

ESRI Working Paper No. 808

August 2025

The Impact of AI Exposure on Labour Market Outcomes and Well-Being: Evidence from Australia

Juan Duran-Vanegas^{a,b} and Dora Tuda^{a,b}

a: Economic and Social Research Institute (ESRI)

b: Department of Economics, Trinity College Dublin

* Corresponding author:

Juan Duran-Vanegas

Economic and Social Research Institute (ESRI)

Whitaker Square, Sir John Rogerson's Quay, Dublin 2

Email: juan.duranvanegas@esri.ie

ESRI working papers represent un-refereed work-in-progress by researchers who are solely responsible for the content and any views expressed therein. Any comments on these papers will be welcome and should be sent to the author(s) by email. Papers may be downloaded for personal use only.

The Impact of AI Exposure on Labour Market Outcomes and Well-Being: Evidence from Australia*

Juan Duran-Vanegas^{1, 2} and Dora Tuda^{1, 2}

¹The Economic and Social Research Institute (ESRI)

²Department of Economics, Trinity College Dublin

August 2025

Abstract

This paper examines the impact of artificial intelligence (AI) exposure on labour market outcomes and well-being in Australia. Using longitudinal microdata from the Household, Income and Labour Dynamics in Australia (HILDA) survey, we examine the impact of AI exposure on employment and job satisfaction outcomes from 2002 to 2022. We find that AI exposure led to a reduction in usual and desired work hours and an increase in hourly wages. AI exposed individuals also experienced lower satisfaction with job security and overall job satisfaction, but increased perceived autonomy. These adverse well-being effects are most pronounced among casual workers, whereas AI exposure appears to increase employment and wages for women and work hours for older workers. No significant effects are found on self-reported health. Our findings are robust to alternative AI exposure definitions, sample restrictions, and controls for pandemic-related labour shocks.

Keywords— Artificial Intelligence, Labour Market, Well-being

JEL Codes— I10, J01, J28, O30

*This paper uses unit record data from Household, Income and Labour Dynamics in Australia Survey [HILDA] conducted by the Australian Government Department of Social Services (DSS). The findings and views reported in this paper, however, are those of the author[s] and should not be attributed to the Australian Government, DSS, or any of DSS' contractors or partners. DOI: <https://doi.org/10.26193/R4IN30>. Corresponding author: Juan Duran-Vanegas, email: juan.duranvanegas@esri.ie

1 Introduction

The rapid advancement of artificial intelligence (AI) is transforming labour markets worldwide, raising critical questions about its impact on employment, wages, job satisfaction, and overall worker well-being. AI, as a general-purpose technology, has the potential to affect nearly every occupation, reshaping the nature of work and redefining skill requirements ([Albanesi et al., 2024](#)). Unlike previous waves of automation, which predominantly affected routine and low-wage jobs, AI exposure is expected to increase with wage, education, and age ([Webb, 2019](#)). However, evidence on the effects of AI exposure on labour market outcomes and well-being is still scarce.

This paper examines the effects of AI exposure on employment, work hours, wages, job satisfaction, and perceived health in Australia, using longitudinal microdata from the Household, Income and Labour Dynamics in Australia (HILDA) survey spanning 2001-2022, combined with AI exposure index developed by [Tolan et al. \(2021\)](#). The main advantage of this index is that it accounts for cognitive tasks within occupations and is not specifically developed for one country – both resulting in a measure of AI exposure less prone to bias. We follow the same individuals on the labour market, over more than 20 years, therefore accounting for unobservable heterogeneity constant over time.

The paper contributes to the literature on the labour market effects of AI in several ways. First, we study the effects of AI exposure on labour market outcomes using within-individual variation over time, thereby accounting for time-invariant unobservable heterogeneity through the inclusion of individual fixed effects. Second, our work offers a novel contribution by leveraging a panel of individuals to estimate these effects at the individual level. Australia presents a compelling case study, consistently ranking above average in global indices of AI influence and potential. For example, it ranks 8th among the 26 leading AI creator countries according to [Chakravorti et al. \(2023\)](#), based on metrics such as data volume and complexity, regulatory frameworks, capital investment, and innovation. Furthermore, Australia performs well in AI-related indicators, including patents granted per capita, employment rates in AI roles, and the prevalence of university-level AI pro-

grams, as reported by [Maslej et al. \(2024\)](#). In addition, HILDA provides rich data on workers' satisfaction and opinions ranging from work hours, pay, job security, work-life balance, stress levels, complexity, intensity, and autonomy - a feature that allows us to deepen our understanding on the well-being effects of AI compared to the existing literature.

We find that individuals in AI-exposed occupations report a decrease in usual and desired hours worked, with an increase in hourly wage rate. In addition, we find that AI exposure also negatively affects satisfaction with job security and overall job satisfaction, despite increasing perceived autonomy in how work is performed. Heterogeneity analysis reveals that the largest negative effects on work-related satisfaction are concentrated among casual workers. In contrast, AI exposure appears to benefit certain groups: it increases employment probabilities and wages for women, and raises work hours for workers over the age of 30. Finally, we find no significant effects of AI exposure on job opinions or self-reported health, relative to non-exposed individuals.

