Australian Government
Department of Employment and Workplace Relations

AI and employment in Australia

Monitoring framework and evidence to date

July 2026

Contents

Executive summary

Key points

- There is no evidence to date of broad labour market upheaval driven by Artificial Intelligence (AI) in Australia. Overall labour market conditions remain strong by historical standards, youth outcomes have mostly held up, and occupational reshuffling has not accelerated.
- However, occupations more exposed to potential automation by generative AI have grown more slowly than other occupations. Between November 2022 and February 2026, employment in the most-exposed fifth of occupations grew by 5.6%, compared with 9.5% in the least-exposed fifth.
- We find a small negative relationship between AI exposure and employment growth. For an occupation with AI exposure one standard deviation above average, our model implies employment was about 2% lower by February 2026 than it would have been under its pre-ChatGPT trend.
- This is not definitive evidence of job loss caused by AI. The finding is sensitive to some modelling choices. In particular, the negative employment relationship is not statistically significant when we use 2 alternative AI exposure measures, nor when the COVID-19 period is excluded from the pre-treatment sample. The result should therefore be read as justification for ongoing monitoring rather than clear evidence that AI has reduced employment.
- The evidence is an early indication of some modest slowing in employment growth in some highly exposed occupations, not proof of large AI-driven job loss.
- DEWR will continue monitoring the relationship between AI exposure and employment growth.

AI could affect employment, but the direction and timing are uncertain

AI could reshape the labour market, but its effect on employment is highly uncertain. AI can perform tasks previously done by workers. But it can also raise productivity, create new tasks, and increase demand for labour in other parts of the economy. The net effect on jobs cannot be settled by theory or prediction. It needs to be monitored with data.

This report examines whether there is evidence that AI is affecting employment in Australia. It seeks to answer one main question: **since the public release of generative AI tools in late 2022, has employment in more AI-exposed occupations grown more slowly than employment in other occupations?** The report does not make predictions about potential future effects.

This report focuses on employment, not every effect of AI on work

We focus on the employment impact of generative AI tools, including large language models such as ChatGPT and Claude, rather than all technological change.

This report does not examine every way AI could affect work. It focuses on employment, hours worked and job advertisements by occupation. It does not measure the relationship between AI and wages, productivity, job quality, or work intensity. It does not examine firm-level AI adoption, changes in tasks within jobs or changes in skill requirements for occupations.

We use a monitoring framework, not a single indicator

We use a multi-dimensional monitoring framework. The framework combines descriptive labour market indicators, a core statistical model and robustness tests of that model.

Each part of the framework is important. The descriptive indicators show whether there are broad signs of labour market disruption. The core model tests whether employment growth shifted away from more AI-exposed occupations after ChatGPT was released. The robustness tests assess how much confidence to place in that model result.

The framework aims to measure what we observe in the data regarding the relationship between AI and employment. It does not make any predictions about potential future effects.

The labour market does not show signs of major AI-driven change

The labour market remains resilient. As at February 2026, the unemployment rate was 4.2%, lower than at any point in the decade before COVID.[1] The employment-to-population ratio was 64.0%, above its pre-COVID average. The volume underutilisation rate was 5.5%, also below its typical pre-COVID level. The hiring rate has slowed and warrants monitoring, but this is not enough to conclude that AI is disrupting the labour market.

The labour market for young workers and young graduates is also not showing broad deterioration. This matters because international evidence suggests young workers, and especially young graduates in exposed occupations, could be among the first groups to be affected by AI. Employment for people aged 20 to 24 has grown slightly faster than employment for people aged 25 and over since ChatGPT was introduced. The unemployment rate for young tertiary graduates remains low, the gap between young and older graduates' unemployment rates remains relatively low, and the share of young graduates working in degree-level occupations has risen.

The evidence does not show unusually rapid reshuffling across occupations. If AI were already causing widespread disruption, we might expect the occupational mix of employment to be changing faster than usual. We do not see that pattern. The pace of occupational compositional change has fallen since late 2022 and remains within its typical pre-COVID range. This is true for the labour market overall and for young workers.

Highly exposed occupations are growing more slowly than others

Using Jobs and Skills Australia's (JSA) occupation-level AI automation exposure scores,[2] we find that employment, hours worked and job advertisements have grown more slowly in highly exposed occupations than in less-exposed occupations. The most-exposed occupations include routine cognitive roles such as 'Filing and Registry Clerks', 'Keyboard Operators' and 'Telemarketers'.

Since November 2022, employment in the most-exposed fifth of occupations has grown by 5.6%, compared with 9.5% in the least-exposed fifth. The unemployment rate for workers previously employed in the most-exposed occupations has also risen more than for other workers, narrowing the gap with the unemployment rate for the least-exposed occupations.

These descriptive results are suggestive, but not conclusive. Some of the occupations most exposed to AI are routine cognitive jobs, such as clerical and administrative occupations, which were already on a long-running downward trajectory before the introduction of ChatGPT. The share of employment in routine cognitive work has continued to decline broadly in line with pre-existing trends. Slower growth in highly exposed occupations may therefore not reflect AI-driven disruption.

Our statistical model finds a modest negative signal

We use a statistical model to test whether the relationship between occupational AI exposure and employment growth changed after November 2022, compared with the pre-ChatGPT period. The

[1] February 2026 is the most recent available employment-by-occupation data from the ABS as at June 2026. (ABS 2026a) We use other data as at February to align with this primary data source.
[2] See the overview of the JSA exposure scores in the next chapter.

model uses quarterly employment data for 355 occupations from February 2015 to February 2026, along with JSA's AI automation exposure scores.

Employment has grown at a slower relative pace in more AI-exposed occupations since ChatGPT was released. For an occupation with an AI exposure score one standard deviation above the average, the model implies employment was about 2% lower by February 2026 than it would have been under its pre-ChatGPT trend, although this estimate is sensitive to modelling choices.

This estimate is much smaller than some public predictions of large AI-driven job losses, and it should not be interpreted as evidence of large-scale labour market disruption. The model does not observe whether firms adopted AI. It does not identify specific jobs created or destroyed. It estimates whether employment patterns changed in a way that is correlated with occupations' potential exposure to AI automation. Other factors, including post-COVID adjustment and longer-running occupational change, may also have affected more-exposed occupations.

The quarter-by-quarter version of the model suggests the negative exposure–employment relationship is clearer in recent quarters than immediately after ChatGPT was released. This is consistent with the possibility that any employment effects of AI may take time to emerge, as firms experiment with AI tools, adopt them in production and adjust staffing. But the quarter-by-quarter estimates are variable, so they should be interpreted as a tentative signal rather than definitive evidence.

Robustness tests suggest caution

The results of the statistical model are sensitive to some modelling choices, notably the choice of AI exposure measure. This means that we should interpret the model's findings with caution.

The result does not appear to be driven by a small number of occupations. The result also survives several timing and model-design variations, including a version that controls for occupations' historical sensitivity to the business cycle.

But the main result is not robust to every test. The AI exposure–employment relationship weakens or disappears when we use alternative measures of occupational AI exposure, or when we estimate the statistical model in a different way. These inconsistencies in the exposure–employment relationship are the strongest reason for caution. If future updates to this analysis find a statistically significant negative exposure–employment relationship when using alternative measures of exposure, that will be a sign that we can be more confident that AI is weighing on employment.

The evidence does not show large AI-driven job loss

Taken together, the evidence points to some modest and tentative signs that employment may be softening in more AI-exposed occupations. But there is no evidence to date of broad AI-driven labour market upheaval in Australia. Aggregate conditions remain strong, youth outcomes have mostly held up and occupational reshuffling has not accelerated. There is a modest and concentrated signal that employment in more AI-exposed occupations has softened relative to less-exposed occupations.

DEWR will continue to monitor the data. The framework in this report can be used to track whether the current signal grows, fades or persists. Future work should extend the evidence base by using administrative data and updating exposure measures as AI capabilities evolve.

Introduction

Key points

- Expert predictions about the potential future effect of AI on the labour market differ dramatically. A wide range of outcomes is possible.
- In theory, technologies such as AI can displace workers. But they can also increase employment by boosting productivity and creating new tasks. The net effect of AI on jobs could be either negative or positive. The overall effect is an empirical question – it can only be estimated with data.
- Existing evidence on the effect of AI on jobs is mixed. This evidence is mostly from overseas. Some studies find that youth employment in more-exposed occupations has grown more slowly in the period since generative AI became widely available. But most studies find no evidence of a negative overall effect of AI on jobs so far.
- This report measures whether occupations more exposed to AI have grown more slowly than other occupations in the period since ChatGPT was introduced. It also examines a range of other labour market indicators to see whether there are signs of disruption.
- The report sets up a framework that could be used for ongoing monitoring. The approach is multi-dimensional: descriptive indicators, a core statistical model and variations on the core model.
- The analysis in this report builds on JSA's GenAI capacity study, in particular using JSA's measures of each occupation's exposure to automation by AI tools.
- The focus of this report is employment and AI. It does not consider the labour market effects of other technologies, such as self-driving cars. It also does not examine other potential ways that AI could affect the world of work, such as work intensification or surveillance. Nor does the report examine the effect of AI on wages or any other outcomes beyond employment, hours worked and job advertisements. It does not include predictions or forecasts of potential future effects of AI.

The future effect of AI is highly uncertain

AI is a general-purpose technology that could potentially reshape the world of work.[3] But there is a wide range of views about the possible impact of AI on the labour market, including among well-credentialled observers. For example, Daron Acemoglu, the Nobel Prize-winning economist, expects minimal labour market disruption, predicting that 'only 5% of jobs will be heavily affected by AI over the next 10 years'.[4] At the other extreme, Dario Amodei, co-founder and CEO of Anthropic, has stated that 'AI could wipe out half of all entry-level white collar jobs and spike unemployment to 10% to 20% in the next one to five years'.[5] There are predictions at many points on the spectrum between these 'employment optimist' and 'employment pessimist' positions.[6]

This report does not add to the speculation about what may happen to employment in the future as a result of AI. Instead, it examines the data to February 2026 to answer the following core question: **since the public release of generative AI tools in late 2022, has employment in more AI-exposed**

[3] It is generally accepted that AI is a general-purpose technology. See, for example, Eloundou et al. (2023).
[4] Wittenstein (2024). See Acemoglu (2025) for the analysis underlying this media commentary.
[5] Comments reported by VandeHei and Allen (2025).
[6] Other papers, such as JSA (2025a), examine the potential future impact of AI and acknowledge the uncertainty around this impact.

occupations in Australia grown more slowly than other occupations? As well as answering this core question, we scan a broader range of descriptive indicators that may reveal other AI-related changes to the labour market. We are focused in particular on generative and agentic AI tools based on large language models (LLMs).

AI could reduce or increase employment

AI can already perform many tasks that are undertaken by human workers. As AI models improve, and adoption increases, it is possible that some workers may be displaced. But displacement is not an inevitable consequence of the ability of AI to perform productive tasks that have traditionally been undertaken by humans.[7]

New technology's net effect on work can be divided into 3 broad components:[8]

- **Displacement:** When technology can do tasks previously undertaken by humans, some firms may substitute technology for human labour, putting people out of work in the process. This is the effect that is often focused on in public discussion.
- **Productivity:** Automation increases the amount that can be produced for each hour of work, on average. This can be through automation that reduces the amount of human labour needed for a task, or augmentation which improves output quality for the same amount of labour.[9] This rise in productivity increases the demand for labour in non-automated tasks. As Deming et al. (2025, 181) describe, 'if productivity gains from automation greatly increase the size of the pie, workers may still get a larger slice even if one or two slices now get eaten by robots.'
- **Reinstatement**: Technology can create new tasks in which humans have a comparative advantage. Past technological changes have created a range of new tasks and occupations that did not previously exist. Most Americans are employed in jobs that did not exist in 1940,[10] and the Australian labour market has been similarly reshaped by past technological change.[11]

AI can reduce employment by replacing tasks, but it can also raise employment by increasing productivity and creating new tasks.[12] The net effect of AI on employment is ambiguous in theory. We will not know the effect until we can measure it with data.[13]

If we do not see job displacement from AI, this would mean the reinstatement and productivity effects have outweighed displacement. This could occur if AI has been used more to augment rather than automate tasks.[14]

Most studies find little overall employment effect so far

A range of studies, mostly overseas, have examined the effect of AI on employment. Most find limited effects so far. Here we highlight some of the more notable recent studies that focus on employment specifically.

[7] As Deming et al. (2025, 180) put it, 'the automation of individual job tasks does not necessarily reduce employment and may even lead to job gains in some sectors of the economy.'
[8] This framework is set out in Acemoglu and Restrepo (2019) and has been extremely influential in labour economics.
[9] JSA (2025b, 17) notes that AI can augment, as well as automate, tasks. They find there may be 'far greater potential for occupations to undergo augmentation than automation, based on current technology.'
[10] Autor et al. (2024)
[11] Borland and Coelli (2022) (pp. 19-20) observe that in Australia, in the past, the 'displacement effect of technological change has been consistently offset by reinstatement and productivity effects.'
[12] The productivity effect includes any situation in which certain tasks can be performed with a lower amount of human labour input, which in turn creates more demand for other non-automated tasks.
[13] Data alone is necessary but not sufficient to measure the effect of AI on employment. It also requires a clear framework for monitoring.
[14] The net effect of new technology on employment also depends on the price elasticity of demand for goods and services that have experienced automation in their production. See Bessen (2018).

A number of international studies find no or minimal impact. Gimbel et al. (2026), using US data, find that ‘the broader labor market has not experienced a discernible disruption since ChatGPT’s release’. Chandar (2025) and Humlum and Vestergaard (2026) also do not find significant decreases in employment for more exposed occupations in the US and Denmark, respectively. However, these papers generally do note that they cannot rule out some negative employment effects for some occupations or groups that may be outweighed by positive outcomes for others.

Some studies find negative effects, notably the American ‘Canaries in the coal mine’ paper by Brynjolfsson et al. (2025), particularly on the employment of young workers in highly exposed occupations. Massenkoff and McCrory (2026) also find small negative effects only for young workers in highly exposed occupations in the US. Klein Teeselink (2025) finds some negative impacts in the UK, again mostly for younger exposed workers. However, Lambert and Schindler (2026) find the decline for exposed young workers could be explained by AI-exposed occupations also having high rates of working from home, which could benefit established workers over new workers.

The Australian evidence is still limited. A recent study by CSIRO researchers (Mason et al. 2026) found no negative effects. A recent working paper using linked Australian administrative data (Gross 2026) finds that earnings and hiring outcomes in more-exposed occupations are overall no worse than for those less-exposed, though younger exposed workers show some signs of lower hiring rates.

Whether AI has reduced employment for exposed occupations is still an unsettled empirical question and a very active area of research. A survey by Kolko (2026) notes that early findings on AI’s effect on labour demand are inconclusive, and importantly only a weak signal for what may happen in future. Future studies will benefit from additional data, new data sources, and will reflect the evolving capabilities of AI models and their increased adoption.

We use a multi-dimensional approach

DEWR’s framework for monitoring AI and employment has 3 components:[15]

- **Descriptive indicators**: We examine a range of labour market indicators to assess whether they are moving in a direction consistent with AI-driven upheaval in the labour market.
- **Core statistical model**: We assess whether the relationship between occupations’ AI exposure and employment has changed in the period since November 2022.
- **Variations on the core model** (‘robustness tests’): We try a range of alternative modelling approaches to assess whether the findings from our core model are robust.

Each of these components is important. Descriptive indicators show whether there is a broad labour market disruption. The core model provides the central statistical test. Robustness tests determine how much confidence we should place in that result.

We treat the date on which ChatGPT was introduced to the public – 30 November 2022 – as the ‘treatment date’, the date that separates ‘before’ and ‘after’ when it comes to AI. This is consistent with a range of other studies.[16] There is nothing ChatGPT-specific in our analysis other than this timing.

We use data up to February 2026, which is the most-recently-available release of employment-by-occupation data from the ABS Labour Force Survey.[17] We use February 2026 as the reference date,

[15] The idea of a multi-dimensional approach draws from Eckhardt and Goldschlag (2025).
[16] For example, Brynjolfsson et al. (2025) and Klein Teeselink (2025)
[17] ABS (2026a)

even for series such as the unemployment rate that have more recent monthly data available. We do this so that our analysis captures a snapshot of a single point in time.[18]

We use JSA's estimates of AI exposure

This report builds on the JSA capacity study *Our Gen AI Transition*.[19] We use the GenAI automation exposure scores developed by JSA to classify each occupation's level of potential exposure to automation by AI. Further details on the construction of these scores and their limitations, as well as information about alternative exposure measures, can be found in Appendix A at page 63.

Figure 1 shows the distribution of AI automation exposure scores, as estimated by JSA, across individual occupations as defined by the Australian and New Zealand Standard Classification of Occupations (ANZSCO).[20] In theory, the scores range from 0 to 1, but in practice no occupation is either fully exposed or fully insulated from generative AI, so the actual scores run from 0.10 to 0.81.

The most-exposed occupations are generally routine cognitive jobs like Filing and Registry Clerks or Keyboard Operators, as their tasks are most able to be automated by generative AI. The least exposed jobs are typically manual, such as Domestic Cleaners and Handypersons.

Figure 1: Routine cognitive occupations like 'Keyboard Operators' are most exposed to automation by generative AI

Distribution of occupations' JSA AI automation scores

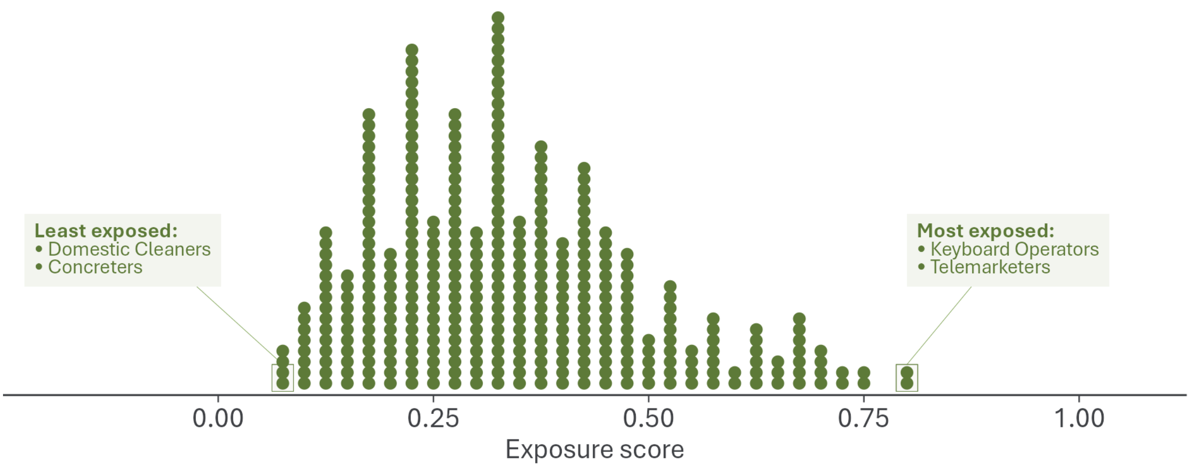

Notes: Each 'dot' is a 4-digit ANZSCO occupation.
Source: JSA (2025a).

JSA's exposure scores are based on work by researchers at the International Labour Organisation (ILO).[21] They use a robust, defensible methodology. But any measure of the potential exposure of individual occupations to future automation by AI is inherently speculative and imprecise. For this reason, our analysis includes a range of descriptive indicators that do not rely on any exposure scores, such as the rate of compositional change in the jobs market. We also estimate our statistical model using alternative measures of AI exposure as a robustness test — see page 52.

For our descriptive analysis, we divide occupations into 5 equal groups ('quintiles') based on their AI exposure.[22] Table 1 shows the 10 largest occupations in each quintile of AI exposure. The group with

---

[18] A small number of indicators are not available for February 2026. For these, we instead use the latest data available prior to February 2026. This is generally November 2025.
[19] JSA (2025a) and JSA (2025b)
[20] JSA's scores are assigned at the level of four-digit ANZSCO occupation unit groups. See ABS (2022b).
[21] Gmyrek et al. (2023), which builds on the methodology of Eloundou et al. (2023).
[22] This approach has been taken by others, including Eckhardt and Goldschlag (2025).

the highest exposure mostly contains cognitive roles – clerks, programmers, and administrative roles. The least exposed quintile consists mostly of manual roles – carers, tradespeople, and drivers.

Table 1: The largest occupations in each exposure quintile

Least exposed	Second least exposed	Middle exposure	Second most exposed	Most exposed
Aged & Disabled Carers	Kitchenhands	Registered Nurses	Sales Assistants (General)	General Clerks
Electricians	Waiters	Storepersons	Advertising, Public Relations & Sales Managers	Retail Managers
Truck Drivers	Motor Mechanics	Secondary School Teachers	Office Managers	Software & Applications Programmers
Child Carers	Bar Attendants & Baristas	Primary School Teachers	ICT Managers	Accountants
Commercial Cleaners	Nursing Support & Personal Care Workers	Chefs	Other Hospitality, Retail & Service Managers	Receptionists
Carpenters & Joiners	Drillers, Miners & Shot Firers	Construction Managers	Welfare Support Workers	Contract, Program & Project Administrators
Metal Fitters & Machinists	Shelf Fillers	Education Aides	Finance Managers	Accounting Clerks
Plumbers	Police	Solicitors	Production Managers	Purchasing & Supply Logistics Clerks
Gardeners	Chief Executives & Managing Directors	Human Resource Managers	Cafe & Restaurant Managers	Advertising & Marketing Professionals
Forklift Drivers	Early Childhood (Pre-primary School) Teachers	General Practitioners & Resident Medical Officers	University Lecturers & Tutors	Checkout Operators & Office Cashiers

Notes: Table lists top occupations for each AI exposure quintile, ranked by employment level as at February quarter 2026.
Source: JSA (2025a), ABS Labour Force (Detailed) and DEWR calculations.

Table 2 shows some characteristics of workers in occupations at different levels of exposure to potential automation by AI. Nearly 70% of workers in the least-exposed group are male, whereas men represent only 43.7% of workers in the most-exposed group. Relatively few workers in the least-exposed group have tertiary qualifications (14.9%),[23] whereas 43.7% of the most-exposed group are tertiary-qualified. The average age of each group is broadly similar, as is the proportion of the group that were born overseas.

The exposure scores represent an estimate, as at 2025, of the potential exposure of occupations to displacement due to AI automation. If AI model capabilities advance rapidly, these scores could become out of date,[24] which would mean any monitoring framework built on these scores could mismeasure the AI exposure–employment relationship. For this reason, we incorporate a range of descriptive indicators that do not rely on the JSA scores, and we will incorporate alternative scores into the monitoring framework if and when they become available.

[23] In this report, 'tertiary' qualifications refer to bachelor degree or higher.
[24] Changes in model capability would only be an issue for monitoring of employment effects if model development is not uniform across different types of capabilities, such that the relative ranking of occupations' exposure is changed.

Table 2: Workers in the most exposed occupations are more likely to be women and have tertiary qualifications than those in the least-exposed occupations

AI exposure quintile	Sex	Educational attainment (highest level)			Age		Migrant status
	% male	% with Year 12 or lower	% with VET quals	% with bachelor degree or higher	Average age	% aged 15–24	% born overseas
Least exposed	69.5%	39.4%	41.4%	14.9%	40.0	17.4%	35.2%
Second least exposed	52.5%	39.3%	29.3%	27.8%	36.6	27.1%	33.2%
Middle exposure	47.6%	17.7%	21.1%	58.0%	41.6	9.8%	36.0%
Second most exposed	49.5%	32.5%	22.4%	41.9%	40.9	16.4%	34.6%
Most exposed	43.7%	29.7%	23.4%	43.7%	41.6	11.4%	36.7%
All workers	**51.9%**	**30.9%**	**27.1%**	**38.6%**	**40.5**	**15.3%**	**35.4%**

Notes: Average of the 4 quarters to November 2025 (2026 data not available). 'VET qualifications' refers to Certificate III / IV and Diploma / Advanced Diploma. The educational attainment columns do not sum to 100, as the totals include 'Certificate not further defined' and 'Level not determined', which are not included in the columns.
Sources: ABS Labour Force (TableBuilder) and DEWR calculations.

What we are not measuring

This report is focused on the relationship between AI and employment. There are many other labour market issues that are beyond the scope of this report. AI could potentially have a range of effects on the world of work. It could affect productivity, job quality or the intensity of work. It could change the types of tasks that particular occupations undertake and thereby alter the skill requirements of jobs. We do not seek to measure any of these potential effects in this report.

We do not measure the effect, if any, of AI on wages. In theory, AI could have a positive or negative effect on wages, depending on the value of the remaining non-automated tasks.[25] Even studies that have found a negative effect of AI on employment for some workers overseas have not found wage effects.[26] Any effect of AI on wages would likely take longer to appear in the data than an effect on employment, given that wages tend to be ‘sticky’ and can take time to adjust. A potential future extension of this work is to examine the relationship between AI and wages, but it is out of scope of this report. This report is focused on the central question of whether AI is reducing employment.

Other technologies could affect work, including potentially reducing employment in some occupations. For example, self-driving vehicles could be potentially disruptive. This report does not examine all sources of potential technological disruption to the labour market, focusing instead on generative AI specifically.

We do not examine specific demographic groups, other than young people. Young people are thought to be specifically susceptible to displacement by AI. This is because AI models are generally best at performing tasks that require codified knowledge (‘book-learning’) but ‘less capable of replacing tacit knowledge, the idiosyncratic tips and tricks that accumulate with experience’,[27] which younger workers disproportionately lack. Young people are therefore more susceptible to automation *within occupations*. Other demographic groups, such as women, are thought to be susceptible to automation

[25] See Gans and Goldfarb (2026) and Autor and Thompson (2025).
[26] See Brynjolfsson et al. (2025).
[27] Brynjolfsson et al. (2025)

because they are disproportionately employed in exposed occupations (see Table 2), not because they are inherently more susceptible. Examining employment trends for young people can give some early indication of potential disruption to the labour market for other workers, whereas the effect on other demographic groups is likely to be downstream of occupation-level trends.

We do not use data on AI adoption by firms.[28] This report assesses whether AI appears to be reducing employment. If it is not, a potential explanation is that AI has not yet been deployed by a sufficient number of employers. Firms' plans for future adoption of AI could also affect employment, as it could influence current hiring decisions. This report does not include measures of planned adoption.

We are focused on using data to examine the effect of AI on employment in Australia to date. This report does not include any predictions of potential future change. We do not measure upstream technological capability, nor the cost or accessibility of AI tools.

Most of our analysis relies on survey data, specifically from the ABS Labour Force Survey. The limitations of a sample survey mean we are restricted in our ability to do detailed analysis of occupational employment trends by demographic characteristics, such as age. Future updates to the monitoring framework may incorporate administrative data, such as the ATO Single Touch Payroll (STP) data, which would increase our ability to do this sort of analysis.

Monitoring cadence

We plan to use the monitoring framework set out in this report as the basis for ongoing quarterly monitoring of the Australian labour market.

Updates will not be possible until late 2026, due to a transition in the ABS Labour Force Survey.[29] This transition will include ceasing to publish employment by ANZSCO occupations, which will require us to convert to the new Occupation Standard Classification for Australia (OSCA) framework.[30]

[28] See DISR (2026) for some analysis of employment in Australian firms with different levels of AI adoption.
[29] ABS (2026c)
[30] ABS (2024).

Descriptive indicators

Key points

- Overall, the Australian labour market remains resilient as at February 2026. Most measures show that, while conditions have eased since late 2022 when ChatGPT was introduced, the jobs market is stronger than in the decade prior to COVID.
- The unemployment rate, at a 3-month average of 4.2% as at February 2026, is lower than at any point in the decade prior to COVID. 64.0% of Australians aged 15 and over are in work, well above the average employment-to-population ratio in the decade before COVID (61.6%) and around the level in late 2022 when ChatGPT was introduced.
- Other key indicators also remain strong. For example, the volume underutilisation rate, which compares the total hours worked in Australia with the total hours both employed and unemployed people say they want to work, is at 5.5%, well below the pre-COVID average of 7.6%.
- The hiring rate has fallen. This rate, measured as the proportion of workers who have recently commenced with their employer, is at a relatively low level. But there are strong alternative explanations for the low hiring rate other than an AI effect.
- The labour market for young people has held up well, with youth (aged 20–24 years) employment growth outpacing that of older (25 years and above) people in the period since November 2022 when ChatGPT was introduced.
- Employment outcomes for young tertiary graduates have been positive, despite expectations young graduates could be among the first groups to be affected by AI. Their unemployment rate remains low relative to older graduates, and the share of young graduates employed in degree-level jobs has risen.
- There is no sign of unusually rapid reshuffling across occupations. If AI were causing broad labour market disruption we would expect faster compositional change, but the pace of change has eased since 2022 and remains within its normal pre-COVID range.
- Some highly exposed occupations have grown strongly. For example, software developer employment ('Software and Applications Programmers') has increased by 25% since November 2022.
- However, an area of emerging concern is that, as a whole, the occupations most exposed to AI automation are growing more slowly than others. Those occupations have seen slower employment and hours growth, larger falls in job advertisements, and a narrowing unemployment advantage relative to low-exposure jobs. These findings may reflect pre-existing trends.

Most labour market indicators do not show broad AI-driven labour market upheaval. But several measures point to softer conditions in occupations more exposed to AI automation.

We compare each descriptive indicator with its November 2022 value and its pre-COVID range. This is to inform a judgement about whether the indicators are in line with 'normal' levels. The descriptive indicators in this chapter should not be interpreted in a 'causal' way – we cannot infer from these indicators whether the introduction of AI has *caused* a change.

These descriptive indicators are intended to complement and cross-check the formal statistical modelling that will come in the following chapter.

The labour market remains resilient

The Australian labour market experienced a rise in unemployment and a fall in employment during 2020 due to the COVID-19 pandemic and the restrictions put in place to contain the spread of the virus. In 2022, the labour market strengthened in Australia (and elsewhere), due to the availability of vaccines, removal of restrictions, and the (in some cases lagged) effect of macroeconomic support put in place during the pandemic, including low interest rates and fiscal support.[31]

The post-pandemic labour market reached its tightest point during 2022. In November 2022, the unemployment rate reached its lowest level since 1974. Conditions have since eased but remain significantly stronger than during the decade before COVID.

The unemployment rate remains well below its typical level in the decade before COVID. Although there has been some moderation in the period since 2022, unemployment remains very low compared to recent decades in Australia. At 4.2% as at February 2026, the unemployment rate remains lower than any level recorded in the decade prior to COVID.[32] While the unemployment rate did begin to rise from late 2022, this is more plausibly due to macroeconomic factors including the effect of global supply shocks and rising interest rates than the introduction of AI tools for widespread use.

Figure 2: The unemployment rate remains well below its typical pre-COVID range

Unemployment rate (3 month average, per cent of labour force)

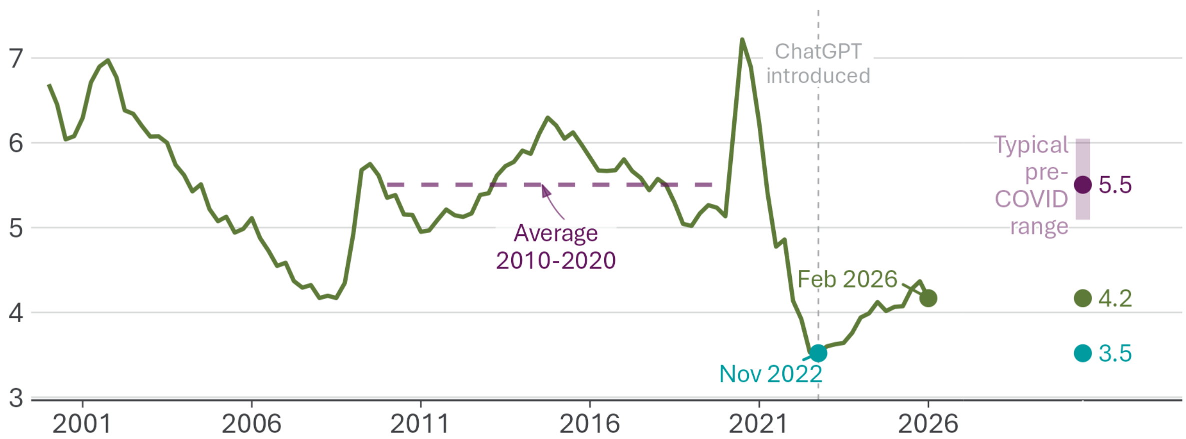

Notes: 3 month average of seasonally adjusted monthly data. Smaller values of this indicator are less consistent with labour market disruption. Pre-COVID average uses data from February 2010 to February 2020. 'Typical' range depicts the 10th to 90th percentiles. Source: ABS Labour Force.

