved from U.S. Bureau of Labor Statistics: https://www.bls.gov/ooh/about/glossary.htm#E