Our findings are robust when we consider alternative definitions of AI exposure, conduct placebo tests, restrict the sample to individuals exposed prior to 2010 to account for potential sorting, and compare the effects of exposure to reduced working hours or redundancy resulting from the COVID-19 pandemic.

Our paper is related to a growing body of research on AI and labour market dynamics, which thus far finds mixed results on the effects of AI exposure on labour market outcomes. Using establishment-level exposure in the U.S. between 2010-18, [Acemoglu et al. \(2022\)](#) finds that the employment and wage growth effects are too small to be detectable. Similarly, [Georgieff & Hyee \(2021\)](#) find no relationship between AI exposure and employment growth using occupation-level exposure measures from [Felten et al. \(2018\)](#) and [Felten et al. \(2019\)](#) for 23 OECD countries over the period 2012-19. In contrast, [Albanesi et al. \(2024\)](#) find that employment shares have increased on average in occupations more exposed to AI, particularly in those occupations with a relatively higher proportion of younger and skilled workers, using occupation-level exposure measures from [Webb \(2019\)](#) and [Felten et al. \(2018\)](#).

The paper also relates to a more limited number of studies on the effects of automation tech-

nologies on workers' well-being and mental health (Giuntella et al., 2025; Abeliansky et al., 2024; Patel et al., 2018). Similar to Abeliansky et al. (2024) and Patel et al. (2018), we find evidence of adverse effects of AI exposure on workers' perception of job security and the work itself. However, our results suggest that the negative impacts on mental and general health are not statistically significant.

The rest of the paper is structured as follows. Section 2 describes the data. Section 3 introduces the empirical strategy. Section 4 presents our results. Section 5 concludes and discusses the policy implications of our findings.

2 Data

In our analysis, we use Household, Income and Labour Dynamics in Australia (HILDA) Survey and AI exposure index developed by Tolan et al. (2021).

2.1 AI Exposure Index

As a measure of AI exposure, we use the index developed by Tolan et al. (2021). It combines 59 generic tasks from the European Working Conditions Survey, Survey of Adult Skills and occupational database O*Net, to 14 cognitive abilities. Together, these are mapped to 328 AI-related benchmarks, which indicate progress in AI techniques.

The main advantage of this approach is that it captures AI's potential to affect and transform tasks. On the one hand, AI may substitute an already standardised task, and on the other, it may transform a task by acquiring new abilities. The authors argue that since AI cannot acquire knowledge or a skill, abilities are a better measure of progress in AI, and are thus less prone to measurement error.

In addition, Tolan et al. (2021) develop the index for multiple European countries, so it is not country-specific, as opposed to, for example Brynjolfsson et al. (2018), Felten et al. (2021), Webb

(2019). An index that is not country specific is less likely to produce biased results when applied to a third country.

Tolan et al. (2021) develop the AI index on the 3-digit occupation level using the International Standard Classification of Occupations (ISCO). Since HILDA collects occupations on the 2-digit level defined by the Australian and New Zealand Standard Classification of Occupations (ANZSCO), we proceed as follows. First, we translate the ANZSCO classification to the ISCO classification, using the tool developed by the Australian Bureau of Statistics.¹ Second, we calculate the average AI index for all the 3-digit ISCO occupations linked to each 2-digit ANZCO occupations. Finally, we rescale the indices so that they take values from 0 to 1 and define AI-exposed ANZCO occupations as those above the median of the index (0.55464). Table 1 shows key differences and metrics of the 3-digit and 2-digit indices, of those that we could match with ANZSCO.² Occupations with the lowest AI exposure are street vendors in the 3-digit classification and its equivalent in the 2-digit ANZSCO classification – van salesperson. Their exposure index is 0.0472 and 0, respectively. Occupations with the maximum exposure, equal to 1, are database and network professionals and ICT security specialists. On average, AI exposure is 0.54285 in the 3-digit classification and 0.53868 in the 2-digit classification. In both classifications, the median is slightly higher than the mean of AI exposure.

Table 1: Comparison of 2- and 3-digit AI exposure index

		ISCO (3-digit)		ANZSCO (2-digit)
Min	0.04702	952 Street vendors	0	62 Van salesperson
Max	1	252 Database and network professionals	1	26 ICT security specialist
Mean	0.54285		0.53868	
Median	0.55986		0.55464	

Notes: This table compares ISCO and ANZCO occupation classifications, when 3-digit classification is averaged up to 2-digit. AI exposure index is re-scaled to values [0,1]

¹<https://www.abs.gov.au/AUSSTATS/abs@.nsf/DetailsPage/1220.02013,%20Version%201.2?OpenDocument>, accessed on 21/02/2025.