[31] See Borland (2026) for a discussion of this period.

[32] ABS (2026b). 3-month averages of seasonally adjusted monthly data. Note that this may not match the headline monthly unemployment rate for the same period.

About these charts

Many of the data visualisations in this report follow the style of Figure 2. It compares the **latest value**, generally **February 2026**, to the value in **November 2022**, the final quarter before **ChatGPT** was introduced. It also shows the **pre-COVID average** value. The pre-COVID average period for Figure 2 and most figures runs from February 2010 to February 2020. Where data is not available for this full period, the longest period available is used and this is made clear in the visualisation. The typical pre-COVID range is shown, which covers the middle 80% of observations over the pre-COVID period. The points on the right-hand side reflect the levels as illustrated in the line chart – these points are then combined in 'dashboard charts', such as Figure 6.

These charts compare current values with historical levels, to assess whether they have moved in a direction that would be consistent with AI-driven changes in employment patterns. The charts are not conclusive in themselves. Judgement is required in interpreting them, as some time series were trending up or down prior to the introduction of ChatGPT in a way that could make the pre-COVID average and typical range an imperfect guide to what would reasonably be expected if generative AI tools had not been introduced. A range of factors other than AI, including macroeconomic shocks and business cycles, affect these series.

Most other headline indicators also suggest ongoing strength in the Australian labour market. For example, as at February 2026, the employment-to-population ratio is 64.0%. This is substantially above the level in the decade prior to the COVID pandemic, when the employment-to-population ratio averaged 61.6%, as shown in Figure 3. This ratio, which measures the proportion of the civilian population aged 15 or above that is employed, provides a clear, high-level measure of overall labour market conditions.

Figure 3: The share of the population in work remains very high

Employment-to-population ratio (3 month average, per cent of population aged 15+)

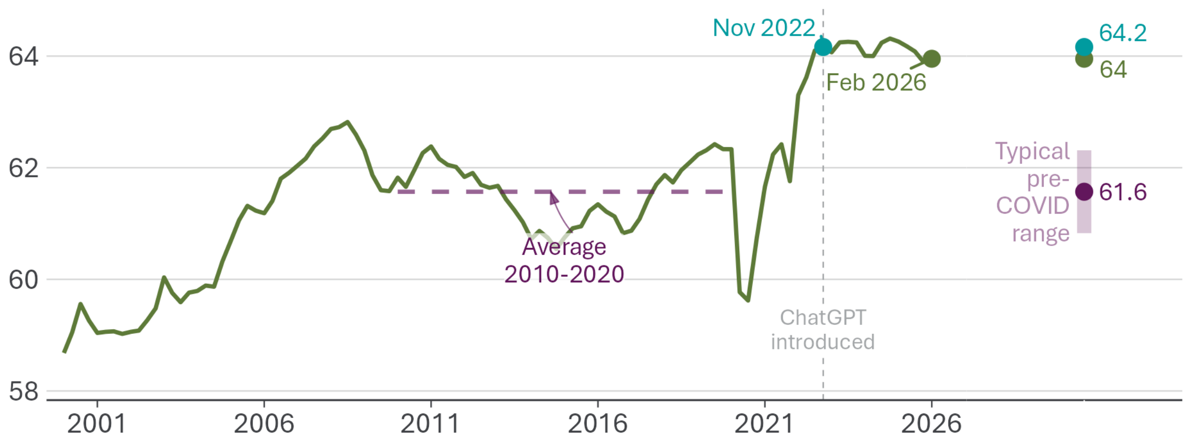

Notes: 3 month average of seasonally adjusted monthly data. Larger values of this indicator are less consistent with labour market disruption. Pre-COVID average uses data from February 2010 to February 2020. 'Typical' range depicts the 10th to 90th percentiles. Source: ABS Labour Force.

The underutilisation rate also suggests that the overall labour market remains resilient. We examine the volume underutilisation rate, which compares the hours that Australians (both employed and unemployed) want to work with the hours they actually work. This is a comprehensive measure of the spare capacity in the labour market, as it reflects both unemployment (people who want to work, but currently do not) and underemployment (people who are employed but would like more hours of work). As with the unemployment rate, a lower number means less labour market slack.

As at February 2026, the volume underutilisation rate is 5.5%, well below the pre-COVID average of 7.6% and outside the typical range experienced in the decade before COVID (Figure 4).[33]

Figure 4: The labour underutilisation rate – the gap between actual and desired hours of work – remains below its pre-COVID average

Volume underutilisation rate (desired unworked hours of employed and unemployed people, as a percentage of total desired hours, 3 month average)

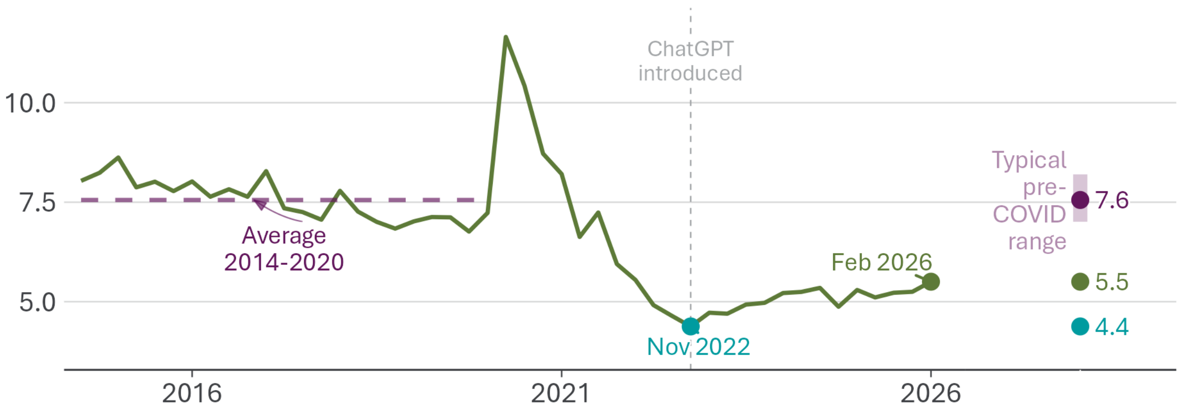

Notes: 3 month average of seasonally adjusted monthly data. Smaller values of this indicator are less consistent with labour market disruption. Pre-COVID average uses data from August 2014 to February 2020. 'Typical' range depicts the 10th to 90th percentiles. Source: ABS Labour Force.

Most headline indicators – the unemployment rate, employment-to-population ratio, and volume underutilisation rate – remain stronger than their pre-COVID typical levels and reflect ongoing resilience in the Australian labour market. But one key headline measure, the hiring rate, has fallen and is below its typical levels.

The hiring rate is 4.2% as at February 2026, as shown in Figure 5.[34] This means that only 4.2% of workers commenced with their current employer in the past 3 months, compared to a pre-COVID average of 5.1% and a November 2022 level of 5.5%. Although a low hiring rate could be consistent with AI-driven effects on jobs, this may not be the case, as:

- Most other headline indicators remain resilient.
- The hiring rate was also low in early 2020, prior to COVID and before the introduction of ChatGPT.
- The period in which the hiring rate fell coincided with tightening monetary policy. Higher interest rates generally reduce hiring.
- The hiring rate has been trending down over the longer term.
- We have seen the hiring rate decline in some previous periods (such as after the Global Financial Crisis).
- The fall in the hiring rate appears to have been driven by a decline in job-to-job transitions rather than a fall in the rate at which people move into employment. Transitions from non-employment to employment remain above pre-pandemic levels.

[33] ABS (2026b). Seasonally unadjusted quarterly data.
[34] ABS (2026a). 4 quarter average of seasonally unadjusted data.

Figure 5: The hiring rate is low, but is around its early 2020 (pre-COVID) level

Hiring rate (employed persons who have been with their current employer for fewer than 3 months, as a percentage of all employed persons)

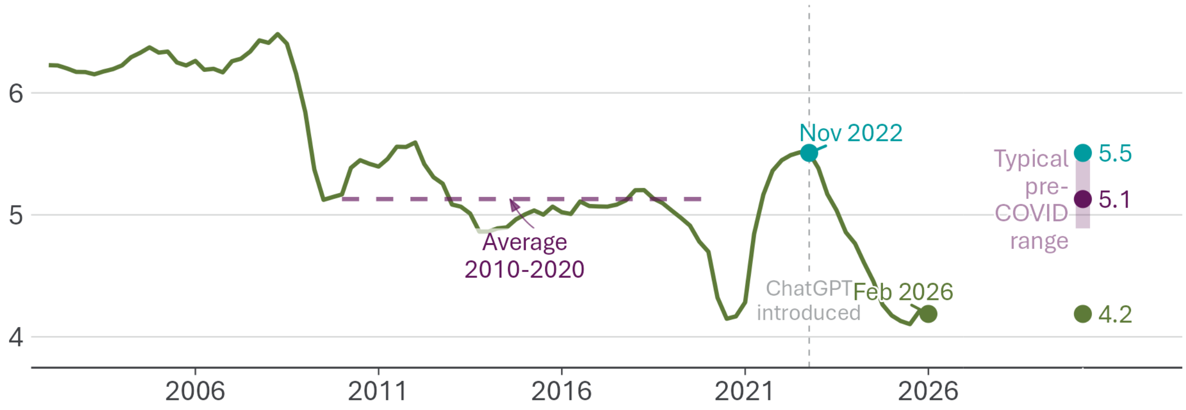

Note: 4 quarter moving average of seasonally unadjusted data. Larger values of this indicator are less consistent with labour market disruption. Pre-COVID average uses data from February 2010 to February 2020. 'Typical' range depicts the 10th to 90th percentiles.
Source: ABS Labour Force (Detailed).

Overall, the data suggests continued resilience in the Australian labour market, with the exception of the hiring rate. Figure 6 provides an overview of these 4 key indicators and how they compare to typical levels experienced in the decade prior to COVID.

Figure 6: Most overall labour market indicators suggest ongoing resilience

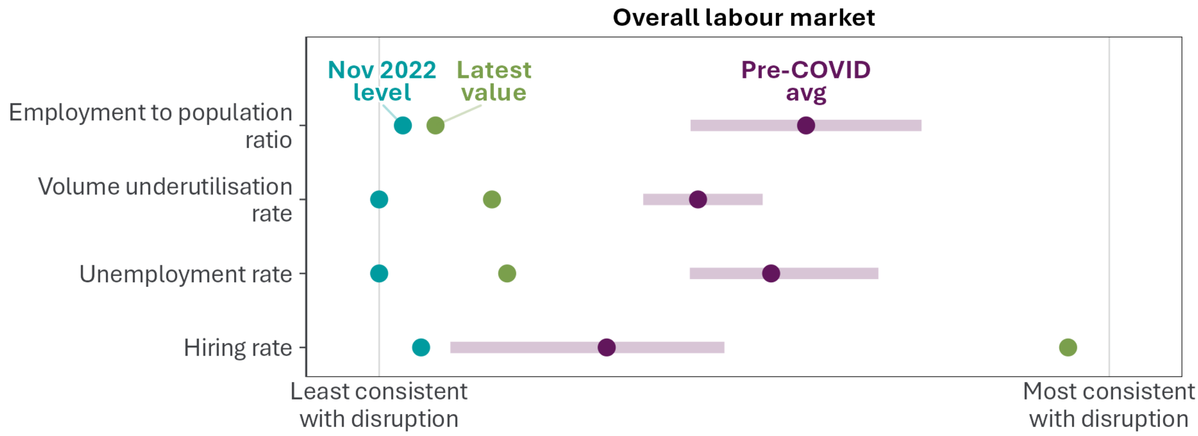

Notes: Hiring rate is a four-quarter moving average of seasonally unadjusted quarterly data. Other indicators are three-month moving averages of seasonally adjusted monthly data.
Sources: ABS Labour Force (Detailed) and DEWR calculations.

About these charts

Figure 6 is a 'dashboard chart' that summarises a range of indicators, comparing their **latest values** to their **November 2022 (pre-ChatGPT)** level as well as their **typical level in the decade before the COVID pandemic**. The vertical lines represent the most extreme values observed in the study period.

Series are aligned so that values towards the left of the chart are less consistent with AI-driven disruption. Values to the right of the chart are more consistent with disruption but are not necessarily 'worse'. In some cases, like the unemployment rate, a lower rate is 'better' and is also less consistent with AI-driven disruption. For some series, such as the rate of compositional change in the labour market (see Figure 15) this is more complicated or ambiguous. Faster compositional change can be 'good' in that it is a sign of a dynamic, well-functioning labour market, but it also could be a sign of AI-related change to the jobs market.

These charts compare time series to their past values. This historical comparison can help inform judgements about whether we are seeing anything 'out of the ordinary' in the jobs market but is not conclusive. Some time series were trending up or down prior to the introduction of ChatGPT in a way that could make the pre-COVID average and typical range an imperfect guide to what would reasonably be expected if generative AI tools had not been introduced. It is also important to remember there are many factors that influence labour demand, both overall and for specific jobs.

Employment outcomes for young workers and graduates remain resilient

Young people are particularly susceptible to negative effects from AI automation. This is for at least 2 reasons:

- When economic conditions change, employers tend to adjust their rate of hiring more than the rate at which they retrench existing staff. Young people are more likely to be affected by a reduction in hiring rates, as they are more likely than older people to be moving into work for the first time.[35]
- Large language models (LLMs) are trained on a corpus of text that includes textbooks and academic articles but does not include the 'tacit knowledge' that workers acquire through on-the-job experience. This can mean that LLMs are more able to substitute for young people, particularly young graduates, than workers with more experience.[36]

For these reasons, young people are sometimes described as the 'canaries in the coal mine' for AI-related automation.[37] Early evidence from the US suggests that young people have experienced slower employment growth than older people since the introduction of generative AI tools – especially in occupations more exposed to AI, though this is contested by some scholars.[38] This section examines the youth labour market, and particularly labour market outcomes for young tertiary graduates, given the expectation that any effects of AI on employment are likely to show up first for this group.

---

[35] See Dhillon and Cassidy (2018) for more discussion of the Australian youth labour market.
[36] See Acemoglu and Autor (2011)
[37] Brynjolfsson et al. (2025)
[38] Brynjolfsson et al. (2025, 1) find that 'early career workers (ages 22-25) in the most AI-exposed occupations have experienced a 13 percent relative decline in employment even after controlling for firm-level shocks.' Lambert and Schindler (2026) find that that this is explained by work from home (WFH) affecting hiring decisions, with a decline in hiring juniors in occupations more likely to allow WFH.

Brynjolfsson et al. (2025, 10) show that in the US, there has been 'some levelling off in employment growth for young workers relative to other age groups, consistent with recent discussion of a worsening job market for entry-level workers.' We do not see this in the Australian data. Headline indicators of the youth labour market remain reasonably solid in historical terms.[39]

Figure 7 shows that, on the contrary, the number of employed people aged 20-24 in Australia has increased slightly *faster* than employment of people aged 25 and above.[40]

Figure 7: Since ChatGPT was introduced, employment for young adults has grown slightly faster than for people aged 25+, unlike in the US

Employment by age group (index, Nov 2022 = 100)

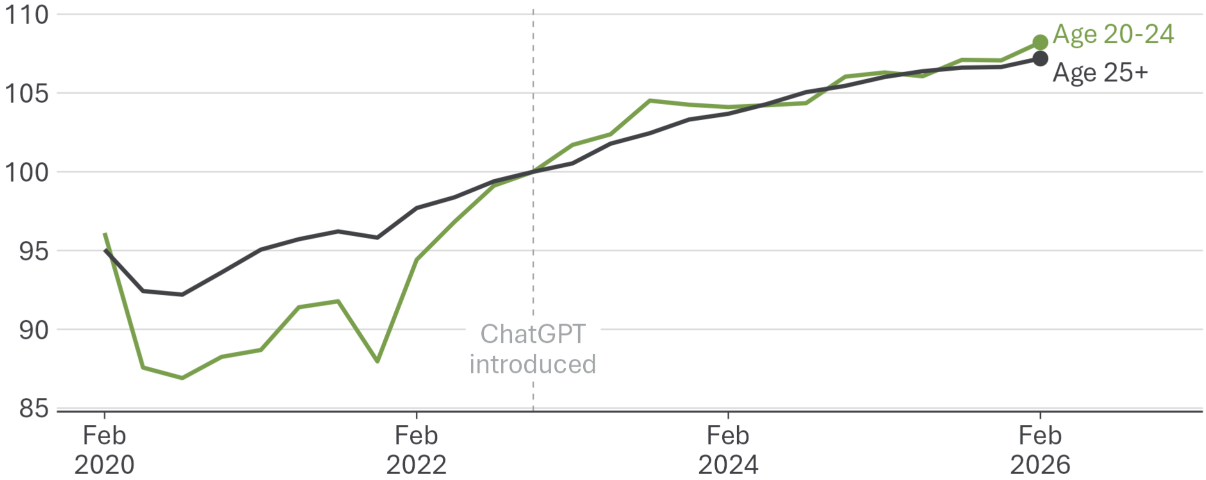

Note: 3-month moving averages of seasonally adjusted data.
Sources: ABS Labour Force and DEWR calculations.

The youth unemployment rate is higher than the overall unemployment rate, as it has been for every month on record. But the *gap* between the youth unemployment and overall unemployment rates is relatively low - below its typical pre-COVID range, as shown in Figure 8.

As at February 2026, the overall unemployment rate is 4.2% and the youth (age 15-24) rate is 9.7%. The gap between the two is 5.5 percentage points,[41] below the pre-COVID average of 6.7 percentage points. Although youth unemployment has risen from its very low 2022 levels – both in absolute terms and relative to overall unemployment – it remains low in historical terms.

[39] In March and April 2026, the monthly youth unemployment rate increased. This is beyond the period considered in this report, as February 2026 is the most recently available employment-by-occupation data, which is the primary data for the report.
[40] This partly reflects strong population growth at younger ages.
[41] This gap is calculated from unrounded figures.

Figure 8: The gap between the youth and overall unemployment rates is lower than pre-COVID

Difference between youth and overall unemployment rate (percentage points)

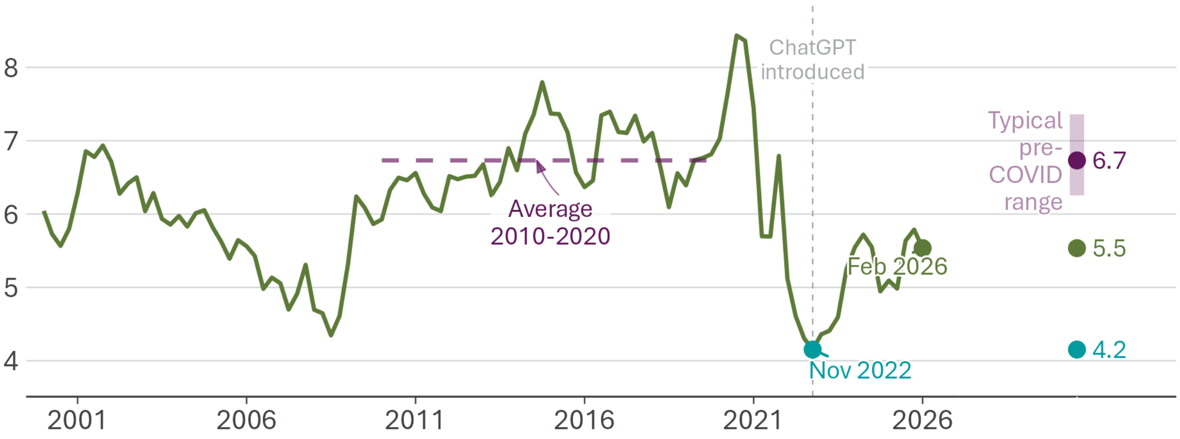

Notes: 3 month average of seasonally adjusted monthly data. Smaller values of this indicator are less consistent with labour market disruption. Pre-COVID average uses data from February 2010 to February 2020. 'Typical' range depicts the 10th to 90th percentiles. Source: ABS Labour Force.

Unemployment rates for young tertiary graduates, specifically, also remain relatively low. The unemployment rate for young (age 20–24) graduates is 5.4%, lower than at any time in the 5 years prior to COVID.[42] Figure 9 compares the unemployment rates for young graduates with young people without a tertiary qualification, as well as the rate for older graduates.

Figure 9: The unemployment rate for young graduates remains low

Unemployment rate by age and level of highest educational attainment, per cent of labour force

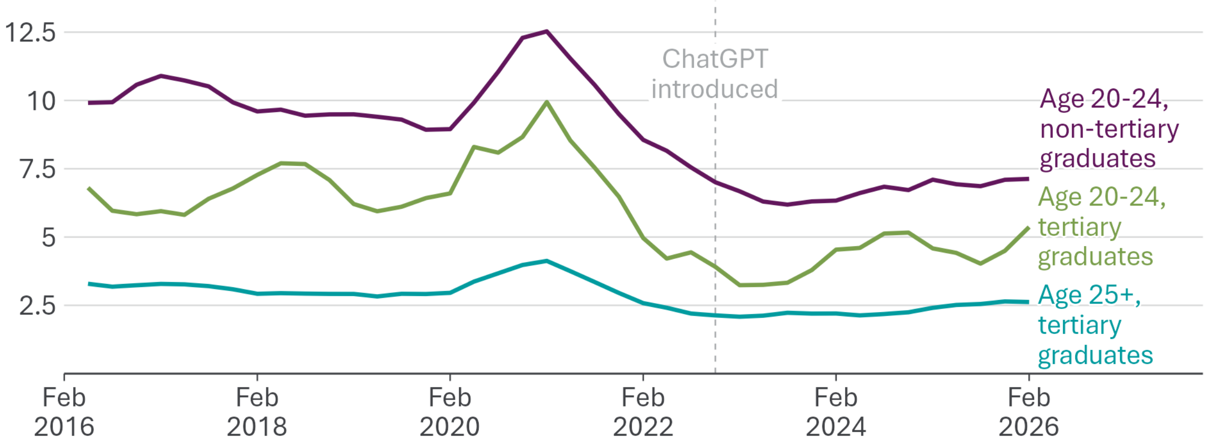

Notes: 4-quarter moving average of non-seasonally adjusted data.
Sources: ABS Labour Force (Detailed) and DEWR calculations.

The unemployment rate of young graduates remains low in historical terms. It is also performing relatively well compared to the rate for older (age 25 and older) tertiary graduates. The gap between the unemployment rates for young and older graduates is 2.7 percentage points. While this is above its immediate pre-ChatGPT level, it is around the bottom of the typical pre-COVID range, as shown in Figure 10. This suggests there has not been a specific deterioration in labour market conditions for young graduates.

[42] Quarterly ABS data on labour force status by age and educational attainment is only available from 2015 onwards. Data cited here are 4-quarter rolling averages of seasonally unadjusted data.

Figure 10: The gap between young graduates' and older graduates' unemployment rates remains very low, suggesting continued strength in the market for young graduates

Difference between young (20-24) tertiary graduates' and older (25+) graduates' unemployment rates (percentage points)

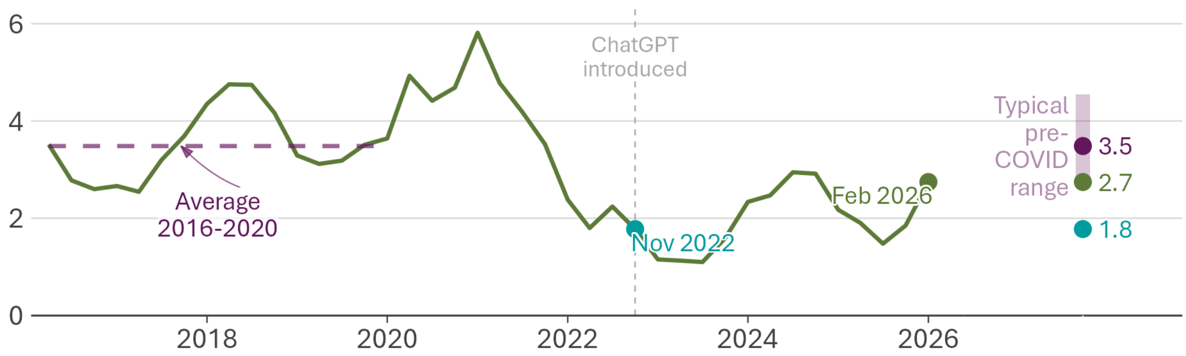

Notes: 4 quarter average of non-seasonally adjusted quarterly data. Series commenced in 2015. 'Graduates' includes all employees whose highest level of educational attainment is a bachelor degree, graduate diploma or certificate, or postgraduate degree. Smaller values of this indicator are less consistent with labour market disruption. Pre-COVID average uses data from May 2016 to February 2020. 'Typical' range depicts the 10th to 90th percentiles.
Source: ABS Labour Force (Detailed).

Another indicator of interest is the gap between the unemployment rates of young graduates and young people without a tertiary degree. This gap has fallen slightly, suggesting conditions for young non-graduates have improved relative to those of young graduates. This gap is now outside the typical pre-COVID range but is not at unprecedented levels – the gap between young graduates and non-graduates was similarly low in 2020 and 2018. This narrowing could be due to cyclical factors rather than technology.[43] Dhillon and Cassidy (2018) point out that the Australian youth-to-overall unemployment rate gap tends 'to increase when economic conditions slow and narrow when conditions improve'. Given that young graduate unemployment remains low in absolute terms, and the young–older graduate unemployment gap remains relatively low, this gap should be treated as a cautionary relative indicator rather than evidence of a clear deterioration in young graduates' labour market outcomes.

AI could have an effect on young graduates' employment outcomes even if it does not result in rising unemployment. One possibility is that more young graduates could end up in occupations that are below their 'skill level' – roles that typically do not require a tertiary qualification. For example, if more young tertiary graduates are unable to find professional jobs, but remain employed in retail or hospitality, this will show up as a decline in the share of young graduates employed in high-skill occupations.

The ABS classifies occupations according to skill level. Those in Skill Level 1 have 'a level of skill commensurate with a bachelor degree or higher qualification'.[44] In November 2022, 51.1% of young (age 20–24) graduates were employed in Skill Level 1 occupations. This figure has risen – in November 2025, 52.2% were employed in Skill Level 1 jobs. A larger share of young graduates is employed in occupations that use their qualifications than before ChatGPT, as shown in Figure 11. We are not seeing a rise in the share of young graduates employed below their skill level.

[43] Berger (2026) notes that, because graduates in their mid-20s are likely to have less labour market experience than non-graduates of the same age, they may be more vulnerable to cyclical labour market downturns unrelated to technological change. Ongoing monitoring will assist in assessing whether a reduction in the unemployment rate gap between young graduates and non-graduates is a temporary, cyclical phenomenon or more enduring and therefore more plausibly associated with technological change.
[44] ABS (2022b)

Figure 11: Most young graduates are still employed in occupations that typically require a degree
Share of young (20-24) tertiary graduates employed in Skill Level 1 occupations (per cent)

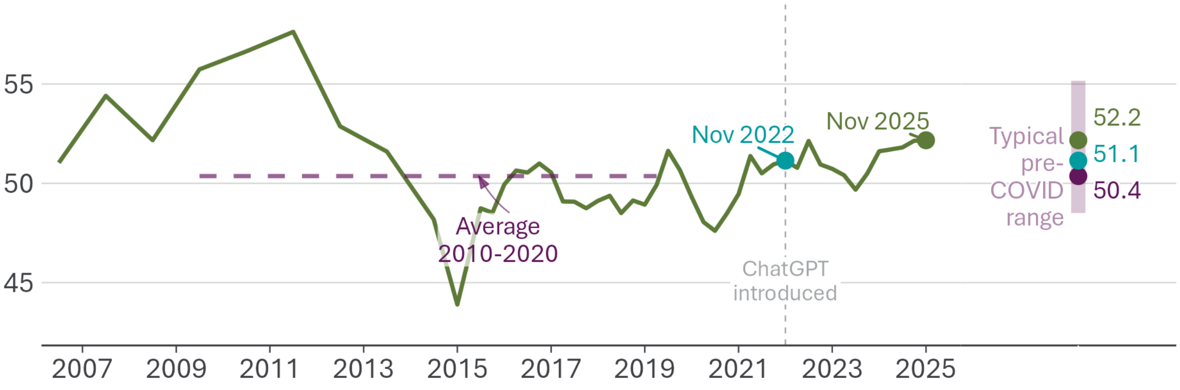

Notes: Annual data to 2015, 4 quarter average of non-seasonally adjusted quarterly data thereafter. 'Graduates' includes all employees whose highest level of educational attainment is a bachelor degree, graduate diploma or certificate, or postgraduate degree. Larger values of this indicator are less consistent with labour market disruption. Pre-COVID average uses data from February 2010 to February 2020. 'Typical' range depicts the 10th to 90th percentiles.
Source: ABS Labour Force (TableBuilder).

Some US-based analysis suggests that young graduates are likely to be the ‘canaries in the coal mine’ for AI-related job loss.[45] Most Australian data as at February 2026 suggests young graduates in Australia are not faring worse than they were pre-ChatGPT:

- The gap between youth and overall unemployment rates is lower than its typical pre-COVID level.
- The unemployment rate gap between young and older tertiary graduates is relatively low.
- The share of young graduates employed in occupations that require a degree has gone up.

The main data point that suggests conditions have worsened for young graduates is the fall in the unemployment rate gap between young graduates and young non-graduates.

[45] Brynjolfsson et al. (2025)

Figure 12: Most youth labour market indicators are within, or better than, their typical pre-COVID levels

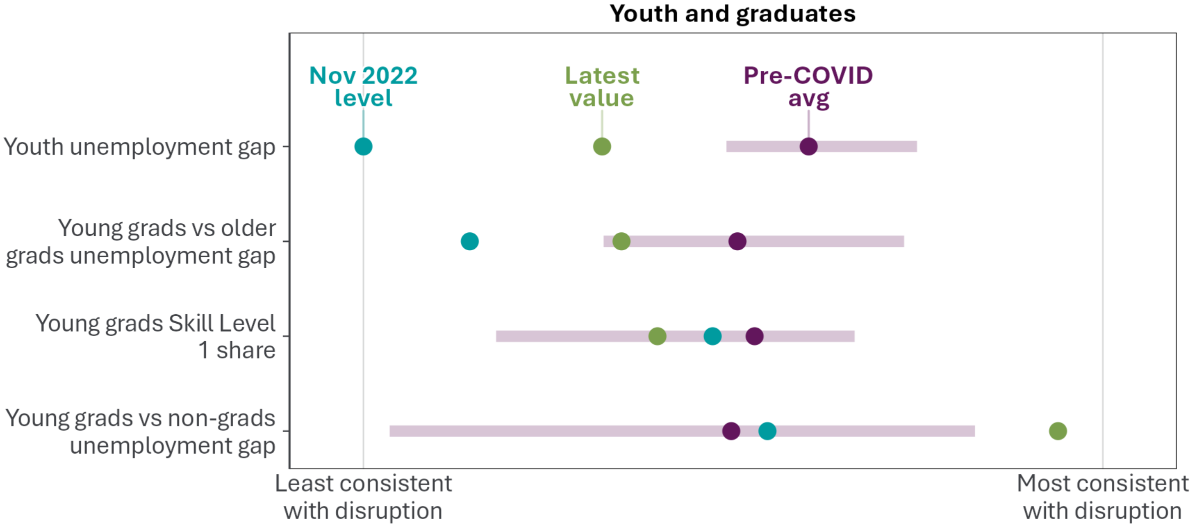

Notes: 'youth unemployment gap' calculated using 3-month moving average of seasonally adjusted monthly data. Other measures are 4-quarter moving averages of seasonally unadjusted quarterly data.
Sources: ABS Labour Force and DEWR calculations.

Occupational reshuffling has not accelerated

If AI starts affecting some segments of the labour market more than others, this is likely to result in an increased pace of compositional change in the labour market. Some occupations would shrink, others may grow, and the rate of this change in the occupational profile would be faster than normal.[46] But we find that the pace of compositional change in the labour market has not accelerated and is slightly below its pre-ChatGPT level.

We measure this pace of compositional change using a dissimilarity index.[47] This metric has been used to monitor potential AI-driven disruption in the labour market in the US,[48] and to measure occupational mismatch between available jobs and job seekers in Australia.[49] Appendix D contains further details on how the index is constructed.

The index ranges from 0 to 100.[50] A higher number means faster compositional change. It measures the share of employment that would need to shift across occupations to match the previous quarter's occupational mix. For example, if 3% of workers would need to move between occupations for this quarter's occupational mix to match last quarter's, the dissimilarity index is 3.

Tracking compositional change in the labour market is a useful cross-check on analyses based on JSA exposure scores – such as in the next section and in our core statistical model – because it does not require assumptions about which occupations are most exposed to AI. The compositional change index measures differences in the mix of occupations that Australians are employed in from quarter-to-quarter, regardless of which specific occupations happen to be growing or shrinking.

Compositional change in the Australian labour market has slowed since late 2022, as shown in Figure 13. The compositional change index has fallen from 5.9% in the November 2022 quarter to 5.5% in the February 2026 quarter, using data on employment by detailed occupations. Note that the pace of

---

[46] While the rate of compositional change is an informative indicator, it should be noted that AI-driven occupational change may be difficult to disentangle from other labour market trends.
[47] A 'Duncan and Duncan' index (Duncan and Duncan 1955).
[48] See Gimbel et al. (2025). Deming et al. (2025) also do a similar exercise.
[49] See Fahrer and Pease (1993) and SEEK (2024)
[50] Duncan and Duncan (1955) indexes typically run from 0 to 1. We multiply by 100 for ease of description.

compositional change has trended upwards over the longer term, so a modest increase of the index from current levels could still be consistent with this longer-run trend. The index calculated with broader (2-digit ANZSCO) occupational categories is also broadly in line with its typical level in the years before ChatGPT was introduced.

Figure 13: The pace of change in occupations has fallen since 2022

Change in occupational composition of employment (share of employment that would need to move between occupations to match previous quarter, per cent)

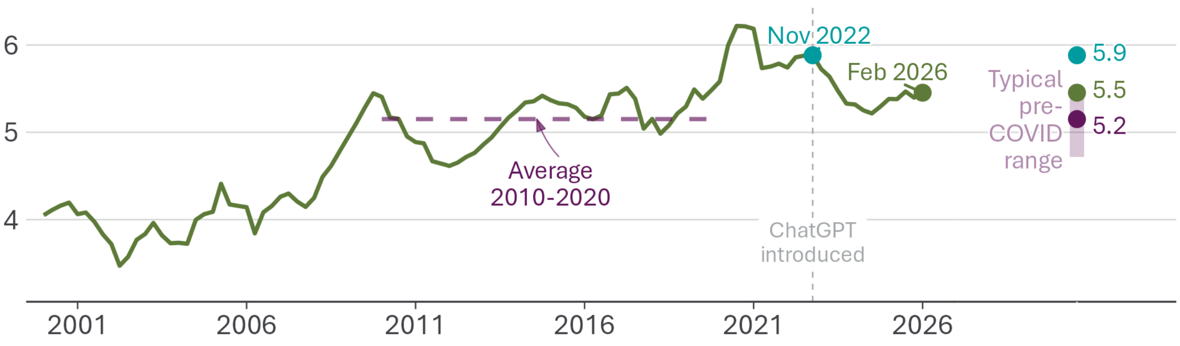

Note: Calculated using a Duncan dissimilarity index of occupations' employment shares compared to the prior quarter. 4-quarter moving average of non-seasonally adjusted data. Smaller values of this indicator are less consistent with labour market disruption. Pre-COVID average uses data from February 2010 to February 2020. 'Typical' range depicts the 10th to 90th percentiles.
Sources: ABS Labour Force (Detailed) and DEWR calculations.

The pace of compositional change for young people is also below its 2022 level and within its typical pre-COVID range, as shown in Figure 14. If AI-driven disruption were showing up first among younger workers, we would expect faster reshuffling in this group. We do not see that pattern. Compositional change for prime-age and older workers is also around similar levels as in the years before late 2022.

Figure 14: The pace of occupational change for young people has fallen since 2022

Change in occupational composition of employment, by broad age group (share of employment that would need to move between occupations to match previous quarter, per cent)

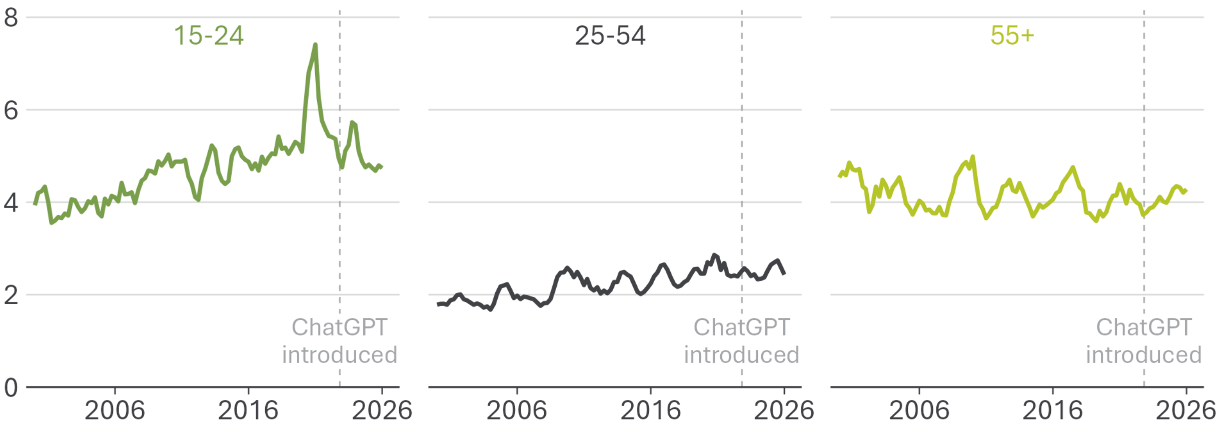

Note: Calculated using a Duncan dissimilarity index of 2-digit occupations' employment shares compared to the prior quarter. 4-quarter moving average of seasonally unadjusted data.
Sources: ABS Labour Force (Detailed) and DEWR calculations.

Figure 15 summarises our metrics of compositional change, with higher rates of compositional change being more consistent with disruption. It should be noted, however, that faster compositional change can also have positive effects.[51]

Figure 15: The pace of compositional change in the labour market is within its pre-COVID normal range

Duncan index measuring the pace of change in the occupational composition of employment

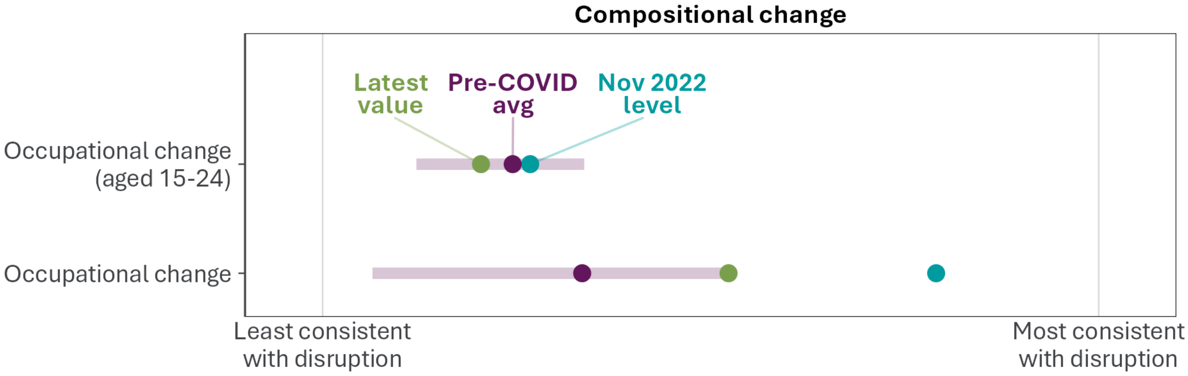

Notes: See Appendix for detailed methodology.
Sources: ABS Labour Force (Detailed) and DEWR calculations.

Conditions have softened in some high-exposure occupations

Key points

- Overall, jobs that are more exposed to automation by AI have seen employment, hours and job advertisements grow more slowly than less-exposed occupations since November 2022.
- However, total employment in the most-exposed occupations has risen by 5.6% in the period since November 2022, while employment in the least-exposed occupations has grown by 9.5%.
- Not all highly exposed occupations have grown slowly. For example, software developer employment is up in Australia in recent years. The ‘Software and Applications Programmers’ occupation employed 199,000 people in February 2026 – up 25% since November 2022.
- The most-exposed occupations' share of total employment has declined modestly since ChatGPT was introduced. However, this partly reflects the continuation of pre-AI trends.
- The unemployment rate for the occupations most exposed to AI has risen more than other occupations. As a result, the gap between the unemployment rate for the most exposed and least exposed occupations has shrunk substantially.

In this section, we examine trends across occupations with different levels of exposure to automation from generative AI. We build on this descriptive analysis with more formal statistical modelling seen in the next chapter of this report.

[51] Labour mobility and churn can enhance productivity and wages – see for example Andrews and Hansell (2021) and Buckley et al. (2024).

This section uses the AI automation exposure scores developed by JSA.[52] See the Measuring AI exposure section at page 11 for an overview of these scores. Further information, including a comparison with alternative exposure measures, can be found in Appendix A at page 63.

Trends in individual occupations

There is no common employment trend across all occupations most exposed to potential automation by AI, as shown in Figure 16. Some highly exposed occupations have seen a slight decline in employment, while others are up. In the 6 most-exposed occupations, the data shows:

- The number of people employed as 'Human Resource Clerks' is above its November 2022 level, after a period of volatility.
- 'Tourism and Travel Advisers' employment has increased after a COVID-era decline and is above November 2022 levels.
- Employment of 'Telemarketers' is also slightly higher as at February 2026 than in November 2022. This is despite a long-run structural decline in this occupation.
- Employment in the 'Keyboard Operators' and 'Filing and Registry Clerks' occupations is around November 2022 levels.
- 'Call or Contact Centre Workers' employment is below November 2022 levels.

Figure 16: There is no clear common employment trend in the most-exposed occupations
Employment in the occupations most exposed to automation by AI (thousands)

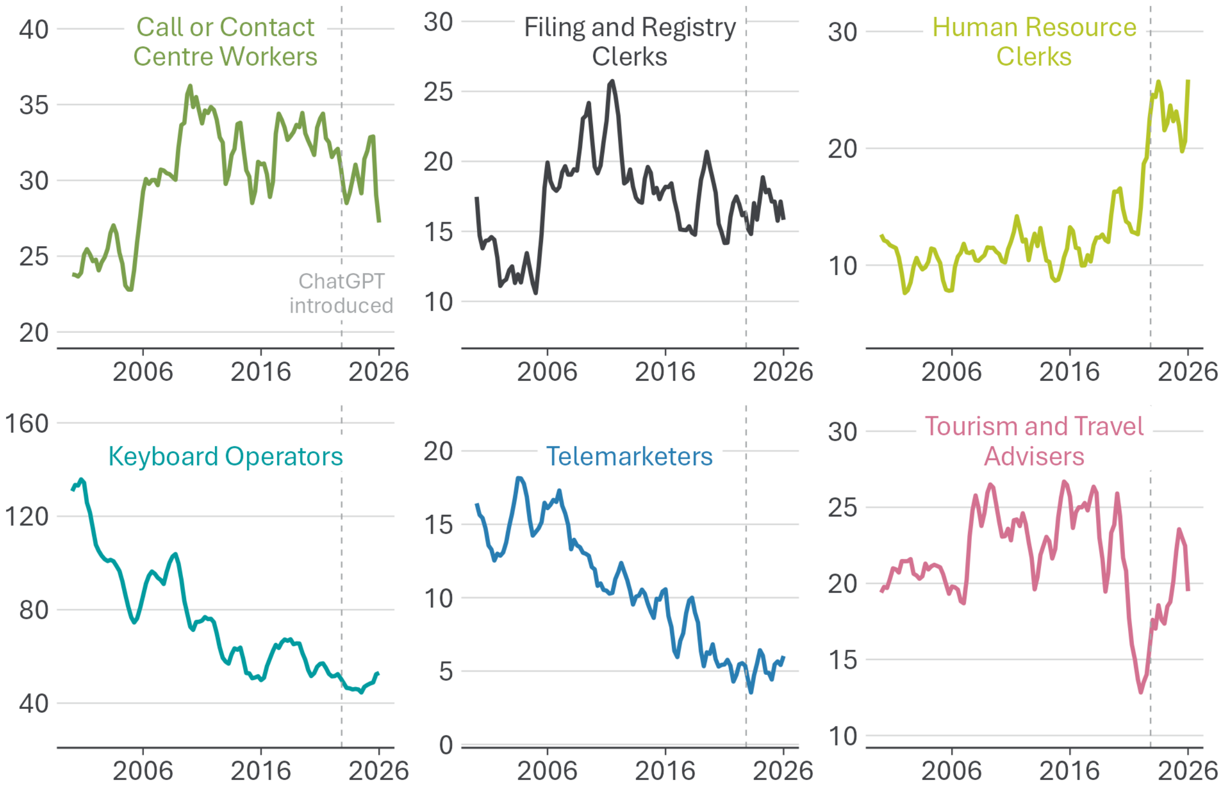

Note: 4 quarter moving average of seasonally unadjusted data.
Source: ABS Labour Force (Detailed).

[52] JSA (2025a)

Much media commentary has focused on employment of software developers, with job cuts at technology firms ascribed to the effect of AI.[53] But the much-discussed difficult labour market for software developers is not evident in the aggregate Australian employment data to date.

Figure 17 shows employment in 4 software developer occupations. The largest of these ('Software and Applications Programmers') employed 199,000 people in February 2026 – up 25% since November 2022. There are 40,000 more people in Australia employed as 'Software and Applications Programmers' than in the quarter before ChatGPT was introduced. Employment in the 'Database Administrators' occupation is also up modestly.[54] 'ICT Business and Systems Analysts' employment is down, but only slightly. The smallest of these occupations ('Multimedia Specialists and Web Developers') has seen a larger proportionate fall in employment but has returned to pre-COVID levels after a surge in 2021. Combined, employment in these occupations has continued to grow at a solid pace since the introduction of ChatGPT.

[53] For example, see Ibrahim (2026).
[54] The full name of the occupation is 'Database and Systems Administrators, and ICT Security Specialists'.

Figure 17: There are more people employed in software-related occupations in 2026 than in 2022

Employment in four ICT professional occupations (thousands)

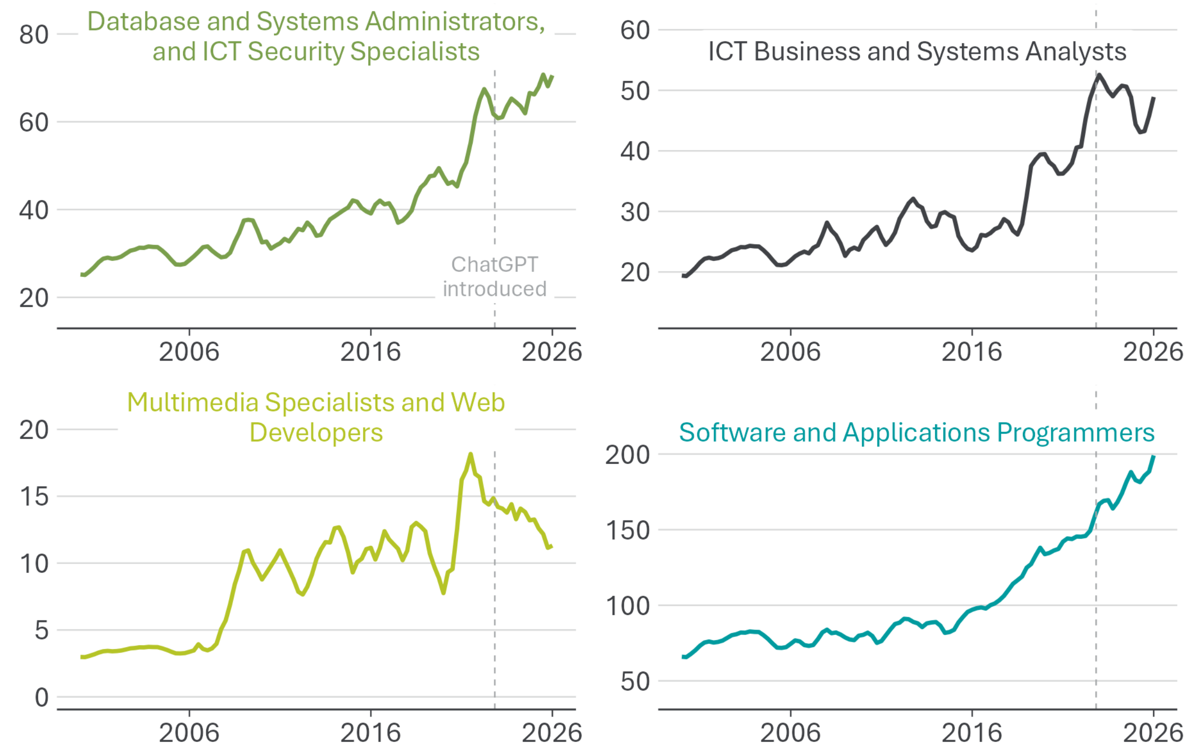

Combined employment in four ICT professional occupations (thousands)

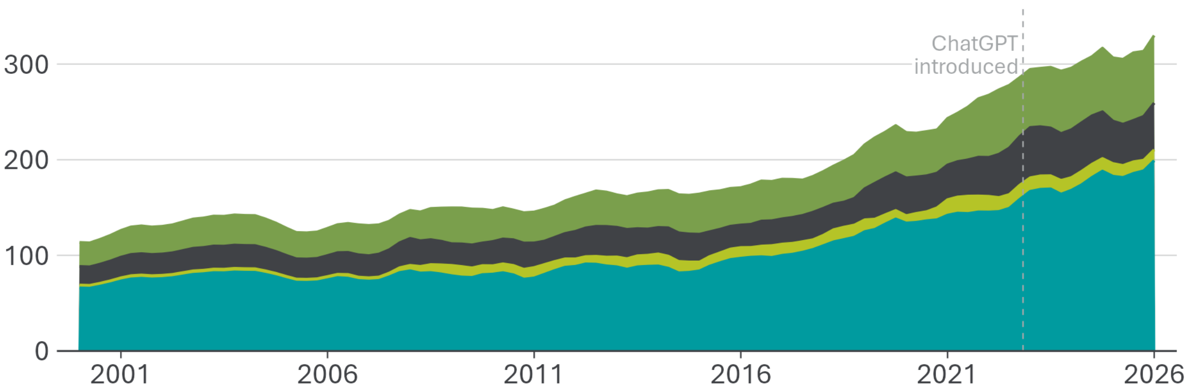

Note: 4 quarter moving average of seasonally unadjusted data.
Source: ABS Labour Force (Detailed).

Employment at the individual occupation level can be volatile for reasons unrelated to AI, including sampling error and idiosyncratic shocks. We can reduce the effect of this volatility by grouping occupations together by their level of AI exposure.

Employment by exposure quintile

Some highly exposed occupations have experienced consistent employment growth in the post-ChatGPT period. However, total employment in the most exposed fifth of occupations has grown more slowly than employment in less-exposed occupations.

We group occupations into 5 groups ('quintiles'), based on their JSA AI automation exposure scores.[55] These quintiles differ in their occupational composition and pre-existing growth trends. The

[55] These quintiles are not weighted by employment size, so they do not contain an equal number of workers.

comparison between them should therefore be interpreted with caution. Slower growth in more-exposed occupations is consistent with AI-related softening but may also reflect continuation of earlier occupational trends or other sector-specific developments.

Since November 2022, employment in the most-exposed quintile of occupations has grown by 5.6%. Over the same period, employment in the least-exposed occupations has grown by 9.5%.

Total employment in the most-exposed occupations has grown more slowly than all other quintiles over the period since November 2022, as shown in Figure 18. All of this sluggishness in employment in the most-exposed occupations has occurred since the beginning of 2024. Although employment in the most-exposed occupations is still above its pre-ChatGPT level, it is slightly below its February 2024 level, despite continued growth in other occupation groups. While the timing is consistent with the introduction of generative AI, other developments could also have influenced these trends. More detailed analysis is contained in the statistical model chapter.

Figure 18: Employment in the most-exposed occupations has lagged behind less-exposed occupations

Employment by quintile of AI exposure (index, Nov '22 = 100)

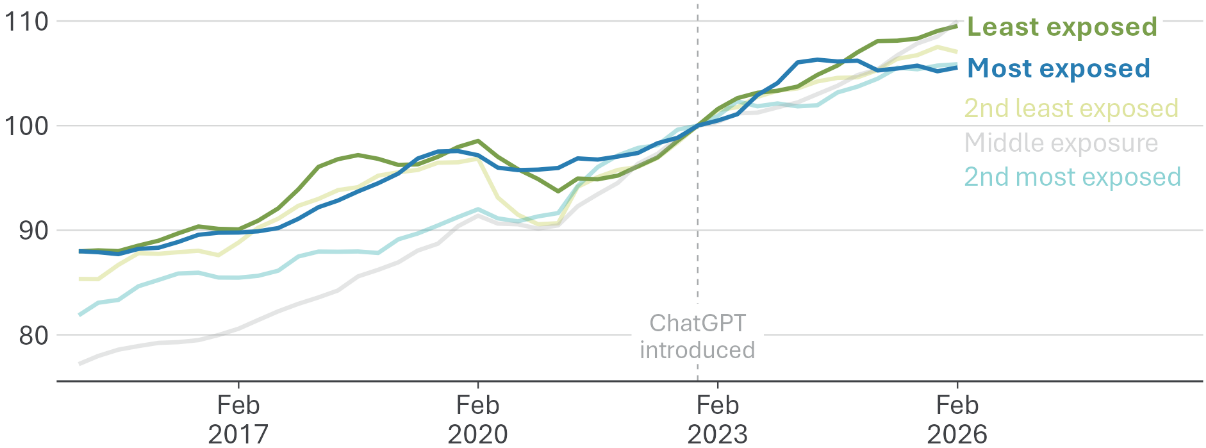

Notes: 4 quarter rolling average of seasonally unadjusted data. Each exposure quintile consists of an equal number of ANZSCO occupations, not weighted by employment size.
Sources: ABS Labour Force (Detailed) and DEWR calculations. GenAI automation exposure scores from JSA (2025a).

We focus on how the gap between outcomes for the most- and least-exposed occupations has changed over time. The employment levels from Figure 18 are converted to a growth rate, and we look at the difference in growth rates between the top and bottom exposure quintiles in Figure 19. The sign of this gap has changed frequently, with periods of faster growth for more exposed occupations (positive gap) and periods which favour the least exposed occupations (negative gap). The latest data favours the least exposed occupations, but the employment growth gap is within its typical range and is moving towards neutral.

Figure 19: High-exposure occupations are growing more slowly than low-exposure occupations, but the gap is within its typical range

Difference in employment growth between most- and least-exposed occupations

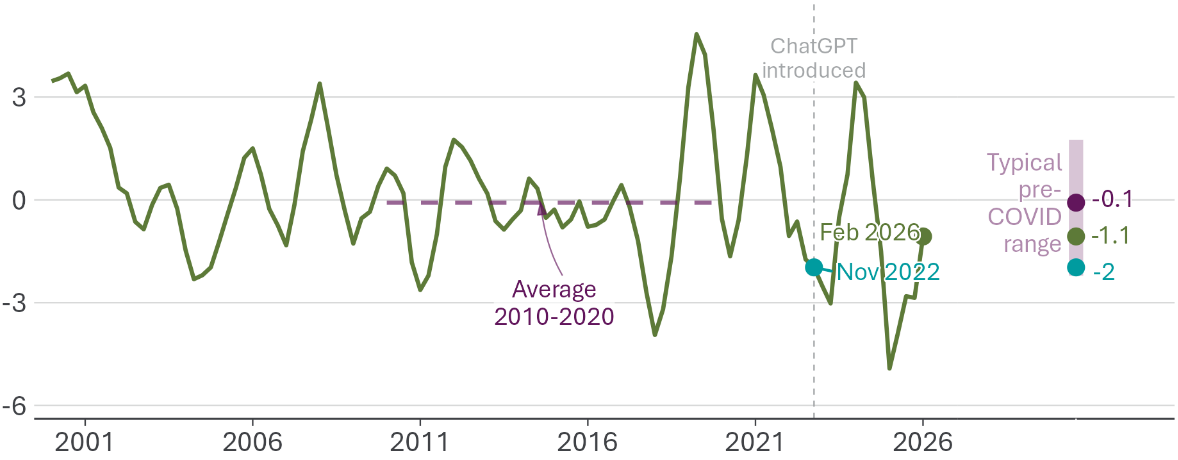

Note: 4 quarter moving average of seasonally unadjusted data. Gap is most exposed minus least exposed.
Sources: ABS Labour Force (Detailed) and DEWR calculations. GenAI automation exposure scores from JSA (2025a).

Headcount employment can be slower to adjust to declines in labour demand than hours worked, due to hiring and firing frictions. Changes in relative demand for different occupations may therefore show up in hours worked before becoming apparent in employment. To assess this, we examine the growth rate in total hours worked and focus on the difference between the most exposed and least exposed occupations (Figure 20). Similar to the results for headcount employment, since November 2022 we have seen hours growth switch from favouring least exposed, to most exposed, and back to least exposed. The latest estimate of the hours gap is negative, and slightly outside its typical pre-COVID range.

Figure 20: Hours worked in high-exposure occupations have grown more slowly since ChatGPT

Difference in worked hours growth between most exposed and least exposed occupations

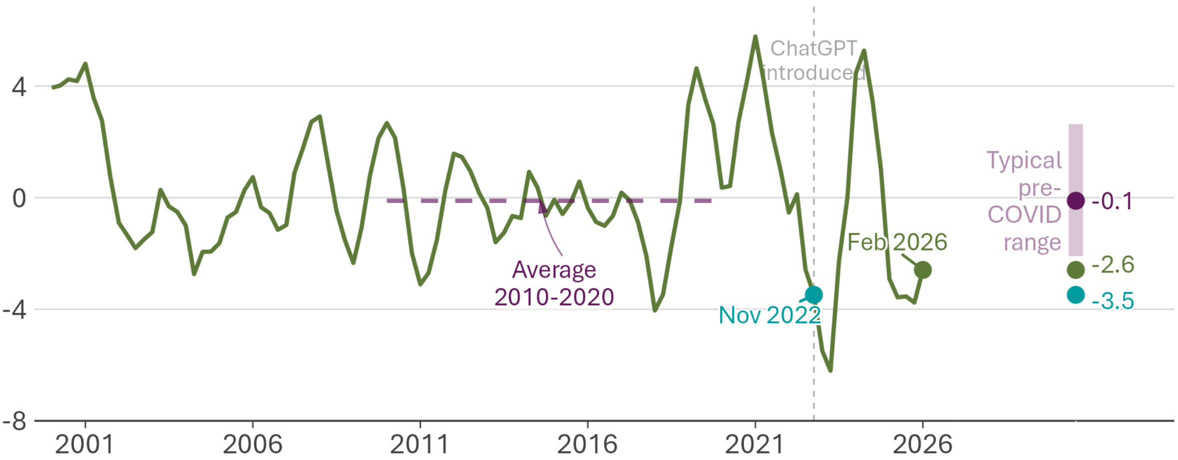

Note: 4 quarter moving average of seasonally unadjusted data. Gap is most exposed minus least exposed.
Sources: ABS Labour Force (Detailed) and DEWR calculations. GenAI automation exposure scores from JSA (2025a).

Growth rates and indices can be useful, but they can also hide differences in the size of each group. For this reason, we also look at the number of people working in occupations in each exposure quintile, expressed as a share of total employment. The exposure quintiles contain equal numbers of occupations, but occupations themselves vary greatly in size, and their employment levels change over time. As a result, the employment shares of the quintiles are not equal or fixed.

The most exposed quintile accounts for more than a quarter of total employment, although its share has been gradually declining. This decline was already under way before the introduction of ChatGPT and has continued at a similar pace since then, falling from 28.7% in February 2015 to 27.0% as at February 2026 (Figure 21).

Employment in the least exposed quintile was also declining as a share of total employment before ChatGPT but has partly recovered in recent years. By contrast, the middle exposure quintile increased its employment share from 19.2% to 21.4%. Taken together, these changes suggest employment shares have declined at both the highest and lowest ends of the exposure distribution, while the middle has expanded.

Figure 21: The employment share of the most-exposed occupations has been declining, before and after the introduction of ChatGPT

Percentage share of total employment by quintile of AI exposure

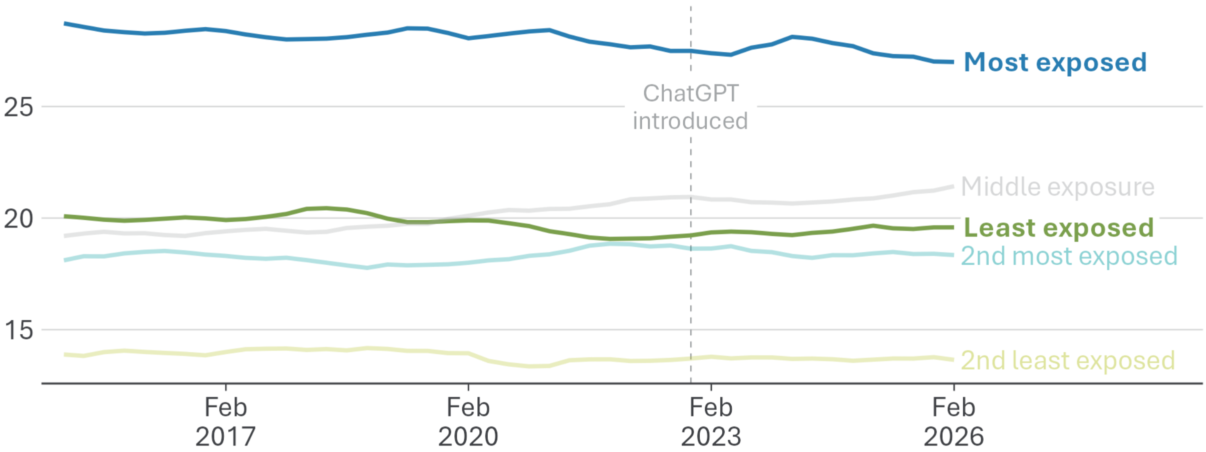

Notes: 4 quarter rolling average of seasonally unadjusted data. Each exposure quintile consists of an equal number of ANZSCO occupations, not weighted by employment size.
Sources: ABS Labour Force (Detailed) and DEWR calculations. GenAI automation exposure scores from JSA (2025a).

One way to summarise these changes is to examine the percentage point gap between employment shares in the most and least exposed quintiles (Figure 22). This measure captures the net movement at the two extremes of the exposure distribution. The gap is currently outside its typical pre-COVID range.

However, this descriptive analysis cannot tell us how much, if any, of the narrowing is related to the introduction of ChatGPT. Because the decline in the most exposed occupations has largely followed its pre-existing trend, it is difficult to distinguish any ChatGPT-related change from trends that were already in train. See the next chapter for more in-depth analysis of the exposure–employment relationship.

Figure 22: The gap between the share of most- and least-exposed occupations has narrowed

Difference in share of total employment between most- and least-exposed occupations

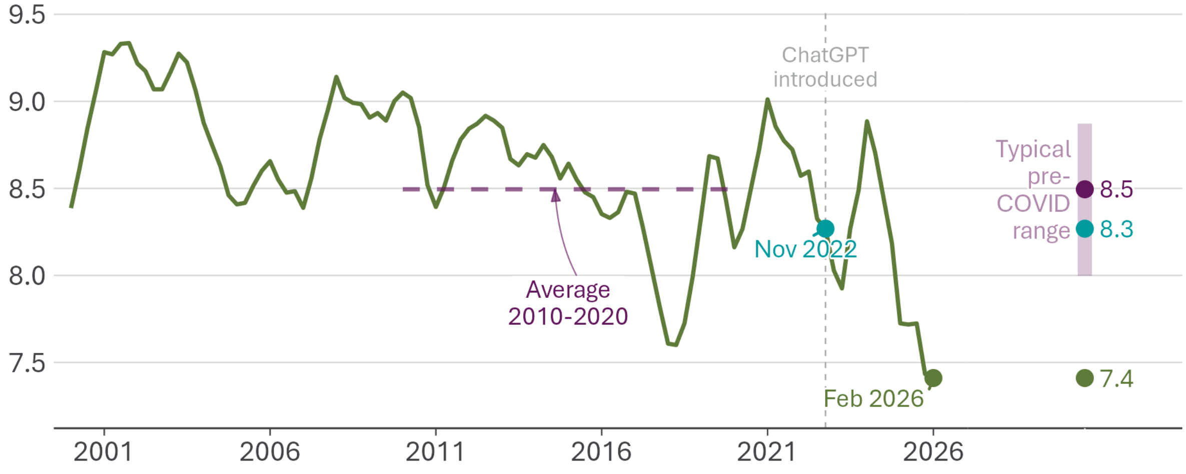

Note: 4 quarter moving average of seasonally unadjusted data. Gap is most-exposed share minus least-exposed share.
Sources: ABS Labour Force (Detailed) and DEWR calculations. GenAI automation exposure scores from JSA (2025a).

Vacancies by exposure quintile

Job advertisements can adjust faster than employment and hours worked and may therefore provide an early indication of future employment effects. We measure the number of job advertisements using the Internet Vacancy Index (IVI)[56] and group occupations into exposure quintiles using the JSA AI automation exposure scores.

In the period before November 2022, vacancies were rising sharply across all exposure quintiles (Figure 23). However, there has been a subsequent broad decline. All quintiles experienced declines of at least 37%. As much of the earlier vacancy growth may be a result of the post-COVID recovery hiring boom, it is not clear how much of the decline may be due to the impacts of AI or simply post-COVID normalisation. However, as at February 2026, vacancies for the most exposed occupations were similar to their pre-COVID levels, while the least exposed occupations are further above pre-COVID levels.

[56] JSA (2026b)

Figure 23: Job advertisements have fallen in all exposure groups, with high-exposure occupations back near pre-COVID levels

Internet Vacancy Index job ads by quintile of AI exposure (index, Nov '22 = 100)

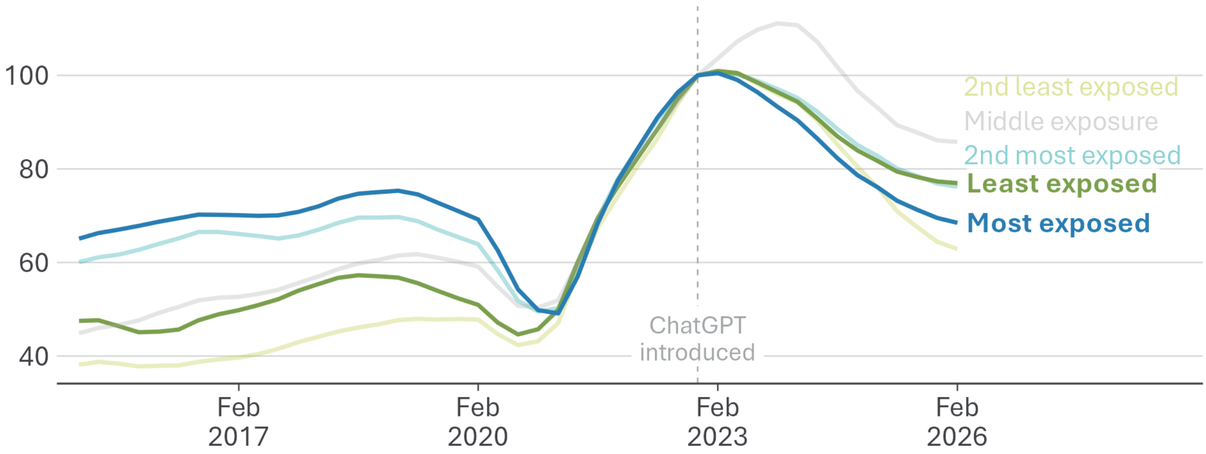

Notes: 4 quarter rolling average of seasonally unadjusted data. Each exposure quintile consists of an equal number of ANZSCO occupations, not weighted by employment size.
Sources: JSA IVI and DEWR calculations. GenAI automation exposure scores from JSA (2025a).

We convert the IVI data to growth rates and focus on the difference in growth between the most and least exposed occupations. Since November 2022 the job advertisements growth gap has been consistently negative (Figure 24). While the size of this gap is within the typical pre-COVID range, persistent gaps in growth can produce effects that accumulate over time. It is important to note that the statistical modelling in the next chapter suggests that this may continue a pre-ChatGPT trend of slower growth in job advertisements for the most exposed occupations.

Figure 24: Job advertisements in high-exposure occupations have fallen more than low-exposure occupations

Difference in job ads growth between most exposed and least exposed occupations

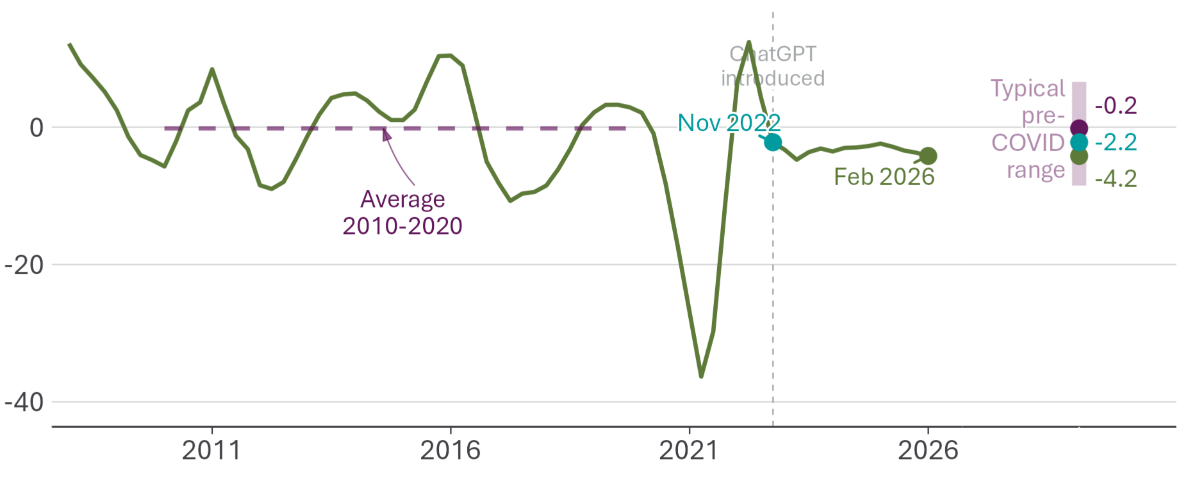

Note: 4 quarter moving average of seasonally unadjusted data. Gap is most exposed minus least exposed.
Sources: JSA IVI and DEWR calculations. GenAI automation exposure scores from JSA (2025a).

Unemployment by exposure quintile

As well as tracking employment trends for different occupational groups, we calculate unemployment rates for each quintile of AI exposure.[57]

Figure 25 shows that the unemployment rate for occupations most exposed to automation by AI has risen since late 2022, to a greater degree than other occupational groups. The unemployment rate for the most exposed occupations remains below that for the least exposed, but this gap has closed considerably.

Figure 25: Unemployment has risen more for the most exposed occupations

Unemployment rate by occupational AI-exposure quintile (per cent)

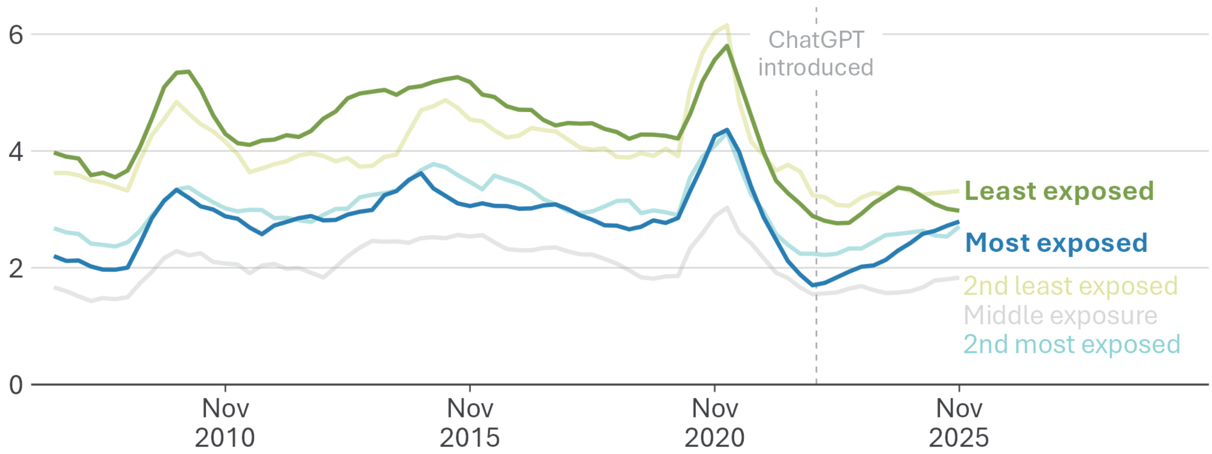

Note: Unemployed people are allocated to quintiles based on occupation of previous job. Does not include new labour force entrants or 'former workers'. Four-quarter average of non-seasonally adjusted data.
Sources: ABS Labour Force (TableBuilder) and DEWR calculations. GenAI automation exposure scores from JSA (2025a).

Figure 26 shows the gap between the unemployment rates for the most- and least-exposed fifth of occupations. In the pre-COVID period, this gap was typically between -1.5 and -2 percentage points but has been moving towards zero since around the start of the pandemic and is now at a level not seen in at least the last 20 years. The gap is still negative, meaning that unemployed workers in the most exposed occupations represent a smaller share of the corresponding workforce than for the least exposed workers. The recent narrowing of this gap is more due to an increasing unemployment rate for most exposed workers, rather than a reduction among least exposed workers.

It is worth noting that the unemployment rate gap between the most- and least-exposed occupations was already shrinking prior to the introduction of ChatGPT. In November 2022, the unemployment rate gap between the top and bottom exposure quintiles was -1.2 percentage points – outside the typical pre-COVID range. It is possible the further falls in the gap are, at least in part, a continuation of pre-AI trends rather than induced by AI.

[57] We classify unemployed people based on the AI automation exposure of their previous occupation. We derive an unemployment rate for each quintile by comparing the number of unemployed people previously employed in occupations in that quintile with the labour force associated with that quintile — that is, currently employed people in those occupations plus unemployed people whose previous occupation was in that quintile. This should be interpreted as an unemployment rate for workers previously attached to each exposure group, rather than a direct measure of unemployment in current occupations. It does not include workers who move successfully between occupations, new entrants without a previous occupation, or changes in the desired occupation of unemployed people.

Figure 26: The gap in unemployment rates between most and least exposed occupations has almost closed

Difference in unemployment rate between most exposed and least exposed occupations

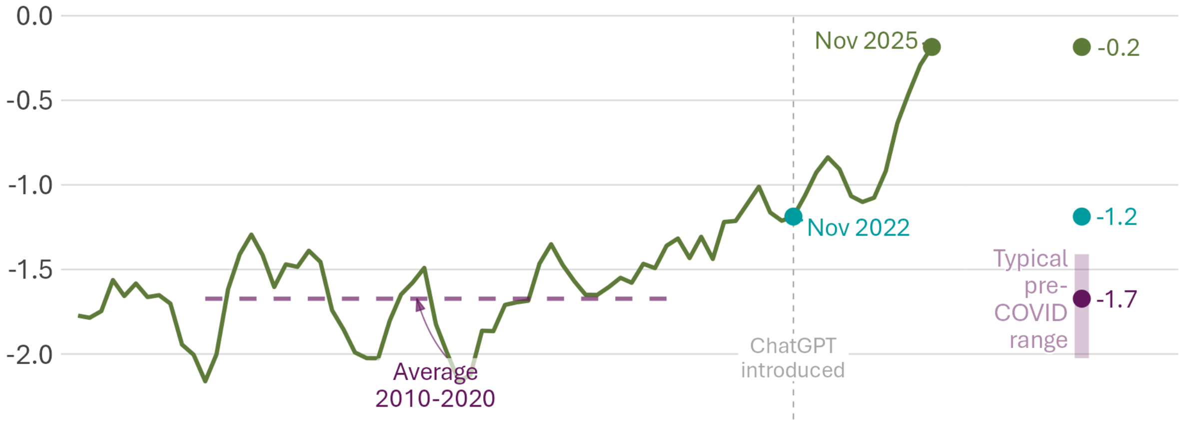

Note: 4 quarter moving average of seasonally unadjusted data. Gap is most exposed minus least exposed.
Source: ABS Labour Force (Detailed).

To better understand these trends in unemployment, we examine the ANZSCO major occupational groups to understand what may be contributing to the rising unemployment rate for the most exposed occupations in Figure 25.[58] Two groups which include a number of highly exposed occupations are Clerical and Administrative Workers and Professionals, both of which have seen their respective unemployment rates return to pre-COVID levels (Figure 27). In contrast, other occupation groups have had stable or more modest increases in unemployment rates since November 2022.

[58] Unemployment rates for individual four-digit ANZSCO occupations are highly volatile, so we focus on major groups.

Figure 27: The unemployment rates for Professionals and Clerical and Administrative Workers have risen to pre-COVID levels

Unemployment rate by occupation in last job (per cent)

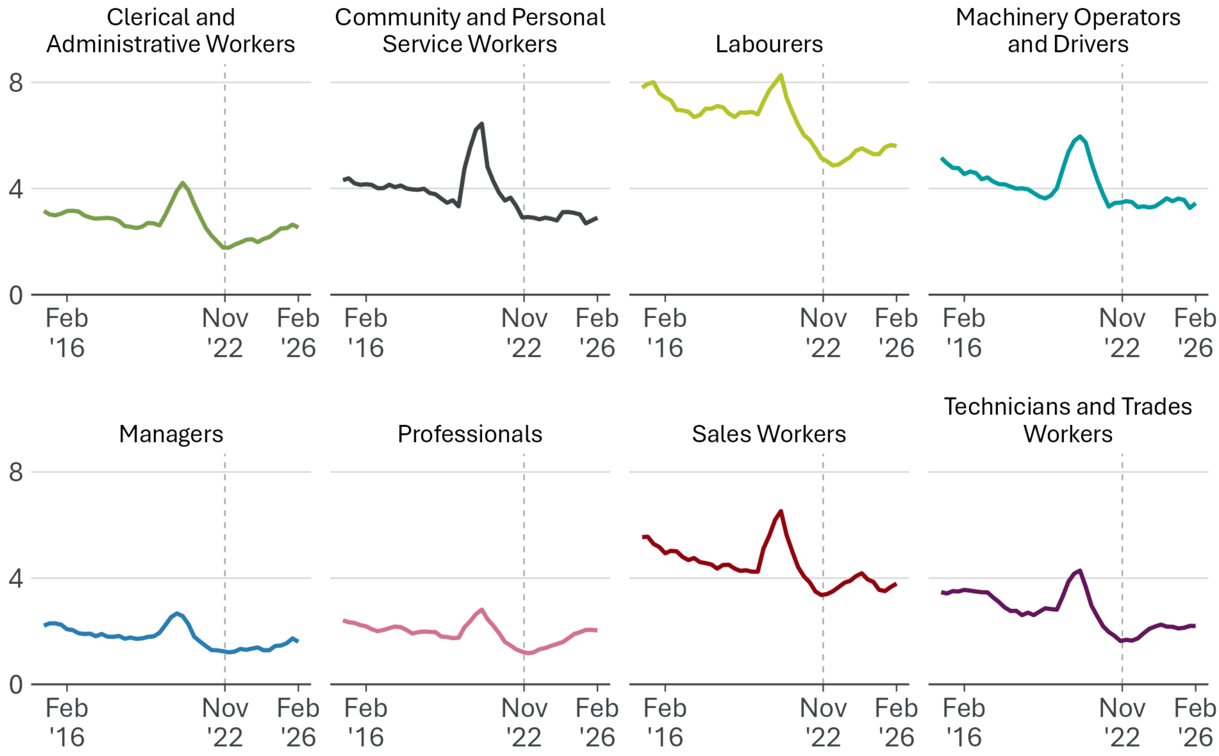

Note: 4 quarter rolling average of seasonally unadjusted data.
Sources: ABS Labour Force (Detailed) and DEWR calculations

Taken together, the data suggests a possible weakening of conditions for the most-exposed occupations, though there was evidence of some decline beginning well before the introduction of ChatGPT. Outcomes are varied – not all highly exposed occupations have grown more slowly in recent years – but in aggregate there is some sign of relative softening in highly exposed occupations relative to others.

A number of indicators comparing most- and least-exposed occupations are summarised in Figure 28. The most concerning signs come from the gap in share of total employment and the gap in unemployment rate being at the most extreme values since 2010. The [statistical model] chapter explores in more detail how employment growth varies across occupations of different exposure scores, and whether this relationship has changed since November 2022.

Figure 28: More-exposed occupations are growing more slowly than less exposed occupations on some measures

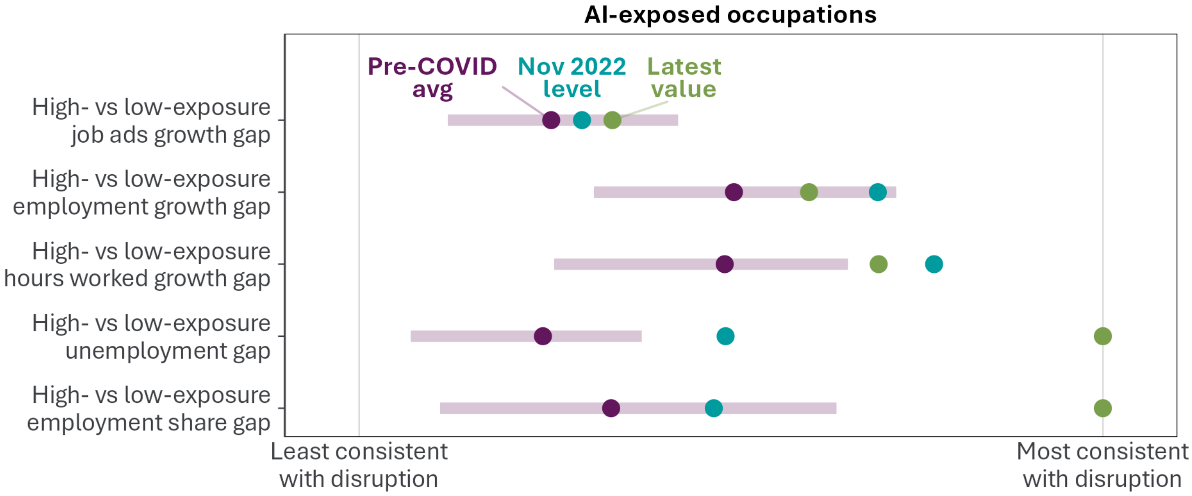

Notes: Four quarter average of seasonally unadjusted data. Gap is most exposed minus least exposed.
Sources: ABS Labour Force (Detailed), JSA IVI and DEWR calculations. GenAI automation exposure scores from JSA (2025a).

Routine cognitive jobs were declining prior to the release of ChatGPT

A widely used framework in labour economics classifies tasks and occupations according to whether they are 'routine' or 'non-routine' on one dimension, and 'manual' or 'cognitive' on another. This framework sorts tasks and jobs into 4 groups.[59] We use this 'tasks framework' to group occupations into 4 categories:

- Routine cognitive (such as receptionists and keyboard operators)
- Non-routine cognitive (such as managers and economists)
- Routine manual (such as plumbers and forklift drivers)
- Non-routine manual (such as cooks and sportspersons).[60]

Routine jobs are generally thought to be more exposed to automation, as the tasks in these jobs are more predictable and codifiable. Over the last 4 decades in Australia and elsewhere, there has been a shift away from routine jobs towards non-routine jobs, with a particularly large rise in the share of non-routine cognitive jobs.[61]

Cognitive jobs have higher AI automation exposure scores on average than manual jobs, as shown in Figure 29. Routine cognitive jobs are more exposed than non-routine cognitive jobs.

[59] See Autor et al. (2003), Autor (2013) and Acemoglu and Autor (2011).
[60] We use the framework as translated to ANZSCO by Borland and Coelli (2022).
[61] Borland and Coelli (2022)

Figure 29: Routine cognitive jobs have the highest average exposure to AI automation

JSA AI automation exposure scores by job category

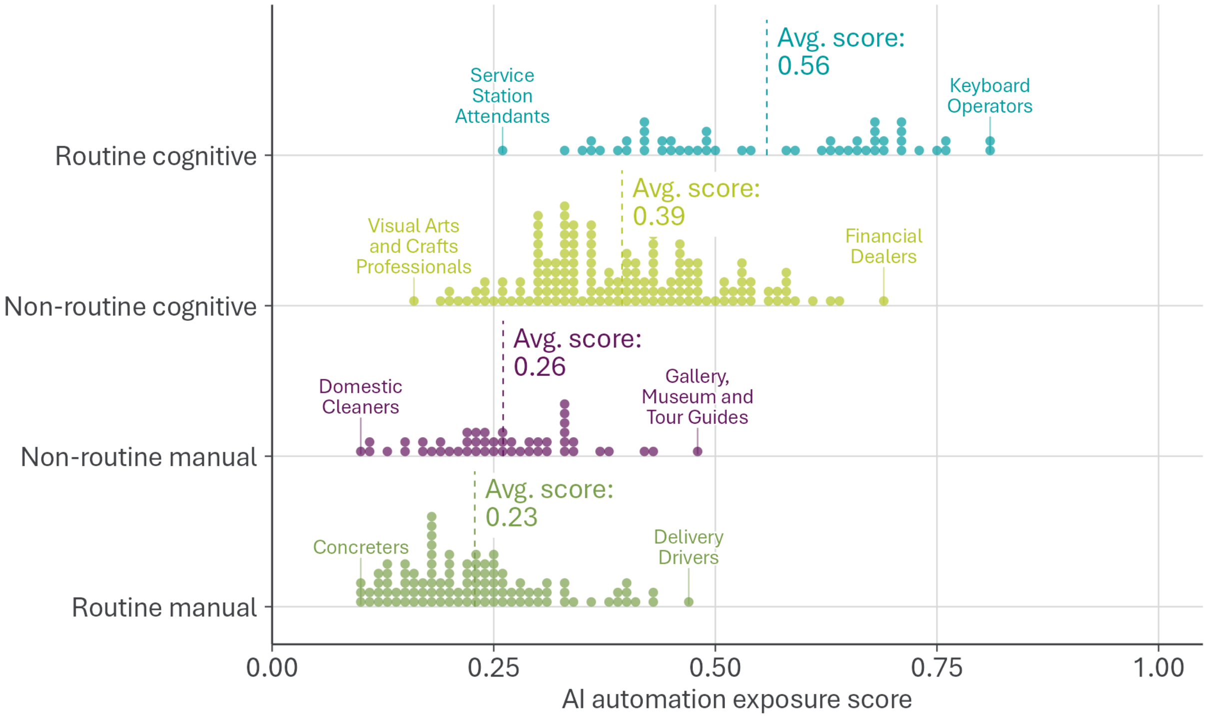

Note: Borland and Coelli classification of job category used with the exception of Tourism and Travel Advisers (assigned to 'routine cognitive') and Hairdressers ('non-routine manual').
Sources: Exposure scores from JSA (2025a). Classification of occupations into routine/non-routine and cognitive/manual categories from Borland and Coelli (2022).

Employment growth by routine and cognitive dimensions

Recent employment growth has been higher for non-routine jobs relative to routine jobs, shown in Figure 30. To a lesser extent, recent growth has favoured manual jobs over cognitive jobs. Routine cognitive jobs have grown the least since November 2022, though employment in routine cognitive occupations remains above its late-2022 level.

Figure 30: Employment growth of non-routine cognitive jobs is similar to pre-ChatGPT trends

Employment by routine/non-routine and cognitive/manual dimensions (index, Nov 2022 = 100)

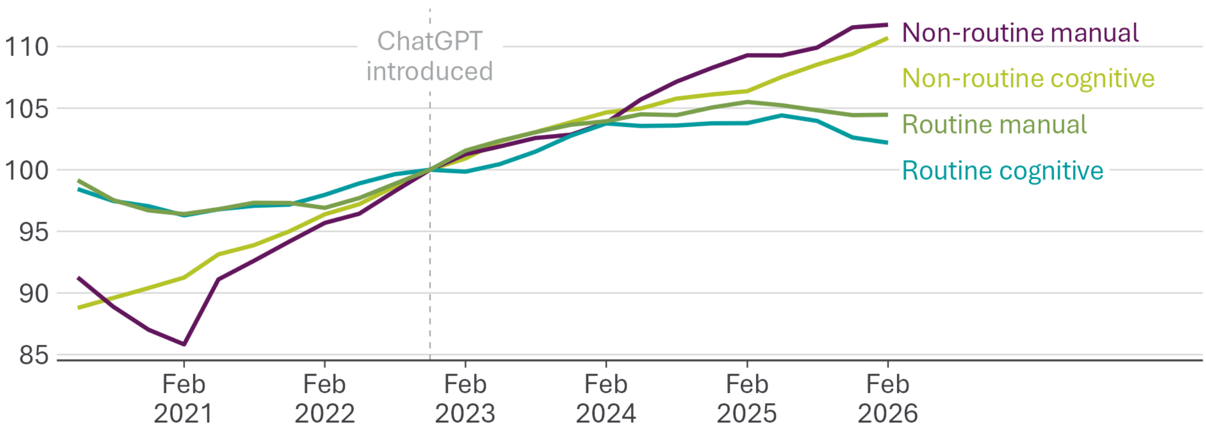

Notes: 4 quarter rolling average of seasonally unadjusted data. Borland and Coelli classification of job category used with the exception of Tourism and Travel Advisers (assigned to 'routine cognitive') and Hairdressers ('non-routine manual').
Sources: ABS Labour Force (Detailed) and DEWR calculations. Classification of ANZSCO occupations by routine/non-routine and cognitive/manual adopted from Borland and Coelli (2022).

Although employment in routine cognitive jobs has lagged other job types since ChatGPT was introduced, this appears broadly consistent with long-run trends documented by Borland and Coelli (2022). Figure 31 shows that the share of routine cognitive jobs has continued to decline broadly in line with pre-ChatGPT trends.

Figure 31: Routine cognitive jobs continue to shrink as share of total employment

Percentage of jobs in routine cognitive occupations

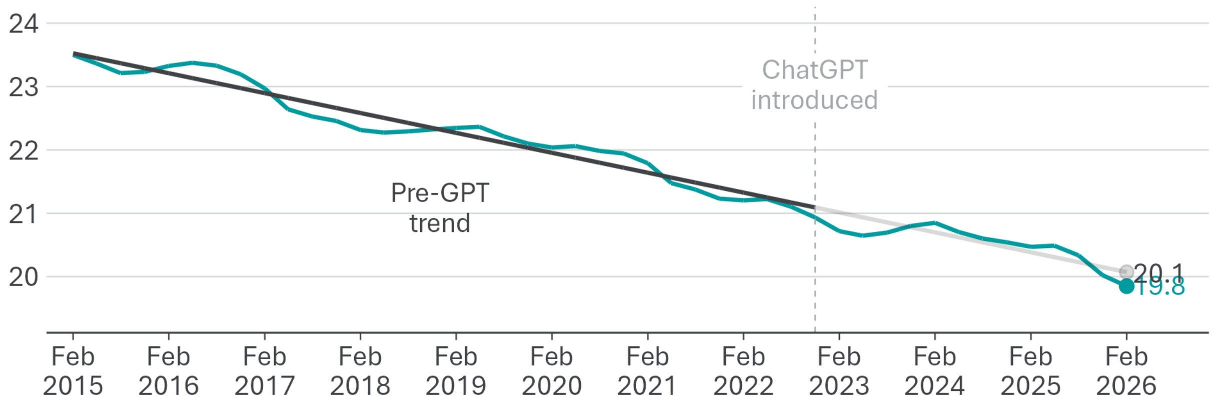

Notes: 4 quarter rolling average of seasonally unadjusted data. Note: Borland and Coelli classification of job category used with the exception of Tourism and Travel Advisers (assigned to 'routine cognitive') and Hairdressers ('non-routine manual'). Total excludes jobs not classified in Borland and Coelli.
Sources: ABS Labour Force (Detailed) and DEWR calculations. Classification of ANZSCO occupations by routine/ non-routine and cognitive/ manual adopted from Borland and Coelli (2022).

Routine cognitive occupations were already declining as a share of jobs prior to ChatGPT. This means the trends examined in the previous section should be interpreted cautiously – they do not provide definitive evidence of a change in the composition of employment away from more-exposed occupations. The statistical modelling in the next chapter is designed to test more directly whether the relationship between AI exposure and employment growth has changed relative to pre-ChatGPT trends.

Summary of descriptive indicators

The descriptive evidence is mixed but mostly shows ongoing labour market resilience. Aggregate labour market conditions remain relatively strong, youth outcomes have mostly held up and occupational reshuffling has not accelerated. However, high-exposure occupations have seen employment, hours and vacancies grow more slowly than in less-exposed occupations. Figure 32 provides an overview of the descriptive indicators considered in this chapter.

Figure 32: Indicators are mixed, but mostly do not suggest major labour market upheaval

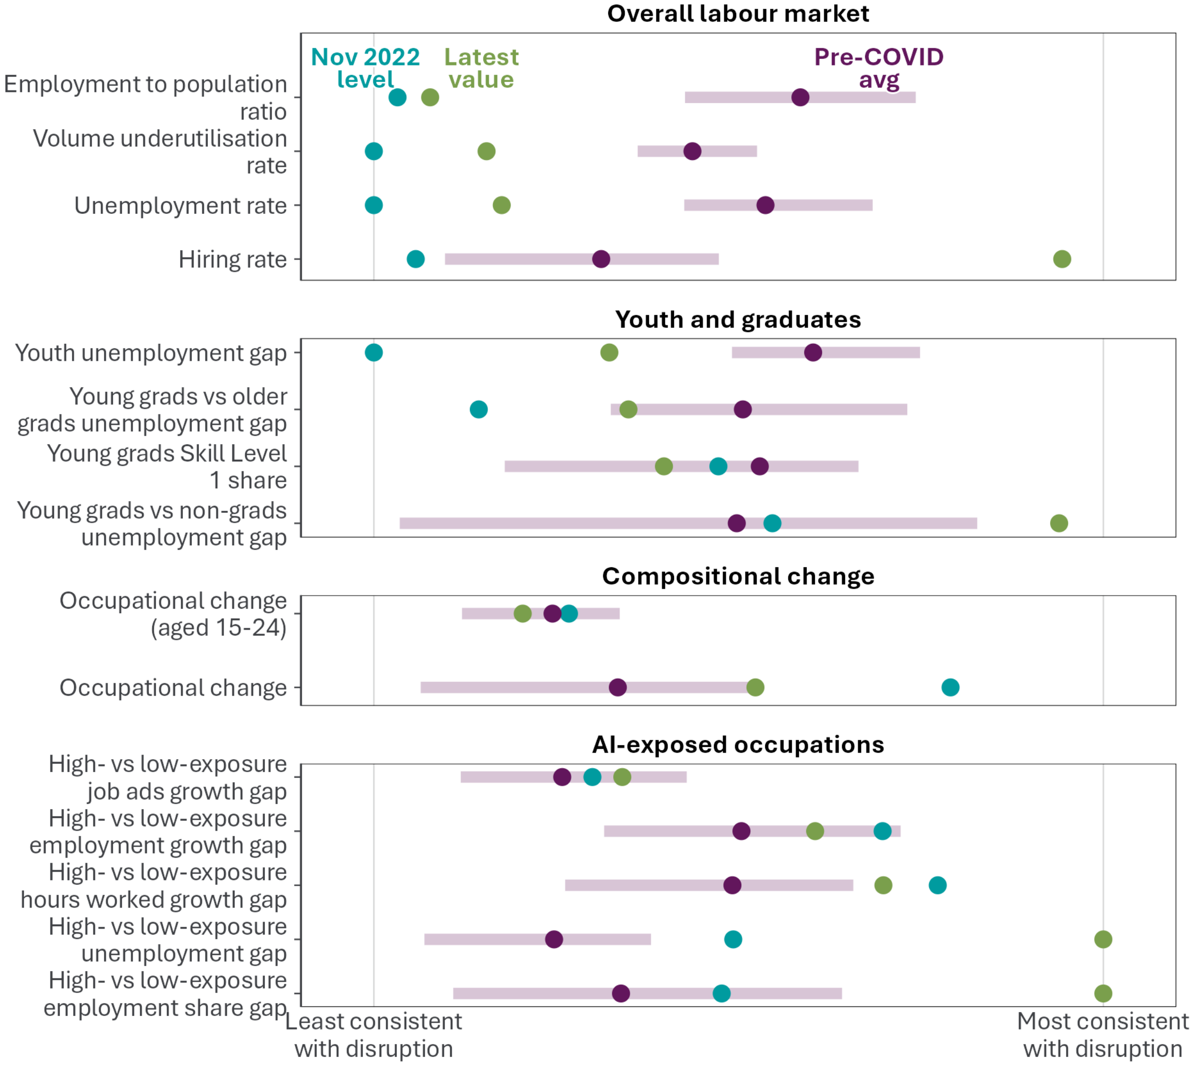

Note: Scales and derivations vary across indicators.
Sources: ABS Labour Force, ABS Labour Force (Detailed), ABS Labour Force (TableBuilder), JSA IVI and DEWR calculations. GenAI automation exposure scores from JSA (2025a).

The indicators we examine using the JSA automation exposure scores suggest that the labour market may have softened for the most exposed occupations. Growth in employment for more-exposed occupations has stalled since 2024 while less-exposed occupations have grown. Job vacancies for the most-exposed occupations have fallen below November 2022 levels, and the unemployment rate for most exposed workers is rising, closing the gap with the rate for less-exposed workers.

It is important to bear in mind, however, that these developments could have been influenced by a range of factors. In particular, comparisons between high- and low-exposure measures are just as sensitive to changes for low-exposure occupations, which may be unrelated to AI technologies, as they are to changes in high-exposure occupations. The next chapter examines more carefully the extent to which impacts have varied by level of automation exposure using statistical models designed to capture this relationship in a more complete and rigorous way.

Core statistical model

Key points

- Our statistical modelling suggests employment has grown more slowly in more AI-exposed occupations since the introduction of ChatGPT, but the signal is weak and the evidence is mixed.
- Our model estimates the relationship between occupations' exposure to AI automation and employment growth, both before and after November 2022. Before ChatGPT was introduced, there was no statistically significant relationship between AI exposure and employment growth across occupations in the model.
- Since November 2022, we find a modest negative relationship between AI exposure and employment growth. Our results suggest that occupations with AI automation exposure one standard deviation above the mean had employment about 2% lower at February 2026 than would be expected based on the pre-ChatGPT exposure–employment relationship. This does not mean that employment has fallen in high-exposure occupations - it means that employment is lower than we estimate would have been the case if pre-ChatGPT trends had continued.
- We also estimate a version of the model that examines the exposure–employment relationship quarter-by-quarter (an 'event study'). It shows the negative relationship between AI exposure and employment is clearer in recent quarters than immediately after ChatGPT was released. This suggests negative effects may be accumulating over time, justifying ongoing monitoring.
- The model's results are tentative, not definitive. Robustness tests in the next chapter assess how much weight to place on the results.

Although the labour market has been resilient, the data shows a softening in high-exposure occupations, with employment growing more slowly than in low-exposure occupations. But the descriptive indicators examined in the previous chapter are a coarse way of looking at the data, grouping together occupations in quintiles and not controlling for other factors that might affect employment. To go beyond these descriptive indicators, we turn to statistical modelling.

We estimate the relationship between occupations' AI exposure and employment growth in Australia, and whether the introduction of ChatGPT in November 2022 led to a change in this relationship. Our modelling results suggest that employment has grown more slowly in more AI-exposed occupations since ChatGPT was released. This is consistent with a small negative AI effect but is not conclusive causal evidence: as outlined in the following chapter, the findings are not consistent across all model variations.

This chapter presents the core statistical model, which is the central test of whether employment growth shifted away from more AI-exposed occupations after ChatGPT. It also presents an event-study version of the model, which examines whether the exposure–employment relationship changed gradually over time or only in recent quarters. The next chapter tests how robust the core result is to alternative assumptions, exposure measures, outcomes and model specifications.

Data and scope

We use quarterly data on employment by occupation from the ABS Labour Force Survey.[62] The model uses data from the February 2015 quarter to the February 2026 quarter. We use data covering 355 occupations for 45 quarters.

ChatGPT was released to the public on 30 November 2022. We use this as the 'treatment date' - data up to and including the November 2022 quarter are 'pre-AI', while everything from the February 2023 quarter onwards is 'post-AI'.[63] There is nothing ChatGPT-specific to this analysis, other than using November 2022 to delineate the pre- and post-AI periods.

We use JSA's occupation-level AI automation exposure scores (see Appendix A). The scores are standardised to have a mean of 0 and standard deviation of 1, so above-average exposure occupations have a positive exposure score and below average exposure occupations have a negative score. These scores do not vary over time.

The core model

The model compares employment growth in more- and less-AI-exposed occupations before and after ChatGPT. It uses variation in exposure across occupations while controlling for each occupation's average size and shocks common to all occupations in a given quarter.[64]

In practice, the model asks a simple question: **did occupations with higher AI exposure start growing more slowly after ChatGPT than occupations with lower AI exposure, relative to their previous relationship?** For example, if high-exposure occupations were already growing more slowly than low-exposure occupations before the release of ChatGPT, the model does not treat that pre-existing gap as evidence of an AI effect. It looks for a change in that relationship after November 2022.

The model's main coefficient of interest ($\Delta\beta$) measures whether the relationship between AI exposure and employment growth changed after November 2022. A negative estimate means employment growth shifted away from more-exposed occupations, relative to the pre-ChatGPT relationship. A positive estimate would mean growth shifted towards more-exposed occupations. An estimate close to zero would mean the relationship between AI exposure and employment growth did not materially change after ChatGPT.

The model does not compare 'AI-affected' and 'not AI-affected' jobs. AI exposure is measured on a scale. The model therefore asks whether more-exposed occupations experienced larger employment changes than less-exposed occupations. Nor does the model show whether firms actually adopted AI, or whether AI caused any particular job loss. It estimates whether employment patterns changed in a way that is correlated with occupations' potential exposure to AI automation.

The model makes one important assumption. Apart from differences associated with AI exposure, it treats occupation growth as reflecting shocks or trends that affect all occupations equally. This assumption is unlikely to hold perfectly, so the results should be interpreted cautiously. We test the sensitivity of this assumption in the next chapter.

A more detailed and technical description of the model and estimation methodology is presented in Appendix B on page 67.

---

[62] ABS (2026a). Our core model uses the 4-digit ANZSCO unit group classification of occupations, sourced from Table EQ08 of the Labour Force Detailed publication. February 2026 was the final issue of this publication, as the outputs from the Labour Force Survey are being revised. Employment-by-occupation data will resume publication in late 2026, using the new Occupation Standard Classification for Australia (OSCA) (ABS 2024) rather than ANZSCO.
[63] Using November 2022 as a date of treatment (based on the timing of the public release of ChatGPT) has been used in other studies on the labour market impacts of AI. Examples include Brynjolfsson et al. (2025), Klein Teeselink (2025) and Gimbel et al. (2025).
[64] The unit of observation is an occupation in a quarter.

Main results: employment has grown more slowly in more AI-exposed occupations

We find employment has grown more slowly in more AI-exposed occupations since the release of ChatGPT. The estimated effect is small, and the evidence is not robust to every modelling choice, so we interpret it as tentative evidence of relative softening, rather than definitive proof of job loss.

The model's main coefficient is negative and statistically significant, as shown in Figure 33. In practical terms, the model finds that an occupation with one standard deviation above average AI exposure had employment about 2% lower by February 2026 than it would have had under its pre-ChatGPT trend. Detailed coefficient estimates are reported in Appendix B.

Figure 33: The core model finds a small, statistically significant negative relationship between occupations' AI exposure and employment growth

Main coefficient of interest in the core model

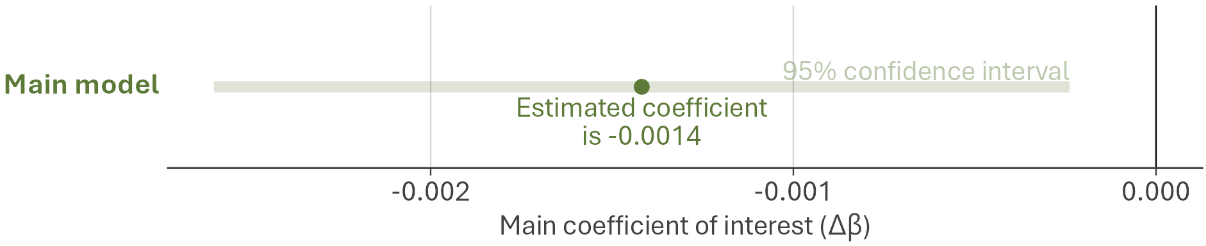

Note: Bars represent 95% confidence intervals. Bars which do not cross zero are at least statistically significant at the 5% level.
Sources: ABS Labour Force (Detailed), JSA (2025a) and DEWR calculations.

This result does not simply reflect the fact that some high-exposure occupations may already have been growing more slowly before ChatGPT. In the pre-ChatGPT period, there was no statistically significant relationship between JSA automation exposure and employment growth. The negative result comes from a change in the exposure–employment relationship after November 2022.

The estimated effect is modest. It is much smaller than some public predictions of large AI-driven job losses, and the confidence intervals only just exclude zero. The result is therefore best read as tentative evidence of a small relative shift away from more-exposed occupations, not as evidence of large-scale labour market disruption.

The event-study section below examines when this relationship appears in the data. The next chapter then tests how sensitive the result is to modelling choices and to the occupations included in the sample.

What the model does not show

The model does not directly observe AI adoption by firms, nor does it identify specific instances of job destruction or job creation. It estimates whether employment growth changed differently in occupations with higher measured AI exposure after ChatGPT was released. That design supports cautious inference, not definitive causal attribution.

There are 3 main reasons for caution. First, all AI exposure scores are imperfect. They measure potential exposure to automation, not actual use of AI in workplaces. If the exposure scores do not capture the occupations most affected by AI, the model may understate or misstate the relationship.

Second, the model assumes that higher exposure is associated with proportionately larger employment effects. If AI affects only a small group of very highly exposed occupations, or if it affects all occupations in similar ways, the model may not capture the effect well.

Third, more AI-exposed occupations may have been affected by other forces around the same time. Routine cognitive work was already declining as a share of employment before ChatGPT, and the post-COVID labour market recovery affected occupations differently. The model is designed to test whether the exposure–employment relationship changed after November 2022, but it cannot rule out all occupation-specific effects other than AI.

For these reasons, the result should be read as evidence of a modest exposure-related employment signal, not as proof that AI caused particular job losses. The next chapter tests how sensitive this signal is to alternative timing assumptions, exposure measures, outcomes and model specifications.

We also cannot distinguish between several potential explanations for a weak or absent relationship between AI exposure and employment. The small relationship we observe in the data could be because of low adoption, or insufficient model capability, or because productivity and reinstatement effects are offsetting displacement.

Event study: variable impact of AI exposure over time

Key points

- The core model averages the relationship between AI exposure and employment across the whole post-ChatGPT period. This could miss effects that took time to emerge.
- There are reasons to think that AI's effect on employment could grow over time. It takes time for businesses to adopt AI, and AI model capabilities have improved markedly over time.
- To test whether the AI exposure–employment relationship has changed over time, we estimate our model quarter-by-quarter, in an 'event study' design using the same data and approach as the core model.
- Before ChatGPT, there was no relationship between occupations' exposure to AI and their employment growth.
- After ChatGPT, the estimates are more variable. But recent quarters are mostly negative, suggesting that the exposure-related employment signal is clearer in the latest data than it was immediately after ChatGPT was released.

Why estimate an event study?

The core model provides a single summary estimate of how the relationship between AI exposure and employment growth has changed, comparing the periods before and after ChatGPT. That is useful for a central estimate, but it imposes a simple timing structure: it assumes the exposure-related relationship changes in a broadly consistent way across the post-ChatGPT period.

That may be too simple. Firms may have taken time to experiment with generative AI, adopt it in production, and adjust staffing. More recent AI models have also been more capable than the tools available immediately after ChatGPT was released. If any employment effect emerged gradually, or only in recent quarters, the core model could understate or blur the timing of the change.

We therefore estimate a version of the model that examines the exposure–employment relationship quarter-by-quarter (an 'event study'). It uses the same data and exposure scores as the core model but allows the relationship between AI exposure and employment to vary quarter by quarter. Quarter-by-quarter estimates are noisier than the core estimate, so individual quarters should not be over-interpreted.

Event study results

We estimate the relationship between AI exposure and employment for each quarter. The results – the main coefficient of interest in each quarter – are shown in Figure 34. Technical information about the event study is in Appendix B at page 72.

Of the 31 quarters before the introduction of ChatGPT, only one has an exposure–employment relationship that is statistically significant at the 5% level.[65] This is reassuring, as it means there was no general tendency for more exposed occupations to grow more slowly (or rapidly) than less-exposed occupations before November 2022.

For the 13 quarters after the introduction of ChatGPT, the results are variable but do turn more negative. Of the 13 quarters post-ChatGPT, there are 6 with a statistically significant negative exposure–employment relationship.

The most recent quarter's coefficient is -0.017 and is statistically significant at the 5% level. This should be read as a relative shortfall, not an absolute fall in employment: for an occupation one standard deviation above average AI exposure, employment is estimated to be about 1.7% lower than the model-implied counterfactual in which the exposure–employment relationship had not changed since the reference quarter.

Figure 34: The exposure–employment relationship is more negative in recent quarters
Relationship between AI automation exposure and employment, by quarter

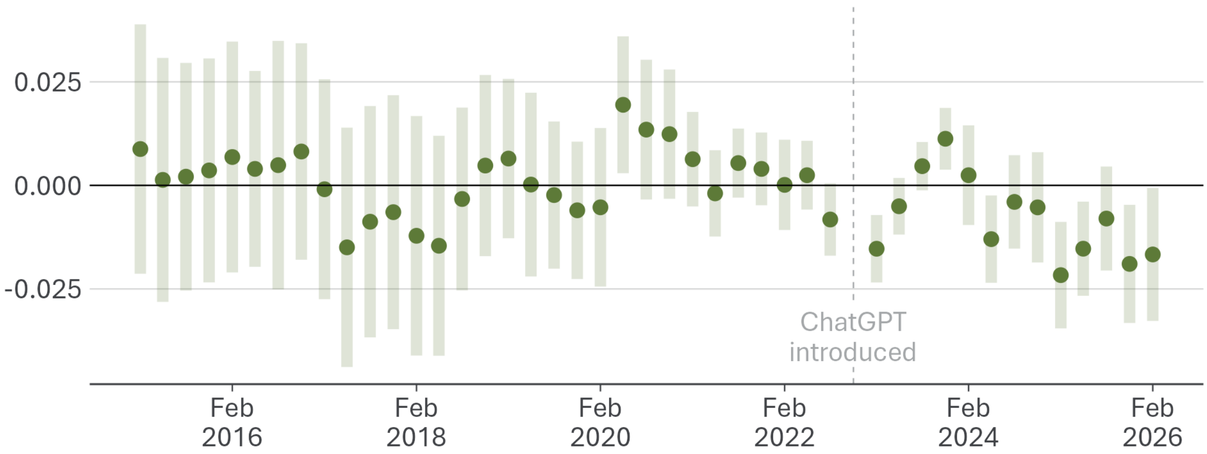

Note: Bars represent the 95% confidence intervals, meaning bars which do not cross the x-axis are at least statistically significant at the 5% level.
Sources: ABS Labour Force (Detailed), JSA (2025a), and DEWR calculations.

The event study complements the core model. Its findings are similar to the core result, while suggesting that the negative exposure–employment relationship is clearer in recent quarters. This could be because AI model capabilities have improved, or because AI adoption has increased over time. The findings of the event study reinforce the need to treat the core estimate as a possible early signal and to continue monitoring whether it grows, fades or persists.

[65] This outlier is for a quarter when COVID-19 response measures were in place and may reflect a genuine but temporary deviation resulting from pandemic related business closures.

How robust are the model estimates?

Key points

- The core model finds a small negative relationship between AI exposure and employment growth after the introduction of ChatGPT. The event study suggests this relationship has turned more negative in recent quarters, although the quarter-by-quarter estimates are variable.
- But these results are not consistent in each variation of the model. These robustness tests mean that the results of our core model should be regarded as a tentative early signal of a possible change in the exposure–employment relationship, not definitive evidence of a negative effect.
- The strongest reason for caution is that the model result depends on how occupational AI exposure is measured. We test three exposure measures and only the JSA measure produces a statistically significant negative exposure–employment relationship. While the JSA automation exposure scores are likely to be the most appropriate for our analysis, this is still reason for caution about the result.
- The results are slightly sensitive to the choice of how the pre-treatment and post-treatment periods are defined. All scenarios we examine produce a negative estimate (weakened employment for more exposed occupations), but not all are statistically significant.
- Our results are not driven by a small number of occupations. Re-estimating the model many times on resampled sets of occupations produces estimates broadly consistent with the main result.
- Modelling choices matter. Some model variations support the core result, while others weaken it or make it statistically insignificant.

Our core model is subject to a range of assumptions and limitations. In this chapter, we check if our results still hold if we make various changes to the model framework. If the estimates from these alternative approaches are similar to the core model, this would increase confidence that the core estimate is not an artefact of one modelling choice.

In this chapter, we answer five questions about the results from our core model:

- **Is the result driven by a few occupations?** No. Re-estimating the model many times on re-sampled occupations gives results broadly consistent with the core model, which is reassuring.
- **Is it sensitive to our choices in defining time periods**? Somewhat. All timing variants produce negative estimates, but the size and statistical significance of the estimate varies.
- **Does the result vary when we use different measures of occupational AI exposure?** Yes. This is the strongest reason for caution: using two alternative exposure measures does not reproduce the statistically significant negative relationship found with the JSA measure.
- **Do the results look different when we model hours worked or job advertisements rather than employment?** Yes. The hours-worked estimate is similar in direction to the employment result but is not statistically significant, while the job-advertisements model does not show the same negative relationship.

- **Does the result depend on model design?** To some extent. Several alternatives produce results similar to the core model, but others weaken the result or make it statistically insignificant.

Together, these tests suggest the core result is not an artefact of a few occupations, but nor is it robust to every reasonable alternative specification. The most important limitation is that the result depends materially on the exposure measure used.

Sensitivity to excluding some occupations

The core model's main estimate combines all occupations. There is a risk that the estimated effect is driven by a small number of occupations. There is also a risk that some occupations' AI exposure may have been mis-estimated by JSA.[66]

To address these concerns, we estimate the main model 1,000 times, each time re-sampling occupations, so that some occupations are omitted from a draw and others appear more than once. This is a process known as 'bootstrapping'.[67]

We do this to see how the estimated change in the exposure–employment relationship pre- and post-ChatGPT (the main coefficient of interest, $\Delta\beta$) varies when we randomly exclude some occupations. The results are shown in Figure 35.

Figure 35: Most bootstrap estimates show a small negative employment signal for AI-exposed occupations

Frequency of Δβ estimates from 1,000 bootstrap draws

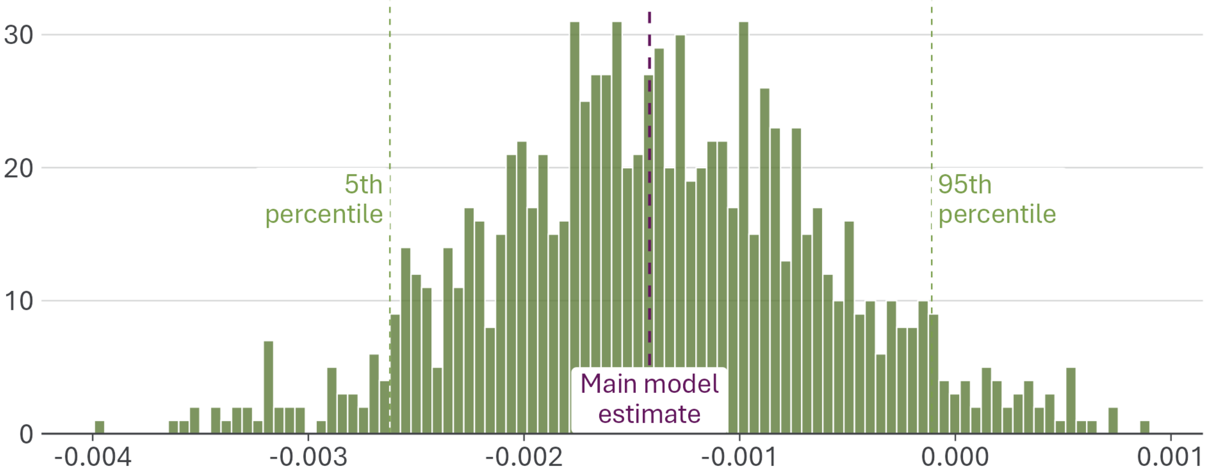

Note: Chart shows Δβ across 100 bins.
Sources: ABS Labour Force (Detailed), JSA (2025a) and DEWR calculations.

The size of the coefficient ($\Delta\beta$) varies across the random draws but is centred on the estimate from the core model. In just over 95% of bootstrap draws, the coefficient is negative. This is reassuring, as it suggests the core model result is not driven by a small number of occupations. But a little under 5% of samples produce a zero or positive estimate, so the sign is not completely stable. This supports a

[66] This risk is common to all occupation-level AI exposure measures and is one reason we test alternative measures. We believe JSA's exposure scores are the most appropriate for our analysis in the Australian context, which is why we use them in our core model.
[67] This occupation-level bootstrap re-samples occupations with replacement. In each draw, some occupations are omitted and others appear more than once. In each random draw, we sample an occupation for all quarters, so each draw preserves the occupation's full time series, and the balanced sampling means all occupations appear the same number of total times across samples.

cautious interpretation: the bootstrap exercise strengthens confidence that the core estimate is not an outlier, but it does not make the result definitive.

Sensitivity to timing

Our event study model suggests that the relationship between AI exposure and employment has varied over time. This raises the question of whether our results may differ if we change the periods of time used to define the 'pre' and 'post' AI era.

To test this, we estimate 4 variations on the main model:

- One variation excludes the COVID-19 pandemic period from the 'pre-treatment' period
- One variation starts the pre-treatment period in 2010, rather than 2015 as in our core model
- One variation starts the post-treatment period one year after ChatGPT's introduction, to allow for the possibility of delayed effects
- One variation combines two of these adaptations, both starting the post-treatment period one year after ChatGPT's introduction and also excludes the COVID-19 pandemic period

Figure 36 shows the estimates from these model variations, along with the main model. All variations, like the main model, produce negative estimates. The size and precision of the estimate varies across specifications. The variation that excludes the COVID-19 period is not statistically significant. The variation that starts the 'post-treatment' period a year after the introduction of ChatGPT shows a slightly more negative effect than the main model, consistent with the findings of our event study which showed effects becoming somewhat more negative over time. When combining these changes, the results are very similar to the main model.

Figure 36: Changing time period definitions affects the size and significance of the coefficient of interest

Variation in the main coefficient of interest (Δβ) across different model specifications

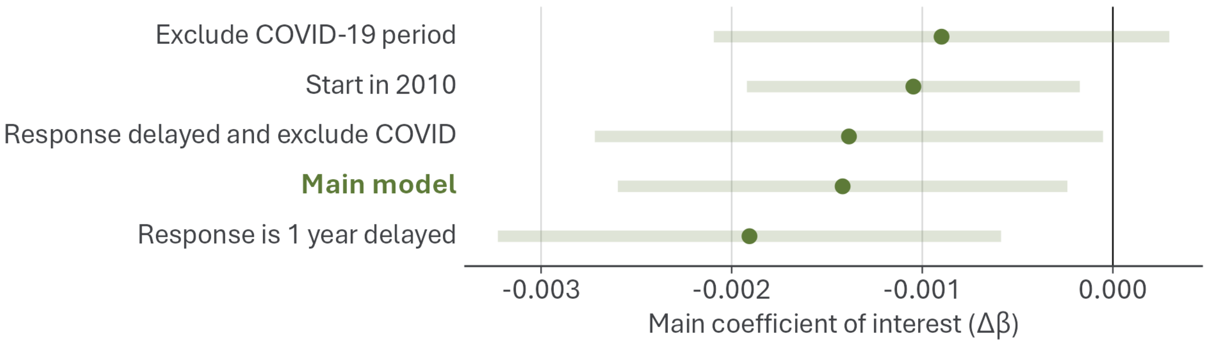

Note: Bars represent 95% confidence intervals. Bars which do not cross zero are at least statistically significant at the 5% level.
Sources: ABS Labour Force (Detailed), JSA (2025a) and DEWR calculations.

Excluding the COVID-19 period

Our model is based around comparing the relationship pre- and post-ChatGPT, but it is possible that the pre-ChatGPT employment-exposure relationship was affected by COVID.

To test this, we estimate a model where the post-treatment period starts as normal from November 2022, but the pre-treatment period ends on February 2020, with data after February 2020 up to the quarter preceding November 2022 excluded from the model.[68] This is labelled 'Exclude COVID-19

[68] The excluded period between COVID-19 and November 2022 will not only capture the immediate impacts of COVID-19, but also the start of the recovery period which includes some changes in labour market composition that have so far persisted, such as expansion of the non-market health and care sectors. Excluding this period may therefore also mean the

period' in Figure 36. The estimated coefficient of interest ($\Delta\beta$) in this model is still negative but is smaller in magnitude than in the main model, and is not statistically significant.[69]

This specification still shows no statistically significant pre-ChatGPT[70] exposure trend and a negative post-ChatGPT exposure trend.[71] But the estimated change between the two periods is not significant. We treat this variation as an important caution about the precision of the core estimate.

Extending the pre-treatment period back to 2010

In our main model, the pre-treatment period starts in 2015. The choice of where to start the pre-treatment period could impact the model results. Going back further means the pre-treatment period has more data to estimate the relationship but also means it may over-weight historic labour market trends that are less relevant today.

To test the model's sensitivity to different start dates, we estimate a version of the model which uses the February 2010 quarter as the starting point of the pre-treatment period. The main coefficient of interest ($\Delta\beta$) from this model is shown in Figure 36. The result is negative and statistically significant,[72] though the magnitude of the effect is slightly smaller than for the core model. The fact that the coefficient remains negative and statistically significant in this model variation provides some reassurance that the core result is not driven by the choice of 2015 as the start date.

More details on the choice of pre-treatment period, including the decision to use 2015 as the start of the 'pre-treatment' period, are presented in Appendix C on page 79.

Delayed employment response

If employers take time to adopt AI, or there are frictions which slow the pace of occupational change, then the impacts of AI on employment may not show up straight away. This could also happen if later releases of generative AI tools outperform earlier releases in a way which causes employers to change their AI strategy. Our event study model suggests that there may be such a delayed response, with a more negative exposure–employment relationship apparent in more recent quarters.

We estimate a version of the core model where the pre-treatment period ends on November 2022, but the post-treatment period only begins from November 2023. The year of data between these periods is not used in this model variation. The main coefficient of interest ($\Delta\beta$) from this variation is shown in Figure 36 as 'Response is 1 year delayed'. It is more negative than the core model and is highly statistically significant.[73] This is consistent with the pattern seen in the event study, where the negative exposure–employment relationship took some time to emerge.

We also estimate a version of the model where the response is delayed and the pre-treatment period ends in the February quarter of 2020, before the COVID-19 pandemic period.[74] This is shown in Figure 36 as 'Response delayed and exclude COVID'. The $\Delta\beta$ term in this model is very similar in scale as for the core model, it is statistically significant,[75] though less so than the core model. Excluding both the COVID-19 period and the year after ChatGPT was released results in a model which produces conclusions that are materially the same as those based on our core model.

---

parallel trends conditional on exposure assumption could be violated due to not capturing some trends which emerged in the years before ChatGPT was introduced.

[69] The p-value for this is 0.14. A p-value is a measure of how likely it is that an estimate could have arisen through pure chance. A p-value of greater than 0.05 means that the estimate is not statistically significant, using the conventional definition of significance.

[70] p=0.35.

[71] p=0.032.

[72] p=0.019.

[73] p=0.0046.

[74] Note that all quarters between May 2020 and August 2023 inclusive are removed from the sample in this specification.

[75] p=0.042.

Sensitivity to AI exposure measures

Our core model relies on the AI exposure scores produced by JSA (JSA 2025a). JSA's scores are derived from work of the ILO[76] and are the main scores produced specifically for Australia. But all estimates of potential automation exposure are imperfect, and the imperfections could affect our results.

We test whether the results from our core model are sensitive to using two alternative measures of AI exposure. We use the same data as for the core model, except that we exclude occupations for which we do not have an alternative exposure score. Appendix A provides more detail on the exposure scores, including the process of converting to the Australian occupational framework.

The first alternative exposure measure we use is the **AI Occupational Exposure (AIOE)** scores by Felten et al. (2021).[77] This is a measure derived in the US, 'using data from 10 selected AI applications'.[78] A key difference between the AIOE framework and the JSA framework is that AIOE does not explicitly aim to identify exposure to automation specifically, instead focusing on exposure in general, and as a result gives higher exposure scores to professional non-routine cognitive jobs.

A second alternative exposure measure we use is **Anthropic's observed exposure** scores.[79] These are based on real-world usage of Anthropic's Claude models. They classify queries into tasks, and map these to US occupations.

Figure 37 shows the main coefficient of interest ($\Delta\beta$) when we use these alternative measures of AI exposure. When we use Felten's AIOE measure, the coefficient is negative but smaller than the main estimate and is not statistically significant.[80] When we use the Anthropic observed exposure measure, our coefficient is near-zero and is again not statistically significant.[81] In other words, when we use exposure measures other than the JSA scores we do not find statistically significant evidence that AI exposure is associated with a decline in employment.

Figure 37: The AI-employment relationship is not statistically significant if we use alternative measures of exposure

Variation in the main coefficient of interest (Δβ) across different model specifications

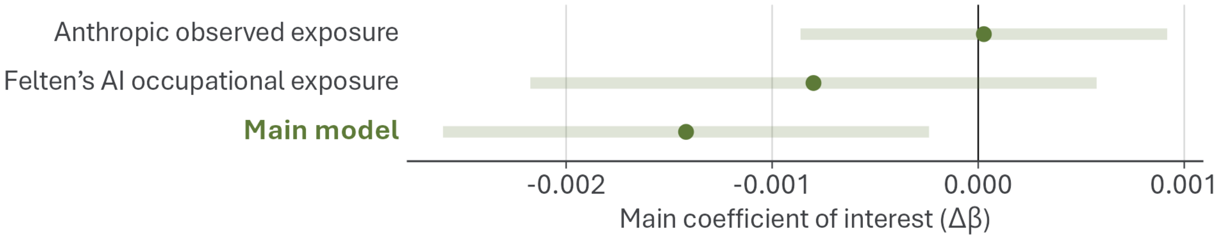

Note: Bars represent 95% confidence intervals. Bars which do not cross zero are at least statistically significant at the 5% level.
Sources: ABS Labour Force (Detailed), JSA (2025a), Felten et al. (2021), Massenkoff and McCrory (2026) and DEWR calculations.

There are reasons to prefer the JSA scoring framework, including that it is based on the Australian occupation structure, focuses on automatability and is not restricted to a specific model. But the

[76] Gmyrek et al. (2023)
[77] See Felten et al. (2021). As with the JSA exposure scores, the AIOE scores are normalised as standard deviations from the mean.
[78] These are 'abstract strategy games, real-time video games, image recognition, visual question answering, image generation, reading comprehension, language modelling, translation, speech recognition, and instrumental track recognition.' (Felten et al. 2021).
[79] Massenkoff and McCrory (2026). As with the other exposure scores, we normalise the Anthropic scores.
[80] p=0.25.
[81] p=0.95.

choice of exposure measure matters. Among the exposure measures we test, only the JSA measure produces a statistically significant negative employment estimate.[82]

This is the main reason for caution in interpreting our main results. When we replace JSA's Australian automation exposure scores with two alternative exposure measures, the negative employment estimate is no longer statistically significant. We also examine how the event study model is sensitive to alternative exposure measures in Appendix B. Neither of these alternative event study models support the finding that more exposed occupations have statistically significantly reduced employment in the most recent quarters compared to November 2022. This does not prove the main result, using the JSA scores, is wrong: the alternative measures may be less well suited to Australian occupations. But it does mean the finding depends materially on how AI exposure is measured, so the core result should be treated as tentative rather than definitive.

Sensitivity to outcome measure

Our core model estimates the AI exposure–employment relationship using headcount employment. But we also test whether our results are similar when we estimate the relationship between occupational AI exposure and two other indicators, namely hours worked and job vacancies. The results are shown in Figure 38.

AI might affect hours worked before, or instead of, reducing headcount. To test this, we run our model with total hours worked in each occupation, rather than employment, as our dependent variable.[83] The coefficient of interest ($\Delta\beta$) from this model variation (see Figure 38) is similar to the main model estimate, but is not statistically significant at the 5% level.[84]

It is plausible that the effect of AI might first become apparent in the flow of hiring rather than the stock of employed persons or hours worked. To test this, we also run a version of our model where, instead of employment, the model's outcome is the total number of job advertisements in each occupation.[85] The data source for job advertisements is the monthly JSA Internet Vacancy Index (IVI) (JSA 2026b), aggregated into quarters that match the timing of the ABS Labour Force Survey.

In contrast to the employment and hours models, the main coefficient of interest in the job advertisements model is positive, though the estimate is not statistically significant.[86] High-exposure occupations had weaker job-advertisement growth before ChatGPT, but the relationship became less negative after ChatGPT.

---

[82] There is no settled best measure of AI automatability for Australian occupations, so results should be interpreted alongside robustness checks using alternative measures. Caution should be taken in assuming that a measure is valid because it produces a statistically significant estimate, especially where the effects of automation on labour demand may be ambiguous (Evans 2026).
[83] The data source is ABS Labour Force Detailed (ABS 2026a), the same as for employment by occupation.
[84] p=0.068.
[85] As with employment and hours worked, the data is at the 4-digit ANZSCO occupation unit group level.
[86] p=0.092.

Figure 38: The results change with hours worked or job advertisements as the dependent variable
Variation in the main coefficient of interest (Δβ) across different model specifications

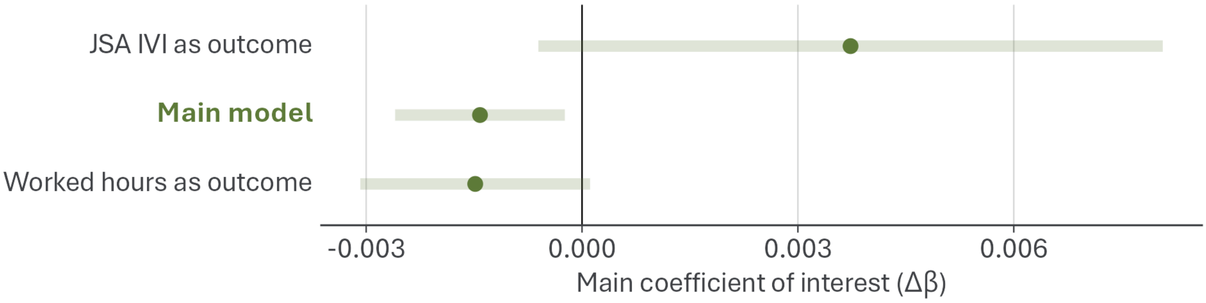

Note: Bars represent 95% confidence intervals. Bars which do not cross zero are at least statistically significant at the 5% level.
Sources: ABS Labour Force (Detailed), JSA (2025a), JSA IVI and DEWR calculations.

Sensitivity to model design choices

The core estimate is sensitive to some methodological choices, especially the choice of AI exposure measure, and to a lesser extent the outcome variable and treatment-period definition. However, reassuringly, variations in the statistical modelling approach mostly produce similar results (Figure 39).

We try a range of variations to the modelling approach, including:

- Using different approaches to estimating the uncertainty around our estimates[87]
- Removing time fixed effects, which means the model no longer controls for shocks common to all occupations in a quarter
- Estimating the model with a different statistical method, Ordinary Least Squares (OLS)[88]
- Weighting occupations by their average employment size before ChatGPT (our main model does not weight by occupation size)
- Removing the level-shift term from our model, so the model allows only for a change in the post-ChatGPT employment growth rate
- Removing from the model any occupations with the largest quarter-to-quarter employment changes, to reduce the influence of extreme observations
- Adding a control for occupations' historical sensitivity to the business cycle

These model variations are discussed in more detail in Appendix C on page 77.

[87] This includes alternative clustering of standard errors and using robust standard errors.
[88] OLS is used rather than our preferred Poisson Pseudo-Maximum Likelihood (PPML) approach.

Figure 39: Most changes to the model have only small impacts on the coefficient of interest

Variation in the main coefficient of interest (Δβ) across different model specifications

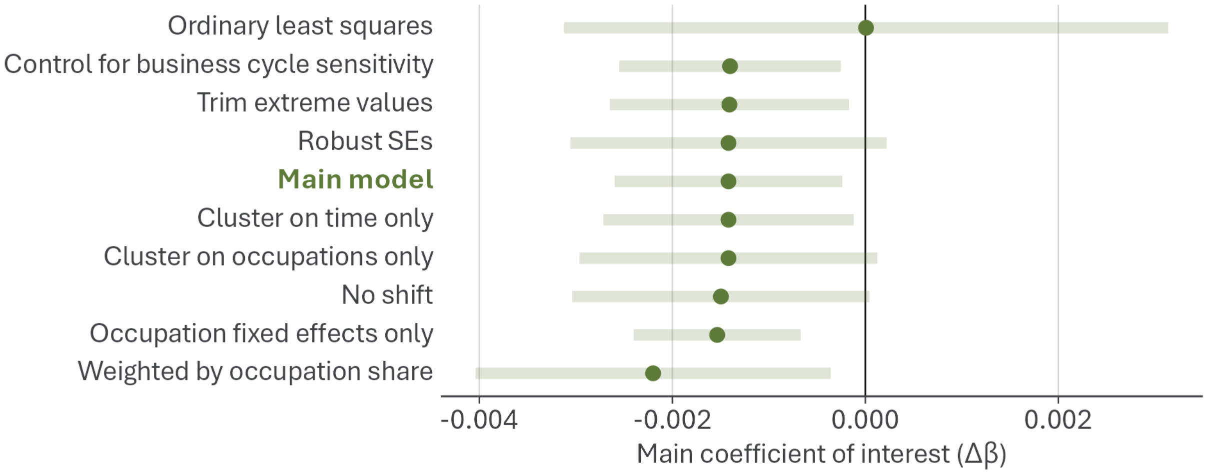

Note: Bars represent 95% confidence intervals. Bars which do not cross zero are at least statistically significant at the 5% level.
Sources: ABS Labour Force (Detailed), JSA (2025a) and DEWR calculations.

It is reassuring that these model variations mostly do not substantially alter our result. The main exception is the OLS model, which finds no exposure–employment relationship. We place more weight on the estimates from our main modelling approach because they are better suited to this setting, for reasons discussed in Appendix B on page 75. But the OLS result remains an important caution.

Another model variation worth noting is one in which we use alternative methods of estimating the uncertainty around the exposure–employment relationship.[89] These variations result in an estimate that is no longer statistically significant, as does a model in which we do not include an immediate variation in the exposure–employment relationship post-ChatGPT.[90] These model variations produce similar estimates of the main coefficient of interest.

Robustness tests suggest caution

Taken together, these checks support a cautious interpretation. The core result is not driven by a few occupations, survives several timing and model-design variations, and is not explained away by our business-cycle control. But it is sensitive to the measure of AI exposure and some other modelling choices. We therefore treat the core model result as suggestive rather than settled causal evidence of AI-driven job loss. Figure 40 shows the results of all robustness tests considered in this chapter.

[89] Specifically, results are insignificant when we cluster standard errors by occupation only, rather than occupation and time. When we estimate the result using robust standard errors, rather than our default two-way clustering, the results are also not statistically significant.

[90] This is the 'no shift' model.

Figure 40: Our results vary with some alternative approaches, which suggests caution in interpreting the results

Variation in the main coefficient of interest (Δβ) across different model specifications

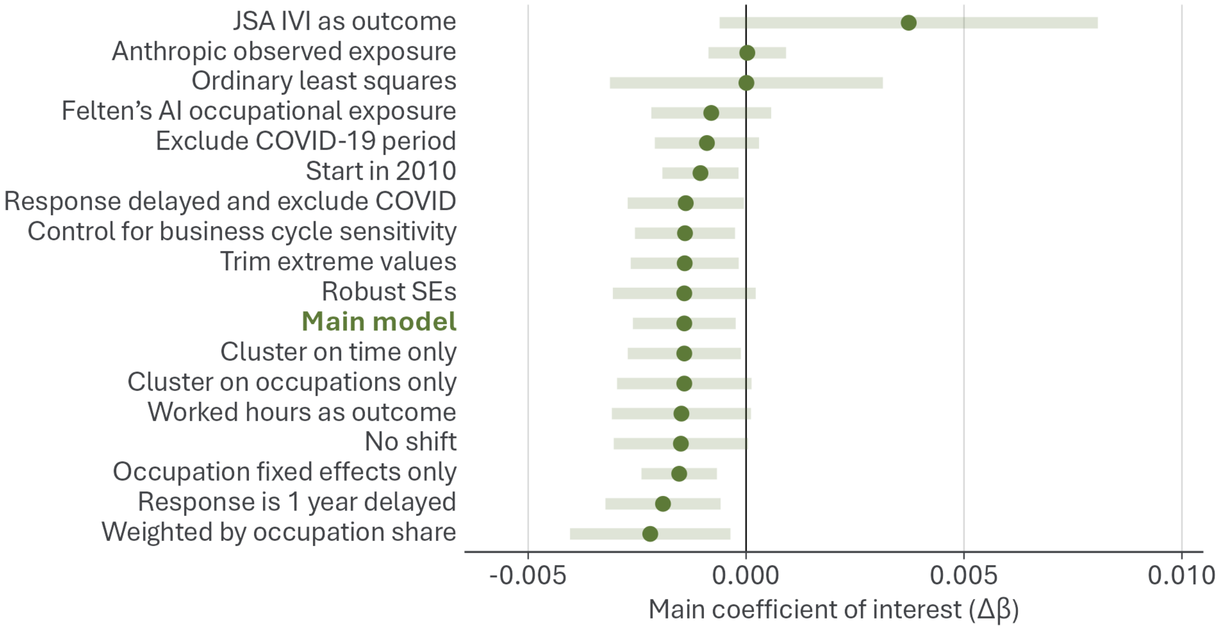

Note: Bars represent 95% confidence intervals. Bars which do not cross zero are at least statistically significant at the 5% level.
Sources: ABS Labour Force (Detailed), JSA (2025a), Felten et al. (2021), Massenkoff and McCrory (2026), JSA IVI and DEWR calculations.

Interpreting the evidence

The evidence in this report does not show broad AI-driven labour market disruption in Australia. But it does show a modest softening concentrated in more AI-exposed occupations.

Overall labour market conditions remain strong by recent historical standards. Youth outcomes have mostly held up. The market for young graduates is not showing significant deterioration. And the pace of occupational reshuffling has not accelerated.

At the same time, employment, hours worked and job advertisements have grown more slowly in occupations more exposed to potential automation. Our modelling suggests the estimated relationship is negative, but small, and should not be read as evidence of large-scale job loss.

The evidence does not show that AI is already causing large-scale job loss in Australia. Continued monitoring will be important to spot any significant effects if and when they occur.

Taken together, the evidence points to limited but non-trivial softening in more-exposed occupations. The descriptive indicators show that employment in high-exposure occupations has softened relative to other occupations. The core model suggests this is not simply a continuation of the pre-ChatGPT exposure–employment relationship. But the robustness tests show that the result is not definitive.

There are other potential explanations for softness in exposed occupations

The main alternative explanation is that more AI-exposed occupations were already on a different trajectory before ChatGPT. Many highly exposed occupations are routine cognitive occupations, including clerical and administrative roles. These occupations have been declining as a share of total employment for many years. Some highly exposed occupations may also have been affected by post-COVID adjustment, changes in the composition of demand and broader business-cycle conditions.

Other phenomena which emerged around November 2022 could also explain an apparent divergence in outcomes by AI exposure if such changes are correlated with exposure scores. It is possible that migration dynamics, sectoral demand shifts, global trade developments or other circumstances could be impacting occupations unequally in a way that has disadvantaged more exposed occupations. The descriptive evidence alone cannot separate AI from these other forces. Slower employment growth in high-exposure occupations could reflect AI exposure, but it could also reflect structural change.

The core statistical model is designed to test this more directly. It asks whether the relationship between AI exposure and employment growth changed after November 2022, rather than simply continuing earlier trends. The model finds a small negative post-ChatGPT shift and does not find a statistically significant pre-treatment exposure trend.

This strengthens the case that there may be something new in the post-ChatGPT period. But it does not rule out all alternative explanations: emerging trends caused by other factors but correlated with AI exposure will also be detected by the model and contribute to the core estimate of how sensitive employment is to AI exposure. The result weakens or disappears under some model variations.

For this reason, we do not interpret the evidence as a settled causal estimate.

What we can and cannot conclude

The evidence supports four conclusions:

- **There is no evidence to date of broad AI-driven labour market upheaval in Australia.** Aggregate conditions remain strong by recent historical standards. The youth labour market has held up better than some international commentary would suggest. Occupational compositional change is not unusually rapid.

- **There is evidence of relative softening in occupations more exposed to potential automation by generative AI.** This is visible across several descriptive indicators, including employment, hours worked, job advertisements and unemployment.
- The core statistical model finds a **modest negative exposure–employment signal after the introduction of ChatGPT**. The estimated magnitude is small. It is much smaller than public claims that AI is already producing large-scale job loss.
- The **model result should be interpreted cautiously.** The model does not observe firm-level AI adoption. It does not identify specific jobs created or destroyed by AI. It estimates whether employment patterns have changed in a way that is correlated with occupational exposure to AI automation. The result is sensitive to some modelling choices, especially the measure of AI exposure.

The evidence does not show that AI is already causing significant job losses in Australia. But it also does not justify complacency. More AI-exposed occupations are showing enough relative softness to warrant close monitoring.

Potential future work

The framework developed in this report provides a basis for tracking whether the tentative signs of slower growth in AI-exposed occupations grow, fade or persist.

Future work could also extend the evidence base. Priorities include:

- Monitoring workers in high-exposure occupations who may have fewer resources or fewer viable pathways to transition into other work
- Using administrative microdata, where possible, to examine worker transitions, firm behaviour and occupation-by-demographics breakdowns
- Updating exposure measures as AI capabilities and adoption patterns change
- Testing whether results differ across industries, regions and demographic groups
- Investigating the AI-wages relationship

Monitoring of employment could be usefully complemented by monitoring of AI capability and adoption.

Conclusion

The Australian labour market is not showing broad AI-driven disruption. But more AI-exposed occupations have softened relative to others, and the core model finds a modest post-ChatGPT negative employment signal.

Because the result is sensitive to exposure measurement and some model choices, it should not be treated as causal proof. Ongoing monitoring of the impacts of AI on the Australian labour market is needed to provide the best foundation to advise government.

References

ABS (2022a) *The labour force framework*, https://www.abs.gov.au/statistics/detailed-methodology-information/concepts-sources-methods/labour-statistics-concepts-sources-and-methods/2021/concepts-and-sources/labour-force-framework.

ABS (2022b) *Conceptual basis of ANZSCO*, https://www.abs.gov.au/statistics/classifications/anzsco-australian-and-new-zealand-standard-classification-occupations/2022/conceptual-basis-anzsco.

ABS (2024) *OSCA - Occupation Standard Classification for Australia*, Canberra, https://www.abs.gov.au/statistics/classifications/osca-occupation-standard-classification-australia/latest-release.

ABS (2026b) *Labour Force, Australia*, Australian Bureau of Statistics, Canberra, https://www.abs.gov.au/statistics/labour/employment-and-unemployment/labour-force-australia/feb-2026.

ABS (2026a) *Labour Force, Australia, Detailed*, Australian Bureau of Statistics, Canberra, https://www.abs.gov.au/statistics/labour/employment-and-unemployment/labour-force-australia-detailed/feb-2026.

ABS (2026c) *Media statement: Modernising the Labour Force Survey from April 2026*, https://www.abs.gov.au/media-centre/media-statements/modernising-labour-force-survey-april-2026.

Acemoglu D (2025) 'The simple macroeconomics of AI', *Economic Policy*, 40(121), 13–58, doi:https://doi.org/10.1093/epolic/eiae042.

Acemoglu D and Autor D (2011) 'Skills, tasks and technologies: Implications for employment and earnings', in D Card and O Ashenfelter (eds) *Handbook of Labor Economics*, Elsevier.

Acemoglu D and Restrepo P (2019) 'Automation and new tasks: How technology displaces and reinstates labor', *Journal of Economic Perspectives*, 33(2), 3–30, doi:https://doi.org/10.1257/jep.33.2.3.

Andrews D and Hansell D (2021) 'Productivity-Enhancing Labour Reallocation in Australia', *Economic Record*, 97(317), 157–169, doi:10.1111/1475-4932.12601.

Autor D (2013) 'The 'task approach' to labor markets: An overview', *Journal for Labour Market Research*, 46, 185–199, doi:10.1007/s12651-013-0128-z.

Autor D, Chin C, Salomons A and Seegmiller B (2024) 'New frontiers: The origins and content of new work, 1940–2018', *The Quarterly Journal of Economics*, 139(3), 1399–1465, https://academic.oup.com/qje/article-abstract/139/3/1399/7630187.

Autor D, Levy F and Murnane R (2003) 'The skill content of recent technological change: An empirical exploration', *The Quarterly Journal of Economics*, 118(4), 1279–1333, doi:https://doi.org/10.1162/003355303322552801.

Autor D and Thompson N (2025) 'Expertise', *Journal of the European Economic Association*, 23(4), 1203–1271, doi:https://doi.org/10.1093/jeea/jvaf023.

Bergé LR, Butts K and McDermott G (2026) *Fixest: A fast and feature-rich framework for econometric estimations in R*, https://arxiv.org/abs/2601.21749.

Berger G (2026) 'Measuring grad unemployment properly', https://macromostly.substack.com/p/measuring-grad-unemployment-properly.

Bessen J (2018) *AI & jobs: The role of demand*, https://www.nber.org/papers/w24235.

Borland J (2026) 'Decline to boom to slowdown: Australia's labour market in the COVID-19 era', *Australian Economic Review*, https://doi.org/10.1111/1467-8462.70047.

Borland J and Coelli M (2022) 'The Australian labour market and the digital economy', in *Proceedings of the Australian Bureau of Statistics & RBA conference*, Reserve Bank of Australia Sydney, Australia, https://www.abs.gov.au/system/files/documents/b148083d79e89d2e287eb9ddd2410f6c/Borland%2C%20Coelli%20-%20The%20Australian%20labour%20market%20and%20the%20digital%20economy.pdf.

Brynjolfsson E, Chandar B and Chen R (2025) *Canaries in the coal mine? Six facts about the recent employment effects of artificial intelligence*, https://digitaleconomy.stanford.edu/publications/canaries-in-the-coal-mine/.

Buckley J, Rankin E and Andrews D (2024) *Non-compete clauses, job mobility and wages in Australia*, https://e61.in/wp-content/uploads/2024/10/Non-competes-RN-20241014.pdf.

Callaway B, Goodman-Bacon A and Sant'Anna PH (2024) *Difference-in-differences with a continuous treatment*, National Bureau of Economic Research, https://www.nber.org/system/files/working_papers/w32117/w32117.pdf.

Chandar B (2025) *Tracking employment changes in AI-exposed jobs*, https://dx.doi.org/10.2139/ssrn.5384519.

Chen J and Roth J (2024) 'Logs with zeros? Some problems and solutions', *The Quarterly Journal of Economics*, 139(2), 891–936, doi:https://doi.org/10.1093/qje/qjad054.

Deming DJ, Ong C and Summers LH (2025) *Technological disruption in the labor market*, https://www.nber.org/system/files/working_papers/w33323/w33323.pdf.

Dhillon Z and Cassidy N (2018) *Labour market outcomes for younger people*, https://www.rba.gov.au/publications/bulletin/2018/jun/labour-market-outcomes-for-younger-people.html.

DISR (2026) *Economic snapshot - April 2026*, Canberra, https://comms.industry.gov.au/pub/pubType/EO/pubID/zzzz69f148bd05a4b578/interface.html.

Duncan OD and Duncan B (1955) 'A methodological analysis of segregation indexes', *American Sociological Review*, 20(2), 210–17, doi:https://doi.org/10.2307/2088328.

Eckhardt S and Goldschlag N (2025) *AI and jobs: The final word (until the next one)*, https://eig.org/ai-and-jobs-the-final-word/.

Eloundou T, Manning S, Mishkin P and Rock D (2023) *GPTs are GPTs: Labor market impact potential of LLMs*, https://arxiv.org/pdf/2303.10130.

Evans B (2026) *Predicting AI job exposure*, https://www.ben-evans.com/benedictevans/2026/5/24/ai-job-exposure.

Fahrer J and Pease A (1993) 'The unemployment-vacancy relationship in Australia', *Australian Economic Review*, (26), 43–57, doi:https://doi.org/10.1111/j.1467-8462.1993.tb00810.x.

Felten E, Raj M and Seamans R (2021) 'Occupation, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses', *Strategic Management Journal*, 42(12), 2195–2217, doi:https://doi.org/10.1002/smj.3286.

Gans JS and Goldfarb A (2026) *O-ring automation*, https://www.nber.org/system/files/working_papers/w34639/w34639.pdf.

Gimbel M, Kendall J and Kulsakdinum R (2026) *Labor market AI exposure: What do we know?*, https://budgetlab.yale.edu/research/labor-market-ai-exposure-what-do-we-know.

Gimbel M, Kinder M, Kendall J and Lee M (2025) *Evaluating the impact of AI on the labor market: Current state of affairs*, https://budgetlab.yale.edu/research/evaluating-impact-ai-labor-market-current-state-affairs.

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

Gross I (2026) *Kookaburras in the coal mine? Measuring AI disruption with Australian administrative data*, https://drive.google.com/file/d/1wbDNJspPly7ZjV1oWpSSX0Ti68KGCeN_/view.

Humlum A and Vestergaard E (2026) *Still waters, rapid currents: Early labor market transformation under generative AI*, https://www.nber.org/papers/w33777.

Ibrahim T (2026) 'Atlassian cuts 1,600 jobs from global workforce as AI slashes labour needs', *ABC*, https://www.abc.net.au/news/2026-03-12/australian-tech-atlassian-cuts-1600-jobs-from-global-workforce/106445334.

Johnston A and Makridis C (2026) *AI, output, and employment*, CESifo Working Paper, https://dx.doi.org/10.2139/ssrn.6460358.

JSA (2025a) *Our gen AI transition: Implications for work and skills, final release*, Canberra, https://www.jobsandskills.gov.au/publications/our-gen-ai-transition-final-release.

JSA (2025b) *Our gen AI transition: Implications for work and skills, final overarching report*, Canberra, https://www.jobsandskills.gov.au/publications/generative-ai-capacity-study-report.

JSA (2026a) *Labour force trending*, Canberra, https://www.jobsandskills.gov.au/data/labour-force-trending.

JSA (2026b) 'Internet vacancy index (IVI)', https://www.jobsandskills.gov.au/data/internet-vacancy-index.

Klein Teeselink B (2025) *Generative AI and labor market outcomes: Evidence from the United Kingdom*, doi:https://dx.doi.org/10.2139/ssrn.5516798.

Kolko J (2026) *Research on AI and the labor market is still in the first inning*, https://www.brookings.edu/articles/research-on-ai-and-the-labor-market-is-still-in-the-first-inning/.

Lambert PJ and Schindler Y (2026) *The broken ladder: AI, remote work, and early-career hiring*, https://ssrn.com/abstract=6787638.

Manning SJ and Aguirre T (2026) *How adaptable are American workers to AI-induced job displacement?*, National Bureau of Economic Research, https://www.nber.org/books-and-chapters/economics-transformative-ai/how-adaptable-are-american-workers-ai-induced-job-displacement.

Mason C, Chen H and Evans D (2026) 'AI adoption and firm demand for workers and skills: New insights from Australia', *The Australian Journal of Labour Economics*, 28(1), 1–28, https://ajle.org/index.php/ajle_home/article/view/303.

Massenkoff M and McCrory P (2026) *Labor market impacts of AI: A new measure and early evidence*, https://www.anthropic.com/research/labor-market-impacts.

Merola R, Ernst E, Samaan D, Rio-Chanona RM del and Teutloff O (2026) *Workers' exposure to AI: What indicators tell us - and what they don't*, doi:https://doi.org/10.54394/00033279.

Rio-Chanona RM del, Ernst E, Merola R, Samaan D and Teutloff O (2025) 'AI and jobs. A review of theory, estimates, and evidence', https://arxiv.org/abs/2509.15265.

Santos Silva J and Tenreyro S (2006) 'The log of gravity', *The Review of Economics and Statistics*, 88(4), 641–658, doi:https://doi.org/10.1162/rest.88.4.641.

SEEK (2024) *SEEK labour market mismatch indicator*, https://www.seek.com.au/about/news/article/labour-market-mismatch-indicator.

VandeHei J and Allen M (2025) 'Behind the curtain: A white-collar bloodbath', *Axios*, https://www.axios.com/2025/05/28/ai-jobs-white-collar-unemployment-anthropic.

Wittenstein J (2024) 'AI can only do 5% of jobs says MIT economist who fears crash', *Bloomberg*, https://www.bloomberg.com/news/articles/2024-10-02/ai-can-only-do-5-of-jobs-says-mit-economist-who-fears-crash.

Wooldridge JM (2002) *Econometric analysis of cross section and panel data*, The MIT press, Cambridge, MA; London, England.

Appendix A: Measuring AI exposure

This appendix sets out further information about the measures of AI exposure used in this report.

It is important to note that exposure measures, in themselves, 'do not offer a prediction about which jobs will be displaced.'[91]

The JSA approach

This report builds on JSA's report *Our Gen AI Transition*.[92] We use JSA's scores that measure each occupation's potential exposure to automation by generative AI tools. JSA's exposure scores are based on work by researchers at the International Labour Organisation (ILO).[93] Like the ILO researchers, JSA scored occupations' exposure by asking a Large Language Model (LLM) to provide a judgement about the level of LLM-automatability of each task associated with the occupation. See page 11 for further information about the JSA scores.

The JSA AI automation exposure scores have limitations, as do all measures of AI exposure.[94] They were produced based on the capabilities of an earlier generation of LLMs, which no longer represents the frontier of performance.[95] The exposure scores measure the *potential* exposure of particular occupations to automation by generative AI, which may differ from actual exposure if adoption of AI is uneven across occupations or regulations impact certain applications more than others. They are based on the tasks currently undertaken by workers in different occupations, which could change over time in response to AI technologies, and do not consider the relative importance of tasks or their substitutability within occupations. Automation measures produced by LLM responses will depend on the specific prompting and version of LLM used and may reflect biases of the LLM model arising from training data or model tailoring.

Selected alternative approaches

To address some of these limitations, we also compare the results of our statistical model using the JSA exposure scores with versions using alternative measures of exposure scores. These are 'AI Occupational Exposure' (AIOE) drawn from Felten et al. (2021) and 'Anthropic observed exposure' drawn from the Anthropic Economic Index Report (Massenkoff and McCrory 2026).[96]

The AIOE scores treat occupations as comprising abilities rather than tasks and consider the relatedness of a broad range of AI technologies (beyond language models) to each work ability. The AIOE scores are based on older iterations of AI models, and do not clearly distinguish between augmentation and automation exposure. The Anthropic observed exposure scores break occupations into tasks, and real-world usage data of Claude is categorised into tasks, with extra weighting placed on usage which is automated. While Anthropic observed exposure scores are the most recent and data-driven, they are limited to a specific AI tool and may understate automation risk coming from other AI applications.

---

[91] Manning and Aguirre (2026, 2)
[92] JSA (2025a).
[93] Gmyrek et al. (2023), which builds on the methodology of Eloundou et al. (2023).
[94] Merola et al. (2026) and Rio-Chanona et al. (2025) each provide overviews of the strengths and limitations of the different frameworks for measuring AI automation exposure.
[95] This is only a problem for our results if model capabilities have developed disproportionately in certain areas that would affect the *relative* ranking of occupations' exposure as opposed to their *absolute* exposure.
[96] For further comparisons of (US-based) exposure scores, see Gimbel et al. (2026). JSA (2025a, 33–34) also contains some discussion of other methods.

Converting US measures to ANZSCO

A limitation of these alternative measures for our purposes is that they are based on US occupation structures and the broader US context. We convert the scores from these US frameworks to Australian occupation structures. As each of these scoring frameworks uses different scales, we standardise each score to have a mean of zero and standard deviation of one.[97]

Not all ANZSCO occupations can be matched to US occupation structures, as the correspondence tables are not perfect one-to-one maps. Further, the US measures do not provide scores for all occupation codes, and there are some ANZSCO occupations that are scored by JSA but only correspond to occupations that were not scored in the US studies. Where ANZSCO occupations are matched to multiple scored US occupations, we average those scores. Figure 41 shows that mapping AIOE scores onto ANZSCO occupations misses 88 occupations that are scored by JSA, and these occupations are reasonably balanced between more and less exposed occupations, only marginally less exposed on average.

Figure 41: Occupations with no AIOE score are marginally less exposed in the JSA framework

JSA scores for occupations with no matched AIOE score (left) and matched AIOE scores (right), per cent of total occupations

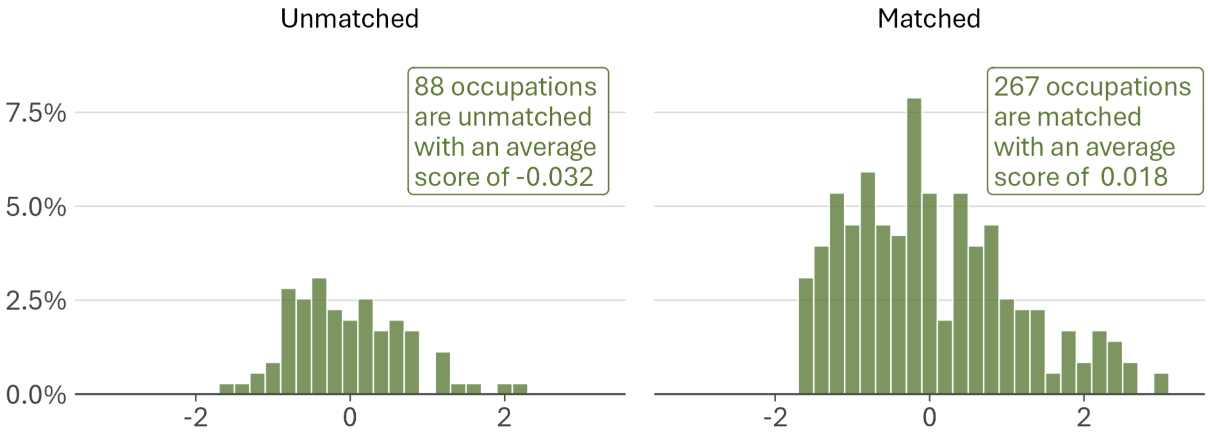

*Note: O*NET SOC (2010) codes are corresponded using a CSIRO provided correspondence table.*
Sources: Felten et al. (2021) and DEWR calculations.

Slightly fewer ANZSCO occupations are able to be assigned an Anthropic observed exposure score. The Anthropic observed exposure scores are published using an updated 2018 O*NET Standard Occupational Classification (SOC) coding framework, which needs to be first mapped to the 2010 O*NET SOC framework for use with the ANZSCO correspondence table. 91 ANZSCO occupations scored by JSA were not matched to Anthropic observed exposure scores, with these occupations having a roughly equal balance between more and less exposed occupations in the JSA scoring framework, only marginally more exposed on average (Figure 42).

[97] This standardisation only partially addresses differences in the distribution of scores. The pre-standardised Anthropic observed exposure scores assign a minimum score of zero to more than half of all occupation codes, so its standardised scores are still not normally distributed.

Figure 42: Occupations with no Anthropic observed exposure score are only marginally more exposed in the JSA framework

JSA scores for occupations with no matched Anthropic score (left) and matched Anthropic scores (right), per cent of total occupations

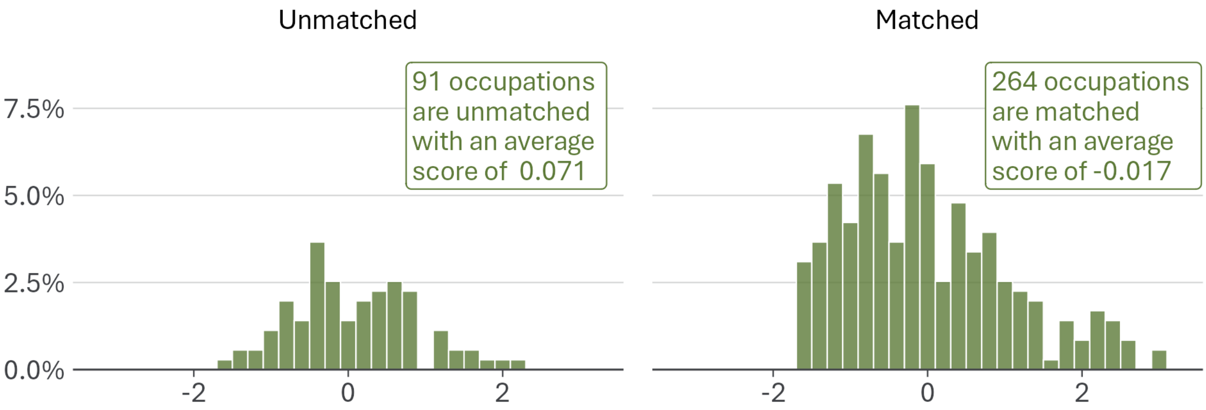

*Note: O*NET SOC (2018) codes are mapped to O*NET SOC (2010) codes via the BLS crosswalk, then corresponded using a CSIRO provided correspondence table.*
Sources: Massenkoff and McCrory (2026), JSA (2025a) and DEWR calculations.

Comparing scores across measures

Each of these approaches uses a different approach to defining 'exposure', leading to substantial differences in which occupations are most exposed. Looking at the most exposed occupations under each framework (Table 3), the JSA and Anthropic observed exposure measures appear more similar, with clerical work ranked among the most exposed. In contrast, the AIOE measure includes highly qualified cognitive roles such as psychologists. These differences may be due to the AIOE measure not emphasising automatability.

Table 3: The most-exposed occupations differ across the exposure measures

Exposure rank	JSA automation exposure	AIOE	Anthropic observed exposure
1	Keyboard Operators	Actuaries, Mathematicians and Statisticians	Call or Contact Centre Workers
2	Telemarketers	Purchasing and Supply Logistics Clerks	Keyboard Operators
3	Filing and Registry Clerks	Accountants	Sales Representatives
4	Human Resource Clerks	Auditors, Company Secretaries and Corporate Treasurers	Software and Applications Programmers
5	Call or Contact Centre Workers	Psychologists	ICT Support Technicians

Sources: JSA 2025(a), Felten et al. (2021) and Massenkoff and McCrory (2026).

While looking at the top 5 gives some insights, there are hundreds of occupation categories, and we use the *tasks framework* (see page 39) to explore the variation in exposure scores across measures in more detail. Figure 43 shows the distribution of exposure scores across the three AI exposure measurement frameworks we consider. Compared to JSA, AIOE assesses routine cognitive jobs as a little less exposed, on average, while non-routine cognitive jobs are marked as more highly exposed. While the average scores are similar between the JSA and Anthropic measures, the distribution of those scores is quite different.

Figure 43: The alternative exposure scores differ from JSA and from each other

Distribution of AI exposure scores across 3 frameworks

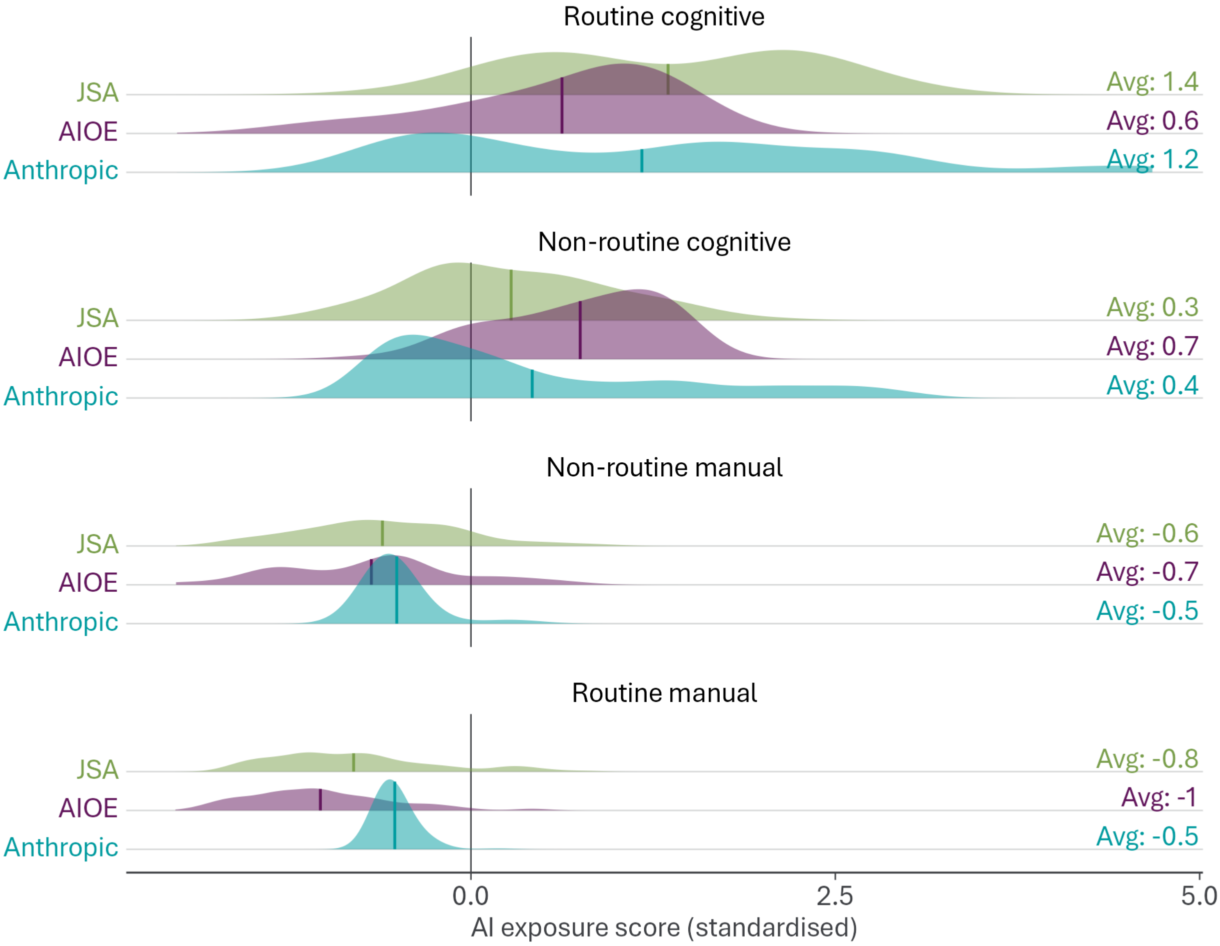

Note: each set of scores is standardised to have a mean of 0 and standard deviaion of 1.
Sources: JSA (2025), Felten et al. (2021), Massenkoff and McCrory (2026), and DEWR calculations.

These differences reflect different methodological approaches to defining AI exposure. For example, AIOE considers routine cognitive jobs relatively less exposed than does JSA. This may reflect that the JSA score we focus on is explicitly concerning AI automation, rather than AI augmentation.

The JSA measure and Anthropic observed exposure scores differ more in their assessment for manual jobs. Since the observed measure is specifically based on Anthropic's Claude application, which is primarily a language model, almost all manual occupations are given an exposure score of zero, which after standardisation is a slightly negative value reflecting relatively less AI exposure. The JSA measure has more variation among manual jobs, and while most are relatively low exposure, some such as delivery drivers have above average automation exposure in the JSA framework.

Appendix B: Additional details for core statistical model

This appendix provides additional details on the approach used for our core statistical model, as well as the event study model, both originally presented in the core statistical model chapter. We describe the data used, the full specification of the model equations, the Poisson Pseudo-Maximum Likelihood (PPML) estimation method, and present a table of the estimated coefficients reporting the pre-trend and post-trend terms as well as $\Delta\beta$. We also present additional versions of the event study model where alternative measures of AI exposure are used.

Later, in Appendix C we also provide further details on the alternative model specifications used in sensitivity and robustness tests.

Data

We use data on employment by occupation from the ABS Labour Force Survey.[98] Our core model uses the 4-digit ANZSCO unit group classification of occupations, sourced from Table EQ08 of the Labour Force Detailed publication.[99] This is quarterly data, not seasonally adjusted.[100] The model uses data from the February 2015 quarter to the February 2026 quarter.[101]

We use 30 November 2022 as the 'treatment date' – data up to and including the November 2022 quarter is included in the 'pre-treatment' period, while everything from the February 2023 quarter onwards is 'post-treatment'.[102] We have 32 pre-treatment quarters and 13 post-treatment quarters of data in our sample.

We use JSA's occupation-level AI automation exposure scores (Appendix A). These time-invariant scores are standardised to have a mean of 0 and standard deviation of 1.

The ABS Labour Force Survey does not include non-civilians, such as defence force personnel.[103] For this reason, we exclude defence force occupations. We also exclude the small number of occupations which were not assigned a score in the JSA study. Occupations that have zero employment in some quarters are retained in our main model. This leaves us with 355 in-scope 4-digit ANZSCO occupations over 45 quarters, for a total of 15,975 occupation-quarter observations in our sample.

The core model specification

The model is summarised in the equation below. It is a Poisson Pseudo-Maximum Likelihood (PPML) panel model with time and occupation fixed effects.

[98] ABS (2026a).
[99] February 2026 was the final issue of this publication, as the outputs from the Labour Force Survey are being revised. Employment-by-occupation data will resume publication in late 2026, using the new Occupation Standard Classification for Australia (ABS 2024) rather than ANZSCO.
[100] The data is collected in the quarters ending in February, May, August and November. JSA (2026a) produces smoothed ('trended') estimates of this data. We use the original, unsmoothed data.
[101] We test alternative 'pre-treatment' windows as part of our robustness tests.
[102] Using November 2022 as a date of treatment, based on the timing of the public release of ChatGPT, has been used in other studies on the labour market impacts of AI. Examples include Brynjolfsson et al. (2025), Klein Teeselink (2025) and Gimbel et al. (2025).
[103] ABS (2022a).

$$\underbrace{E[Emp_{i,t}|X_{i,t}]}_{\text{Employment in occupation } i \text{ at time } t} = \exp(\underbrace{\alpha_i + \gamma_t}_{\text{Occupation and time fixed effects}} + \underbrace{\beta_{pre} s_i \cdot t_{pre}}_{\text{Relationship between exposure and growth pre-ChatGPT}} + \underbrace{\beta_{post} s_i \cdot t_{post}}_{\text{Relationship between exposure and growth post-ChatGPT}} + \underbrace{\beta_{shift} s_i \cdot Post_t}_{\text{Immediate effect of ChatGPT on employment}}),\ E[Emp_{i,t}|X_{i,t}] \sim \text{Poisson}(\mu_{i,t})$$

Each of the terms in the model, as well as our derived estimate $\Delta\beta$, are described in the table below.

Variable	Description	Interpretation
$Emp_{i,t}$	Level of employment	Number of employed persons in occupation i at time t.
α_i	Occupation fixed effects	Controls for time-invariant differences across occupations. Captures the scale effect of occupation i, reflecting that on average some occupation groups are much larger or smaller than others.
γ_t	Time fixed effects	Captures the uniform occupation growth at time t which affects all occupations equally regardless of AI exposure. Controls for common shocks affecting all occupations.
s_i	Normalised AI exposure score	Each occupation i's JSA AI automation exposure score, standardised so that the average is 0 and the standard deviation is 1. More-exposed occupations have positive s_i scores while less-exposed occupations have negative s_i scores. Exposure does not change over time.
t_{pre}	Pre-ChatGPT time index	The negative number of quarters before the end of the pre-treatment period. It is zero in the post-treatment period.
t_{post}	Post-ChatGPT time index	The positive number of quarters since the start of the post-treatment period. It is zero in the pre-treatment period.
$Post_t$	Treatment period dummy	A dummy variable which is 1 in the post-treatment period and 0 in the pre-treatment period.
β_{pre}	Pre-treatment coefficient	Estimated per-quarter relationship between AI exposure and employment in the pre-treatment period. This captures whether employment growth was correlated with AI exposure before ChatGPT was introduced.
β_{post}	Post-treatment coefficient	Estimated per-quarter relationship between AI exposure and employment in the post-treatment period. This captures the extent to which employment growth has been correlated with AI exposure since ChatGPT was introduced.
β_{shift}	Level shift coefficient	Estimated shift in the constant in the relationship between AI exposure and employment in the post-treatment period.
$\Delta\beta$	**Key statistic of interest**	Calculated as $\beta_{post} - \beta_{pre}$, it captures the **change in the per-quarter relationship between AI exposure and employment**, after controlling for fixed effects.

What the model measures

The gradient terms β_{pre} and β_{post} capture the relationship between AI exposure and quarterly employment growth in the pre-ChatGPT and post-ChatGPT periods respectively. As time passes, t_{pre} (or t_{post}) increases each quarter, and the amount of time before (or after) the introduction of ChatGPT is interacted with the exposure score s_i of each occupation so that the effect captured by β_{pre} (or β_{post} respectively) increases with both exposure and with the passage of time. Whereas γ_t captures overall growth in employment across all occupations, the β terms interact time with AI automation exposure scores to capture the extent to which differences in exposure across occupations explains varying employment growth in each of these periods.

The β_{shift} term captures any step-change in employment levels correlated to AI exposure that occurred at the introduction of ChatGPT and persisted. This represents any immediate and persistent deviation related to AI exposure, something which may not occur if the employment response is more gradual due to AI capabilities changing over time or employers taking time to adopt AI and adjust staffing. This differs from the gradual per-quarter effects of β_{pre} and β_{post} which interact through duration of time from the exposure event at November 2022.

While it does not appear directly in the core model equation, our **key parameter of interest** is

$$\Delta\beta := \beta_{post} - \beta_{pre}$$

$\Delta\beta$ tells us whether the relationship between AI exposure and employment changed after November 2022. It is designed to answer the question: **did occupations with higher AI exposure start growing more slowly after ChatGPT than occupations with lower AI exposure, relative to their previous relationship?**

The model assumes that, without ChatGPT, the earlier relationship between occupations' exposure and employment growth would have continued.[104] Under that assumption, $\Delta\beta$ measures the post-ChatGPT shift in employment growth associated with AI exposure.

If $\Delta\beta$ is positive, employment growth since November 2022 has favoured higher exposure occupations more than in the past. If $\Delta\beta$ is negative, employment growth since November 2022 has favoured lower exposure occupations more than in the past. If $\Delta\beta$ is not statistically significant, then the relationship between AI exposure and employment growth since November 2022 is unchanged from the pre-ChatGPT trend.

As the exposure scores s_i are standardised with mean 0 and standard deviation 1, these estimated coefficients should be interpreted as dose-response effects rather than a binary treatment effect. That is, a more highly exposed occupation is expected in the model to have an amplified employment effect compared to a moderately exposed occupation. A negative $\Delta\beta$ could arise if higher exposure occupations had declining employment growth, if lower exposure occupations had increasing employment growth, or a combination of these outcomes.

Relation to causal frameworks

The **core model equation** is not a standard causal framework, but it is closely related to two established frameworks in the literature: 'difference-in-difference with continuous treatment', and 'interrupted time-series'.

A two-way fixed effects difference-in-difference (DiD) with continuous treatment design typically includes group fixed effects, time fixed effects and a continuous treatment (exposure) variable

[104] This is the **parallel trends conditional on exposure** assumption. This directly relates to our **core question** from the Setting the scene chapter.

interacted with a treatment dummy (in our case $s_i Post_t$ and the related coefficient β_{shift}). It is designed to identify the effect per unit of treatment across the whole treatment period, relative to pre-treatment.[105] Dropping the β_{pre} and β_{post} terms from our core model equation would reduce to this framework, and would be appropriate if we thought that AI exposure was likely to impact exposed occupations uniformly over time after the introduction of ChatGPT.

However, technological change can impact labour markets more gradually. Our **core question** asks about the change in rate of growth, rather than the absolute level of employment. So similar to the interrupted time series framework we allow for not only a level shift (β_{shift}) but also estimate pre-trend (β_{pre}) and post-trend (β_{post}) terms and consider how these trends have changed. Whereas interrupted time series methods typically look at the disruption for just one or two time series, we have a time-series for each occupation. The identification using these trend terms comes from how the cross-occupation relationship between employment and the time-invariant exposure score s_i. The difference in trends $\Delta\beta$ is our key coefficient of interest.

Our specification is therefore best understood as a **structural break in the exposure–employment growth gradient** rather than a canonical treated/untreated DiD, or even as a continuous treatment DiD. While it is not strictly a DiD identification, the identifying assumptions from these frameworks are informative.

In the presence of continuous variation in treatment exposure, the standard DiD parallel trends assumption needs to be modified to account for varying treatment intensities. We call this assumption parallel trends conditional upon exposure. In essence, the requirement is that if occupations with some exposure level s were instead exposed with intensity s' then the change in outcome should be expected to be the same as the change which actually occurred for occupations exposed with intensity s'. This is a strong assumption and means that there are no confounding trends unrelated to AI exposure which differentially impact employment levels and that effect sizes depend on the treatment exposure. While not an exhaustive list of possible confounders, two potential confounding trends are the impacts of COVID and the business cycle, each of which we do robustness testing for.

What the core model finds

Our core model shows that **since the introduction of ChatGPT employment growth has favoured less-exposed occupations more than in the past**. The key coefficient of interest $\Delta\beta$ as estimated in the core model is negative and statistically significant at the 5% level (Table 4). However, the estimated effect is small in magnitude and is sensitive to some modelling choices and the choice of exposure measure.

Our model framework focuses on the difference between the pre-treatment effect and the post-treatment effect. We find that the pre-treatment coefficient (β_{pre}) is statistically insignificant, suggesting no detectable pre-existing relationship between JSA automation exposure and employment growth in the pre-treatment period. This provides some reassurance that the post-treatment result (β_{post}) is not simply continuing an obvious pre-existing exposure-related employment trend. The difference captured by $\Delta\beta$ is therefore almost entirely due to the emergence of a negative and statistically significant post-treatment per-quarter effect of exposure dosage on employment (β_{post}).

Table 4: Core model coefficients

Parameter	Estimate	Lower bound	Upper bound	P value	Significance
βpost	-0.001381	-0.0027	-0.000046	0.043	*

[105] Callaway et al. (2024).

Parameter	Estimate	Lower bound	Upper bound	P value	Significance
βpre	0.000037	-0.0011	0.001198	0.951	
βshift	-0.001191	-0.0176	0.015211	0.887	
Δβ	-0.001418	-0.0026	-0.000239	0.018	*

Note: Lower and upper bounds are 95% confidence intervals using the model's clustered standard errors.
Sources: ABS Labour Force (Detailed), JSA (2025a) and DEWR calculations.

Recall that the β_{shift} parameter captured the instant and persistent impact of AI exposure on employment from the introduction of ChatGPT. Our model estimate of β_{shift} is small and not statistically significant, consistent with an employment response that is gradual over time rather than a sudden and persistent deviation in headcount. It may appear to be of a similar magnitude as $\Delta\beta$ and β_{post} in Table 4, but these other terms are gradient effects which accumulate every quarter. With 13 post-treatment quarters, the total impact of β_{post} that would apply in the final quarter is considerably larger.

Model limitations

Like any model, this approach relies on certain assumptions. If these do not hold, the model may give misleading results.

Validity of exposure scores

The main model relies on AI automation exposure scores, which are an LLM-based assessment of automatability of occupations (JSA 2025a). If this measure of AI automatability is not valid, then our identification strategy would be unsuitable. The limitations of JSA's approach to assigning exposure scores are discussed more in Appendix A on page 63, but importantly they capture potential exposure rather than focusing on implemented usage in automation. To partially assess the sensitivity to the scores assigned by JSA we include some robustness tests which instead use alternative measures of occupational exposure to AI automation.

Linear response with increasing exposure

The β_{pre} and β_{post} terms in our model interact linearly with continuous exposure scores s_i, reflecting an expectation that more exposed occupations may be more affected, an approach which is common in the literature.[106] The model treats AI exposure as a scale, not a binary yes/no categorisation. It asks whether occupations with higher exposure scores experienced larger employment change.

This dose-response setup of the model exploits continuous variation in exposure scores, but also means there is no binary assignment of treated and untreated occupations. In the case that only very highly exposed occupations have their employment impacted, this would contribute to the estimated dosage-response effect, but without identifying the relevant dosage threshold. The model estimates the change in the *relative* employment across occupations with different levels of exposure. If the introduction of generative AI tools has had a uniform effect across occupations, the model approach will not detect this. While not part of our modelling approach, our descriptive indicators track some potential impacts of AI that might affect jobs in the ANZSCO classification equally, such as youth labour market outcomes on page 20.

[106] For example, Klein Teeselink (2025) and Johnston and Makridis (2026).

Risk of confounding trends

We assume parallel trends conditional upon exposure. In other words, we assume that the relationship between AI exposure and employment growth across occupations would have remained unchanged if AI tools had not been introduced. This requires that there are no confounding trends unrelated to AI exposure. This is a strong assumption and could easily be violated by trends that existed before the introduction of ChatGPT and persisted, or trends that emerged after ChatGPT. If there are confounding trends which are correlated with AI exposure, they will be detected by the model and misattributed to the terms sensitive to AI exposure. While not an exhaustive list, compositional change favouring non-routine tasks, changes in migration, sectoral demand, business cycle effects and working from home adoption are potential confounding factors. With the data we have available, we design robustness tests to assess for potential confounding by business cycle effects and COVID-19.

The event study model

Our core model assumes that the relationship between employment and AI exposure has a consistent relationship over time in the pre-treatment period (β_{pre}), and a consistent relationship over time in the post-treatment period (β_{post}), as well as a discontinuity (β_{shift}). An advantage of this approach is that change associated with the introduction of ChatGPT can be summarised by just two terms: $\Delta\beta$ and β_{shift}.

A downside is that this assumption may be incorrect, such as if employers anticipated the development of ChatGPT, or had a delayed response. The time-path of response could be complicated, such as an initial response to increase hiring of more automatable roles for the purpose of training AI models, followed by later replacement of automatable roles. Averaging the impacts in the post-treatment period could mask employment reductions if they only arise after some delay.

We present an alternative model that allows for the impact of AI exposure to vary over time, before and after the introduction of ChatGPT. It suggests that any negative association between higher AI automation exposure and employment may be concentrated in the more recent data.

Event study rationale

Allowing the relationship between employment and AI automation exposure to vary will allow us to identify when impacts arise and when they change in magnitude at a finer grain. This may be particularly useful for ongoing monitoring and understanding how the latest data contributes to our overall measured effect. It also allows us to assess the appropriateness of our assumptions around the timing of treatment, parallel trends, and the constant response magnitude.

There are also some downsides to this approach that should be kept in mind. The impacts in each time period are measured based on the deviation from the time period when ChatGPT was introduced, meaning the reference period is small and results will be sensitive to this choice. We are producing estimates for each time-step, but dividing the data up into many time-periods means each time-step contains only a fraction of the data, so the estimates may be less precise. Because the event study estimates many coefficients, some significant results may occur by chance. We should therefore look for a persistent pattern, not isolated significant quarters.

Event study specification

Our event study uses the same PPML panel model framework as the main model, in which employment levels in each occupation are partially explained by AI automation exposure, as well as time-based effects that are uniform across occupations and occupation-specific factors. The data used is identical to the core model.

Where the core model interacts the AI automation exposure score with the amount of time (quarters) before or after the introduction of ChatGPT, the event study model instead interacts exposure with a series of quarter dummies. The quarter which includes the introduction of ChatGPT is considered the reference time-period, with the estimate for other quarters reflecting the deviation from that baseline.

This setup is summarised in the **event study equation** below, where $T \setminus \{0\}$ is the set of all quarters except November 2022 ($t = 0$) when ChatGPT was introduced:

$$Emp_{i,t} = \exp\left(\alpha_i + \gamma_t + \sum_{q \in T \setminus \{0\}} \beta_q \cdot s_i [t = q]\right)$$

This is very similar to the **core model equation**, with employment levels ($Emp_{i,t}$), occupation fixed effects (α_i) and time fixed effects (γ_t) having the same interpretation as before. The quarter dummy for quarter q is represented with the Iverson bracket notation $[t = q]$, so that β_q captures the employment-exposure interaction only for that quarter. In fact, the sum for a specific quarter reduces to the simplified equation:

$$Emp_{i,q} = \exp(\alpha_i + \gamma_q + \beta_q \cdot s_i)$$

Our event study specification is similar to some models used in the international literature concerning the labour market impacts of AI. Klein Teeselink (2025), using UK data, includes an event study model of monthly exposure-related employment effects that is essentially equivalent to our event study equation, however he expresses it in a log-linear form and estimates his model using ordinary least squares with fixed effects. Brynjolfsson et al. (2025) also includes a range of similar event study models which they estimate using Poisson panel regressions with fixed effects, though their models are split by age group and use exposure quintiles rather than a continuous measure of exposure.

Event study interpretation

The key parameters are the β_q coefficients of the score-quarter interactions. These capture the relationship between AI automation exposure and employment levels in a given quarter. A positive and statistically significant coefficient means that, for that quarter, employment levels increased as exposure scores increased, relative to the period when ChatGPT was introduced. A negative and statistically significant coefficient means that, in that quarter, employment levels decreased as exposure scores increased.

We find a pattern of pre-November 2022 coefficients that are small, statistically insignificant and not systematically trending. This supports the assumption that the relationship between AI exposure and employment growth was roughly zero and broadly stable before the introduction of ChatGPT. As these coefficients are measured relative to when ChatGPT was introduced, this does not mean that there was no relationship between AI exposure and employment, only that it was not changing.

The coefficients after November 2022 reveal the time-varying response of employment to the introduction of ChatGPT. The estimates have become more negative in recent quarters.

As with the core model, occupation fixed effects capture the difference in typical size of different ANZSCO occupation codes and the time fixed effects are interpreted as employment shocks which impact all occupations equally.

The results of the event study were presented in Figure 34. The event study complements our understanding from the core model. Whereas the core model showed that there was no statistically significant relationship between exposure and employment for the pre-treatment period (β_{pre} insignificant), the event study provides evidence that this was stable across pre-treatment quarters. The quarterly estimates vary substantially, and there is no single point that stands out as when negative effects began or accelerated, but for at least the last year the data appears consistently negative. The final estimate of -0.017 is consistent the claim based on the core model that employment for an occupation with exposure score 1 standard deviation above the mean would be expected to have declined by roughly 2% since the introduction of ChatGPT.

Event study with alternative measures of AI exposure

Similarly to the main model $\Delta\beta$ estimate, we test whether our conclusions from the event study model using JSA's automation exposure scores are sensitive to using alternative measures. We again use Felten's AIOE scores and Anthropic observed exposure scores as our comparative measures. The results using these alternatives do not agree with our main results, so this provides further reason for caution in interpreting our main estimates as definitive.

An equivalent event study using Felten's **AIOE scores** as a measure of AI exposure shows a break between February and May 2020 (Figure 44). The relationship between exposure and employment appears relatively stable either side of this break, though there are some small and significant positive estimates in the most recent quarters. This is not consistent with the main event study using JSA measures which showed negative estimates in the most recent quarters. The negative but insignificant $\Delta\beta$ estimated using AIOE scores (Figure 37) is consistent with what we see in the AIOE event study: a simple linear slope in the pre-treatment period would be steeper than in the post-treatment period, suggestive that the change in the relationship has disadvantaged more exposed occupations, but simple slopes do not fit the data well.

Figure 44: Using Felten's AIOE exposure measure suggests that a break favouring more exposed occupations started in May 2020

Relationship between AI automation exposure and employment, by quarter

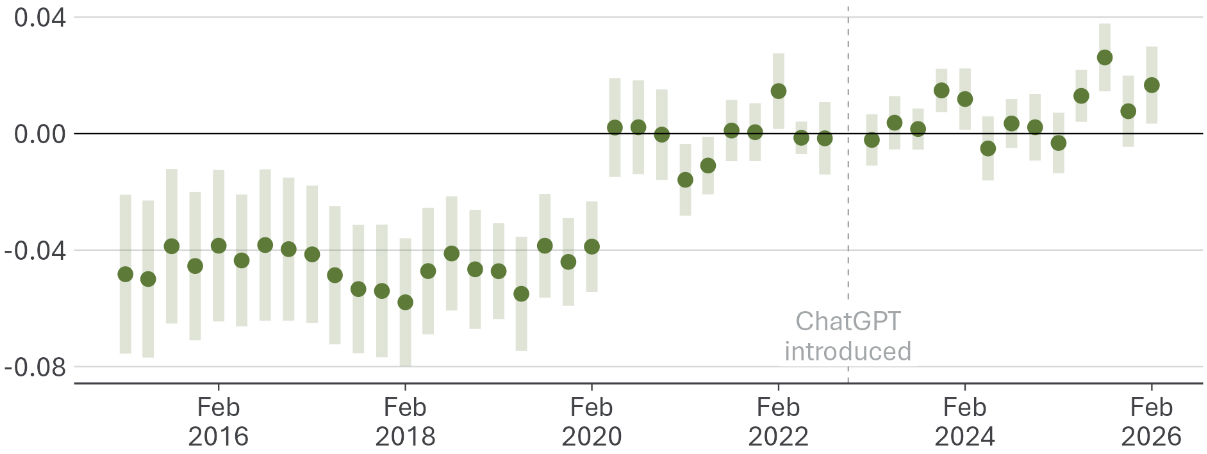

Note: Bars represent the 95% confidence intervals, meaning bars which do not cross the x-axis are at least statistically significant at the 5% level.
Sources: ABS Labour Force (Detailed), Felten et al. (2021) and DEWR calculations.

An equivalent event study using **Anthropic observed exposure** as a measure of AI exposure shows almost all quarters both before and after ChatGPT was introduced have an insignificant relationship between AI exposure and employment, relative to the reference period of November 2022 (Figure 45). This aligns with the insignificant estimate of $\Delta\beta$ when using Anthropic observed exposure scores (Figure 37). This is a different outcome to the main event study, but also different from the AIOE event study.

Figure 45: Using Anthropic observed exposure scores suggests no significant changes over the study period

Relationship between AI automation exposure and employment, by quarter

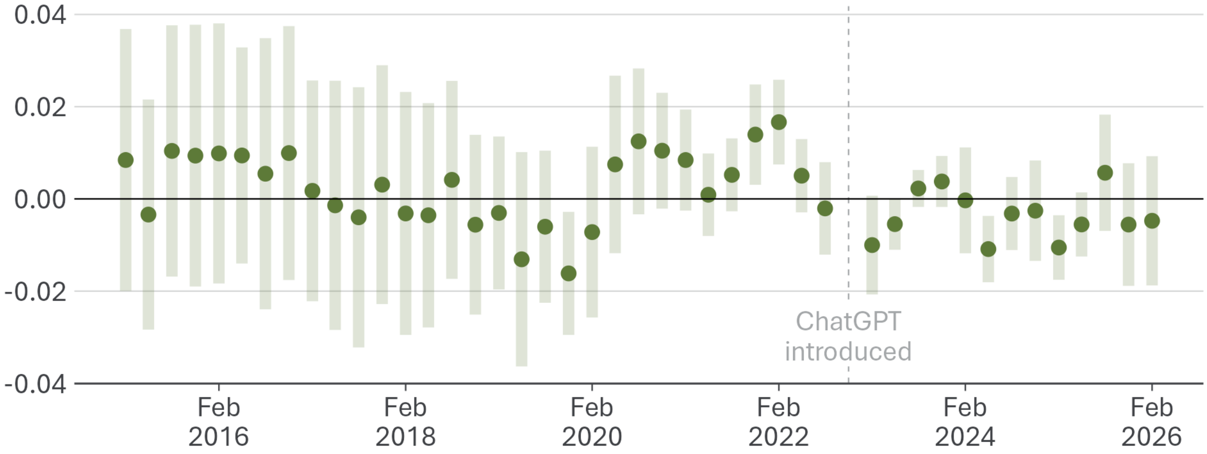

Note: Bars represent the 95% confidence intervals, meaning bars which do not cross the x-axis are at least statistically significant at the 5% level.
Sources: ABS Labour Force (Detailed), Massenkoff and McCrory (2026) and DEWR calculations.

Why we use a PPML fixed effects model

Our core model, event study model and almost all robustness tests are implemented as Poisson Pseudo-Maximum Likelihood (PPML) panel models with fixed effects. Here we will explain more about what the PPML approach and why we use it.

All of the PPML models we use involve treating employment in occupations over time as depending on an occupation effect and an exponential-linear component:

$$E\left[Emp_{i,t}|X_{i,t},\alpha_i\right] = \exp\left(\alpha_i + \beta \cdot X_{i,t}\right) = \exp(\alpha_i) \cdot \exp\left(\beta \cdot X_{i,t}\right)$$

Here α_i is the unobserved time-constant occupation effect for occupation i, $X_{i,t}$ is a vector of observed characteristics which may vary across both time and occupation (such as time dummies, exposure scores or the unemployment rate) and β is a vector of the corresponding coefficients to be estimated. We select independent variables $X_{i,t}$ which we believe may explain the variation in employment $Emp_{it,t}$, and estimate a β that is most likely to produce the observed distribution of $Emp_{i,t}$ given the observed distribution of $X_{i,t}$.

The conditional expectation for each cluster ($E\left[Emp_{i,t}|X_{i,t},\alpha_i\right]$) depends on a multiplicative unobservable cluster effect α_i and a parametric linear term ($\exp\left(\beta \cdot X_{i,t}\right)$). This is the **conditional mean assumption** in the Poisson panel model setup. This is necessary for consistent estimation of β using the standard fixed effects Poisson estimator. Details of the log-likelihood functions conditional on β to be maximised and the form of this estimator can be found in Wooldridge (2002).[107]

Crucially, while the Poisson modelling framework was initially developed for non-negative integer count data, the fixed effects Poisson estimator is consistent whenever the conditional mean assumption holds. It extends to continuous non-negative variables. It is still a consistent estimator when the data does not follow a Poisson distribution, permitting both under-dispersion and over-dispersion which violate variance-mean equality (Wooldridge 2002).

PPML fixed effects models are not the only option to estimate coefficients for equations similar to the conditional mean assumption. Taking logarithms of both sides produces a log-linear equation that

[107] We use the 'fixest' R package implementation of the fixed effects Poisson estimator to perform the estimation. See Bergé et al. (2026).

looks like a typical linear regression model. While logarithms are invalid for $Emp_{i,t} = 0$, techniques using approximating transformations such as $log(1 + y)$ have been used to deal with non-negative data with log-linear relationships.

A number of papers have argued in favour of PPML fixed effects models with non-negative continuous outcome variables, even when the data is not Poisson distributed. Santos Silva and Tenreyro (2006) use simulation evidence to show that ordinary least squares (OLS) regression using log-linearisation with $log(1 + y)$ can produce misleading estimates of β coefficients in the presence of significant heteroscedasticity,[108] while the fixed effects PPML estimator produces more reasonable estimates.[109] This framework has been used in other studies modelling the relationship between employment headcount and AI exposure, such as in Brynjolfsson et al. (2025) who use a PPML fixed effects event study framework to estimate employment effects by month, AI exposure quintile and employee age group, controlling for firm-level confounding effects.

Given the large variations in size across occupations, and the expectations that uptake of AI tools may have occurred at different times and scales across similarly exposed occupations, we anticipate that heteroscedasticity will be substantial and may potentially bias estimates of terms in β if estimated using OLS. We therefore adopt the PPML framework for most of our modelling results, which is also well-behaved for observations with zero employment. Unless stated otherwise, standard errors are two-way cluster-robust standard errors clustered by time and occupation with a small sample correction. This accounts for serial correlation of errors within occupations and across quarters.

108 Heteroscedasticity refers to the fitted values of a model having residuals (difference from actual value) having non-uniform variance (spread). OLS models specifically require no heteroscedasticity for efficient and unbiased estimation, whereas a number of generalised linear models modify this assumption.
109 See Chen and Roth (2024) for more discussion of the 'logs with zeros' issue.

Appendix C: Additional details for model variants

We showed in the robustness tests chapter that our results are mostly consistent across a range of variations to the main model. The key parameter of interest $\Delta\beta$ is usually negative though only sometimes statistically significant. Here we elaborate on the model specifications and provide more detailed results for some of these variations.

Clustering of standard errors and fixed effects

Our central model standard errors are based on two-way clustering (by occupation and time) as is generally best practice for fixed effects panel data models. This accounts for both occupation-level serial correlation in employment levels as well as common shocks (e.g. to macroeconomic conditions) that affect all occupations in a given quarter. However, with a limited number of post-treatment quarters, clustering by time could lead to downward-biased standard errors, which could cause 'false positives' in our model. To test the effects of the modelling choices regarding fixed effects and error clustering on the results, we estimate versions that modify these assumptions.

While still including both occupation and time fixed effects, we adjust the error clustering in three ways: use heteroscedasticity-robust standard errors, cluster only by occupation and cluster only by time. These variations do not change the estimated coefficients, only the standard errors and therefore the confidence intervals. The estimated $\Delta\beta$ for these variants are presented in Figure 39 as 'Robust SEs', 'cluster on occupations only' and 'cluster on time only' respectively. The heteroscedasticity-robust standard errors and cluster only by occupation models are marginally significant at the 10% level, while the clustering only by time model is significant at the 5% level similar to the main model.

We also run a version of the model with only occupation fixed effects, removing the time fixed effects. Standard errors are still two-way clustered by time and occupations. Without time fixed effects, the model estimate of $\Delta\beta$ is a similar magnitude, but is now highly significant at the 0.1% level ('Occupation fixed effects only' in Figure 39). This specification doesn't control for common shocks that affect all occupations, so there are strong reasons to put less weight on this than the core specification.

Ordinary least squares regression

Our core model equation is estimated using a Poisson Pseudo-maximum likelihood (PPML) procedure. We use this rather than a log-linearised ordinary least squares (OLS) regression as PPML is more robust to heteroscedasticity. But, as a robustness test, we include a log-linear style transformation[110] of our core model equation using $log(1+y)$ which we estimate using OLS regression with fixed effects. The OLS model we fit is based on the equation:

$$\log\left(1 + Emp_{i,t}\right) = \alpha_i + \gamma_t + \beta_{pre} \cdot s_i t_{pre} + \beta_{post} \cdot s_i t_{post} + \beta_{shift} \cdot s_i Post_t$$

The terms in the equation for our OLS model have the same meaning and interpretation as in the core model equation, and we again define the key parameter of interest $\Delta\beta := \beta_{post} - \beta_{pre}$. The estimate for $\Delta\beta$ from this model is shown in Figure 39 as 'Ordinary least squares' and is far from statistically significant (p=0.996) with wide error margins using two-way cluster-robust standard errors. The OLS

[110] A true log transformation is not possible as $\log(0)$ is undefined and our data includes quarters with zero employment for some occupations. $\log(1+y)$ is defined for $y = 0$ and is a good approximation of $\log(y)$ for large y.

model does not support the conclusion that the relationship between employment and AI automation exposure has changed since the introduction of ChatGPT.

However, the OLS model is arguably a poor specification in the presence of substantial heteroscedasticity, as discussed in our rationale for using PPML. A plot of residuals against fitted values for this OLS model is shown in Figure 46. The OLS residuals have more variance for smaller values of $\log(1 + Emp_{i,t})$, indicating heteroscedasticity which would violate OLS assumptions and could bias coefficient estimates (Santos Silva and Tenreyro 2006).

Figure 46: The residuals show a fan-shaped pattern, indicating non-constant error variance (heteroscedasticity)

Residuals (vertical) against fitted values (horizontal) for log(1+emp) OLS model

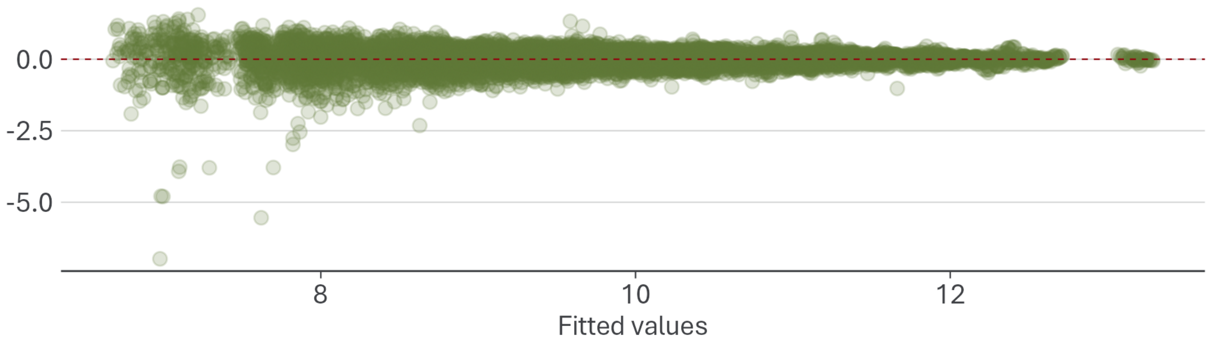

Sources: ABS Labour Force (Detailed), JSA (2025a) and DEWR calculations.

Weighting by pre-treatment share of employment

Our central model is unweighted. It measures the difference in employment trends between high- and low-exposure occupations, placing equal weight on occupations of different sizes, with occupation fixed effects capturing much of the size differences.[111] This is appropriate for our purposes but may understate the effects of AI on workers if effects are concentrated in large occupations. To test this, we run a version of the model that weights occupations by their average employment size in the pre-treatment period. The estimate of $\Delta\beta$ in the weighted model is negative and is about 50% larger than in the main model ('Weighted by occupation share' in Figure 39) and is significant at the 5% level.[112]

To better understand why including weights in the PPML estimation procedure amplifies $\Delta\beta$ we examine how employment size (the weights used) varies with AI automation exposure. Our descriptive analysis included some assessment of the relative size of the exposure quintiles, showing in Figure 21 that the most exposed quintile of occupations represents more than a quarter of employment. However, within each quintile there is a mix of small and large occupations, which will each be weighted based on their occupation share. The most exposed occupations have the rightmost density curve, indicating that many high exposure occupations employed significantly more workers than a typical occupation (Figure 47). The least and second least exposed occupations lean a bit to the left of the median (P50) weight, reflecting the typically smaller size of these lower exposure occupations. Weighting the model means that the $\Delta\beta$ is more sensitive to employment changes for the most exposed occupations and less sensitive to least exposed occupations. Because the main (unweighted) model estimates the average occupation-level association, while the weighted model gives greater influence to larger occupations, the two approaches answer related but different questions.

[111] The variation in size of occupations does enter the core equation through occupation fixed effects and as the outcome, but not in a way to increase the influence of larger occupations in determining the β coefficients.
[112] p=0.019.

Figure 47: The most exposed occupations have a higher weight distribution than other exposure quintiles, as these occupations typically employ more workers

Density of weights across occupations for each AI automation exposure quintile

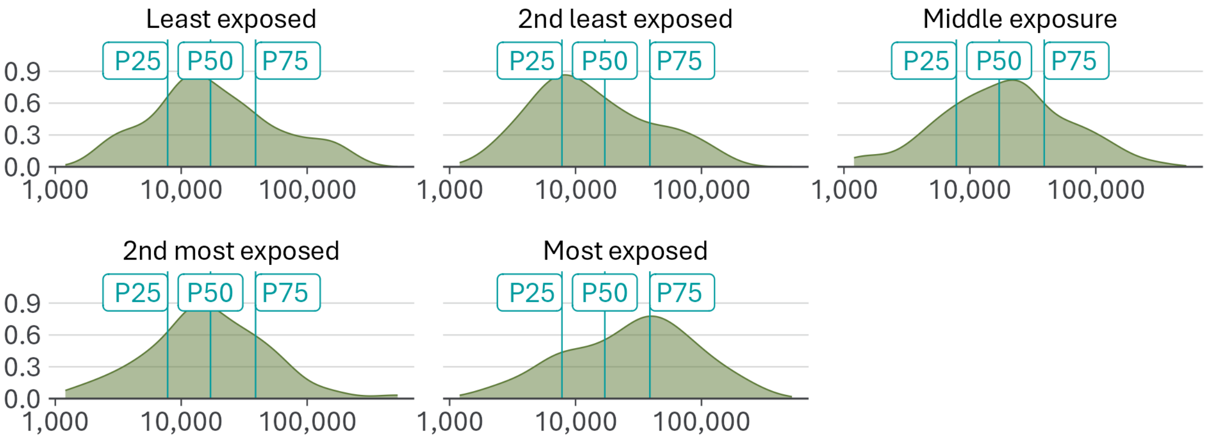

Notes: Weights are calculated as average employment for the pre-treatment period. Weight quartile markers are for all occupations, not by exposure quintile.
Sources: ABS Labour Force (Detailed), JSA (2025a) and DEWR calculations.

Omission of the level-shift term

Our core model allows for a step-change in the relationship between exposure score and employment trends at the time of ChatGPT's introduction, represented by β_{shift} in the core model equation. This is appropriate for our model and if there is no observable step-change, the model will report an insignificant estimate of that term, which our core estimate does.

To test how the inclusion of this insignificant term impacts our key parameter of interest we also test a version that omits β_{shift} and allows only for a change in slope. The estimate of $\Delta\beta$ from this version is shown as 'No shift' in Figure 39. It is of a comparable size as for the main estimate but is not statistically significant at the 5% level.[113] We still prefer the estimate from the core model which includes β_{shift}, as it better isolates the per-quarter impact of AI exposure since ChatGPT's introduction by removing the (small and insignificant) time-constant component.

Filtering out the most extreme values

Employment headcount estimates at the ANZSCO 4-digit level can be somewhat volatile, which can be worse for small occupations. We run the model on a restricted sample, where the largest instances of growth and decline are excluded. This is performed by calculating quarterly growth rates, and for each quarter removing the bottom 1% and top 1% occupations by growth rate. Occupations excluded due to extreme growth in one quarter are still included in other quarters when their growth is not at the extremes. The estimated $\Delta\beta$ for this model is 'Trim extreme values' in Figure 39, which is negative and of comparable in magnitude to the main estimate and still significant at the 5% level.[114]

Selecting the start of the pre-treatment period

There is a lot of flexibility in the choice of when to start the pre-treatment period, as the Labour Force Survey employment-by-occupation time series data goes as far back as August 1986 (ABS 2026a). Going back further means that any estimated pre-trend will draw from a larger pool of data but may also capture confounding trends where historic data contains time-limited effects.

[113] p=0.057.
[114] p=0.026.

Our main model (and most variations) uses data from 2015 onwards. Starting with 2015 balances the need to have sufficient pre-ChatGPT data to estimate the existing exposure–employment relationship, while not going so far back in time that the data is less relevant and more affected by previous shocks (such as the Global Financial Crisis). Figure 48 examines the pre-treatment relationship between AI exposure and employment, estimated using variants of our core model framework which adjust the starting threshold for inclusion in the model. Regardless of the choice of starting year the β_{pre} pre-treatment per-quarter impact of AI exposure on employment is insignificant. Starting from 2015 produces one of the smallest point estimates. We do not consider models that start the pre-treatment period after 2020, as this leaves too few quarters of data.[115]

Figure 48: The pre-trend coefficient is statistically indistinguishable from zero, whenever we start the pre-treatment period

Variation in the estimated β-pre coefficient across different choices of model start year

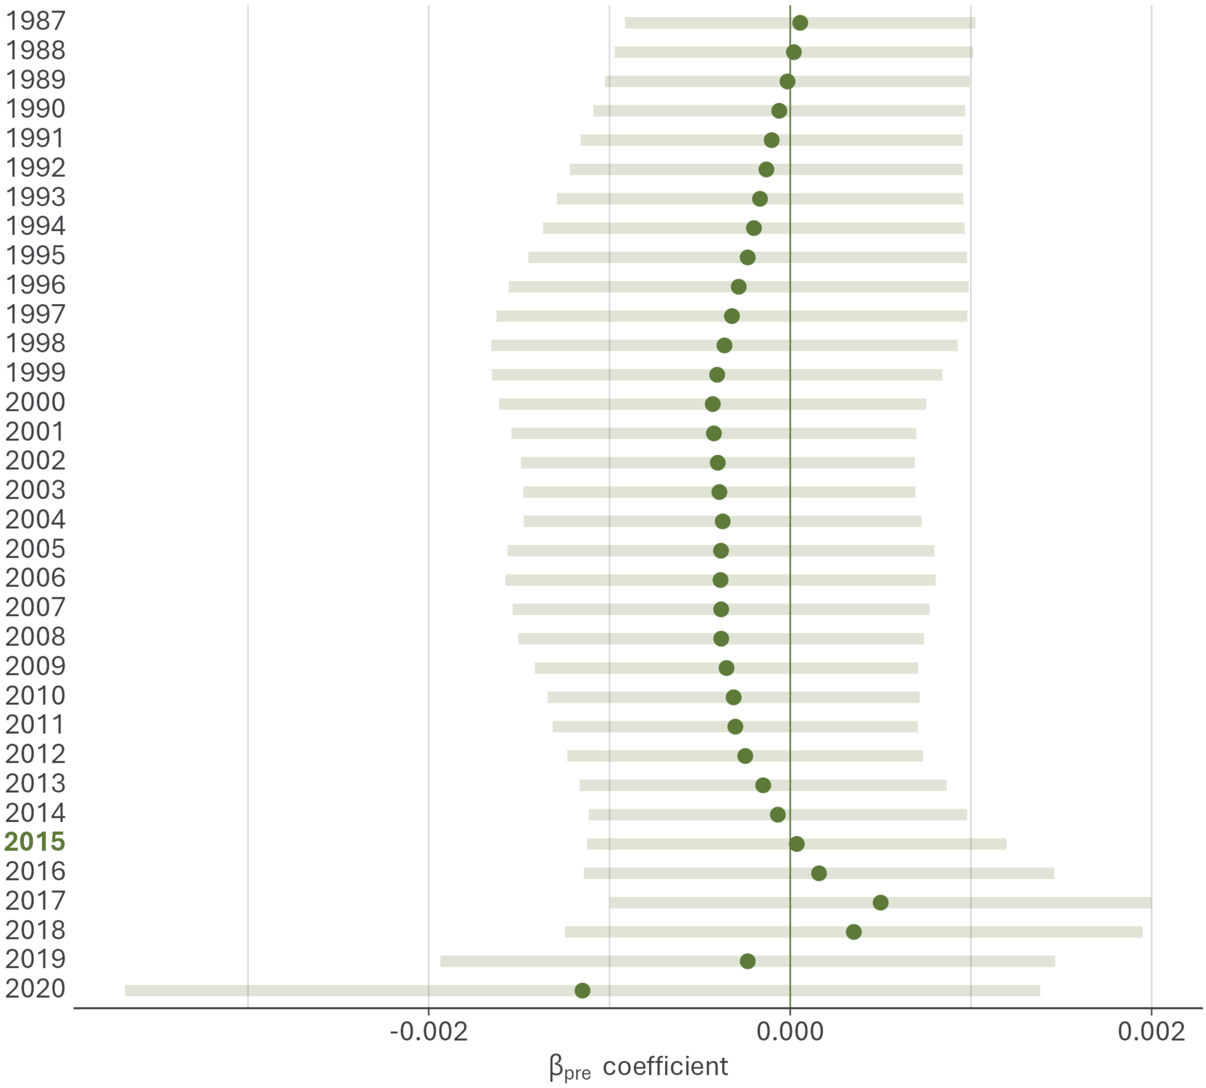

Note: Bars represent 95% confidence intervals. Bars which cross zero are not statistically significant at the 5% level.
Sources: ABS Labour Force (Detailed), JSA (2025a) and DEWR calculations.

[115] There are only 10 quarters of data after February 2020 and before November 2022, meaning that it is not feasible to use a pre-treatment period which starts after the COVID-19 period while still retaining sufficient data for reliable estimation. We include a sensitivity test which truncates the end of the pre-treatment period at February 2020, and it produces estimates similar to the core model but with lower statistical significance and magnitude of the $\Delta\beta$ term. The estimate is larger and more precisely estimated when the COVID-19 period is included.

Controlling for variable sensitivity to business cycle movements

Our model includes time fixed effects, which means that macroeconomic shocks that affect employment across-the-board will not affect our estimate of the AI effect on jobs. But there could be macroeconomic shocks that affect some occupations more than others. If sensitivity to macroeconomic shocks is correlated with AI exposure, this will bias our estimates. For this reason, we test whether controlling for variable sensitivity to business cycle movements impacts our core model estimate of $\Delta\beta$.

First, in a separate model, we estimate a set of pre-ChatGPT business cycle sensitivity scores by fitting a model of employment growth rates dependent only upon the unemployment rate and occupation interactions. Given the existence of zero employment quarters for some occupations, we approximate employment growth for occupation i at time t by $G_{i,t}$ defined as:

$$G_{i,t} := \log(1 + Emp_{i,t}) - \log(1 + Emp_{i,t-1})$$

The first-stage model, which is estimated using ordinary least squares with fixed effects, simply involves occupation-level interactions between the unemployment rate U_t at time t with an occupation dummy Occ_i, as well as occupation fixed effects α_i:

$$G_{i,t} = \alpha_i + \sum_i \beta_i^U \cdot Occ_i \cdot U_t$$

The key parameters from this model are the β_i^U coefficients, which capture how changes in the unemployment rate impact the growth rate for occupation i. The estimated β_i^U coefficients reflect the variable sensitivity to the business cycle[116] in the pre-treatment period. These interaction terms are relative to the ANZSCO unit group 1111 which corresponds to 'Chief Executives and Managing Directors'. A negative β_i^U means that occupation i has growth which declines more than that of Chief Executives and managing Directors for a given increase in the unemployment rate. In other words, negative β_i^U occupations are relatively pro-cyclical. Positive β_i^U occupations are relatively counter-cyclical and have relatively higher employment growth when the unemployment rate increases. The variation of estimates of β_i^U is shown in Figure 49.

Figure 49: Estimates from the business cycle sensitivity model show variable employment growth responses to the unemployment rate

Distribution of estimates of occupation-unemployment rate interaction terms

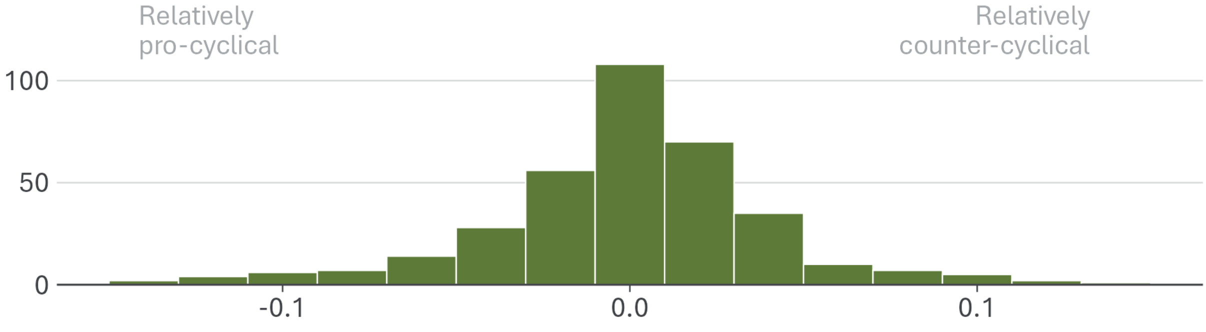

Sources: ABS Labour Force (Detailed), JSA (2025a) and DEWR calculations.

[116] While the unemployment rate may seem to be a crude proxy of the business cycle, the estimation procedure is only sensitive to movements in the unemployment rate. Assuming any constant non-accelerating inflation rate of unemployment (NAIRU) and measuring the unemployment rate relative to the NAIRU would produce identical estimates of β_i^U.

The second stage of the robustness test controlling for variable sensitivity to the business cycle then involves taking the core model equation and adding in an extra term where the β_i^U terms are interacted with the unemployment rate U_t for all time periods:

$$Emp_{i,t} = \exp\left(\alpha_i + \gamma_t + \beta_{pre} \cdot s_i t_{pre} + \beta_{post} \cdot s_i t_{post} + \beta_{shift} \cdot s_i Post_t + \beta_{cycle} \cdot \left(\beta_i^U \cdot U_t\right)\right)$$

The extra term β_{cycle} captures how the previously estimated business cycle sensitivity scores β_i^U together with the unemployment rate U_t contribute to employment ($Emp_{i,t}$). This aims to control for the expected changes in employment levels due to business cycle effects, in both the pre-treatment and post-treatment periods, to better isolate the impacts of AI automation exposure.

The $\Delta\beta$ estimate for this robustness test is shown in Figure 39 as ‘Control for business cycle sensitivity’. It is of a comparable size to the $\Delta\beta$ from the main model and still significant at the 5% level.[117] The added β_{cycle} term, which captures how changes in the unemployment rate and occupation-specific sensitivity to the business cycle explain employment changes, is statistically significant[118] and has the expected sign (positive). Taken together, this means that while changes in the unemployment rate explain some of the changes in employment levels in both the pre-treatment and post-treatment period, the $\Delta\beta$ effect is in addition to the effect of β_{cycle}. In other words, the negative association between AI exposure and employment since the introduction of ChatGPT is not being driven by business cycle effects, at least not in the way occupations historically responded to the business cycle.

[117] p=0.017.
[118] p=0.043.

Appendix D: Dissimilarity index

To summarise how quickly the occupational composition of employment changes over time, we use a *dissimilarity index* (sometimes called a Duncan index).[119] The results are shown and discussed on page 25. This appendix sets out further details about the construction of this dissimilarity index.

Definition

The index compares the distribution of employment across occupations in quarter t with the distribution in quarter $t-1$.

Let $Emp_{i,t}$ be employment in occupation i in quarter t, and define the employment share:

$$S_{i,t} = \frac{Emp_{i,t}}{\sum_{i=1}^{I} Emp_{i,t}}$$

The dissimilarity index between consecutive quarters is:

$$D_t = \frac{1}{2}\sum_{i=1}^{I}\left|S_{i,t} - S_{i,t-1}\right|$$

We report $100 \times D_t$, so the index ranges from 0 to 100%.

Interpretation

The index is the minimum share of employment that would need to shift across occupations for the occupational distribution in quarter t to match the distribution in quarter $t-1$. A value of zero implies no change in the occupational mix from the previous quarter; higher values indicate faster reshuffling across occupations.

The index does not measure the rate at which individual workers change from one occupation to another – it simply compares each occupation's share of total employment in one quarter with the previous quarter.

Data

We use employment-by-occupation data from the ABS *Labour Force (Detailed)* release. Further information about this data source is set out at Appendix A.

We compute the index using both:

- 4-digit ANZSCO occupation unit groups (more detailed)
- 2-digit ANZSCO occupation sub-major groups (broader).

The index is based on employment shares, so it is not driven by changes in total employment levels. However, the level of occupational detail matters: more detailed occupation definitions (e.g. 4-digit) yield higher index values. The dissimilarity index calculated with 4-digit data reflects changes both across and within 2-digit occupations, whereas the index calculated with 2-digit data only reflects the former.

We calculate the dissimilarity index on a balanced panel of occupational employment data. In other words, where there are quarters in which an occupation does not appear in the data, we add it in with a value of 0, so all occupations are included in all quarters. The quarterly employment-by-occupation data is not seasonally adjusted.

[119] See Duncan and Duncan (1955).

Smoothing

Employment by detailed occupation is subject to sampling variability, particularly at the 4-digit level and for demographic disaggregations (such as age). It is also subject to seasonal variation.

To reduce short-term volatility, we present a 4-quarter moving average of the index, calculated from the unsmoothed quarter-to-quarter index values:

$$\widetilde{D}_t = \frac{1}{4}\sum_{k=0}^{3} D_{t-k}$$

We do not apply occupation-level suppression or trimming rules. All occupations in the data are included in the index as reported.

Limitations

The dissimilarity index is a descriptive summary of changes in the occupational distribution. It is presented as a diagnostic cross-check on other measures included in this report.

It does not identify which occupations drive changes, does not measure within-occupation task change, and does not provide causal evidence about the drivers of occupational reshuffling (including AI).