²Of 321 3-digit ISCO occupations, we are able to match 304 to the ANZSCO classification.

2.2 Survey Data – HILDA

HILDA is a longitudinal household survey, available from 2001, when 7,600 households (more than 17,000 individuals over the age of 15) were interviewed. In 2011, there was a 2,000 sample top-up, with retention rate in wave 18 of 76%.

We focus on a sample of employed and unemployed individuals aged 18-67, from 2001 to 2022.³ HILDA collects detailed socio-economic information, including hourly wage, self-reported job satisfaction and security, as well as self-reported health based on the 36-Item Short Form Survey (SF-36), which we use to answer our research question.

In our sample, we keep individuals who are observed as employed in at least 2 consecutive waves. The sample median of participation waves is 7. Hourly wage is constructed from the self-reported weekly wage and weekly usual hours worked. We exclude individuals for which we observe hourly wage rate above the 99th percentile of the wage rate distribution. Figure A1 in Appendix A shows a histogram of the constructed hourly wage.

We categorize individuals' occupation codes into exposed and non-exposed given the previously defined groups we constructed based on the AI exposure median value. For employed individuals, we categorize using the current occupation, whereas for unemployed we do it with their latest employment.⁴

Table 2 shows descriptive statistics of key outcome and control variables in our sample, for both AI-exposed and non-exposed individuals. Overall, around half of our sample is female and married, and on average have 0.86 children in their household. Of those that are AI-exposed, 48% are female, they are on average 2 years older than those non-AI exposed, and have a higher number of children in the household. AI-exposed are less casually employed, they work 4 hours per week more than non-exposed, and have 47% higher hourly wage rate compared to non-AI exposed. In terms of health and job satisfaction, AI exposed and non-exposed are fairly similar, with the

³Eligibility for the Australian Age Pension begins at age 67.

⁴In particular, we use the last reported occupation whenever this information is available within the three previous years before the beginning of the unemployment spell.


exception of social functioning, where AI-exposed have higher social functioning.

Table 2: Summary Statistics

	(1) Full sample			(2) AI exposed			(3) Non-exposed		
	Mean	SD	Obs.	Mean	SD	Obs.	Mean	SD	Obs.
Female	0.517	0.50	203,159	0.481	0.50	117,433	0.566	0.50	85,726
Age	39.606	12.87	203,159	40.740	12.10	117,433	38.053	13.72	85,726
Married	0.494	0.50	203,159	0.554	0.50	117,433	0.412	0.49	85,726
Bachelor or higher	0.303	0.46	203,159	0.448	0.50	117,433	0.103	0.30	85,726
Secondary or below	0.179	0.38	203,159	0.102	0.30	117,433	0.286	0.45	85,726
Number of kids in the household	0.855	1.12	203,159	0.926	1.12	117,433	0.758	1.10	85,726
Employment	0.963	0.19	203,159	0.999	0.03	117,433	0.913	0.28	85,726
Casual employment	0.166	0.37	203,159	0.081	0.27	117,433	0.282	0.45	85,726
Weekly usual hours	36.814	14.34	195,191	38.070	12.87	117,111	34.931	16.11	78,080
Weekly desired hours	34.641	10.97	78,521	33.995	10.14	46,529	35.581	12.02	31,992
Hourly wage	29.981	15.76	171,761	34.036	16.78	105,502	23.524	11.30	66,259
Satisfaction - overall	7.696	1.59	195,289	7.729	1.51	117,180	7.647	1.69	78,109
General health	71.164	18.53	181,894	72.198	18.27	106,653	69.697	18.81	75,241
Mental health	74.253	16.28	182,683	75.045	15.60	107,069	73.131	17.13	75,614
Role-emotional	92.531	18.59	169,440	92.827	18.27	99,826	92.106	19.03	69,614
Social functioning	85.835	19.94	182,680	87.073	19.09	107,075	84.081	20.95	75,605
Vitality	60.658	18.66	182,159	60.801	18.59	106,761	60.455	18.76	75,398
Observations	203159			117433			85726		


Further, Figures 1 – 3 plot the mean value of our binary AI exposure variable by wage percentile, age and educational attainment. AI exposure is lowest at the bottom of the wage distribution, for youngest individuals and lowest educational attainment. In contrast, AI exposure is highest at the top of the wage distribution, in the late 20s and early 30s and for university graduates and higher.

Figure 1: Exposure to AI by Occupational Wage Percentile




Notes: The graph plots mean occupational-level exposure to AI by occupational wage percentile rank using a locally weighted smoothing regression. Wage percentiles are measured as the employment-weighted percentile rank of an occupation's mean hourly wage in the HILDA 2016 wave.

Figure 2: Exposure to AI by Age



Notes: The graph plots a binscatter of mean occupational-level exposure to AI by individuals' age groups and a quadratic fit line (dashed).

Figure 3: Exposure to AI by Education Attainment



Notes: The graph plots the mean occupational-level exposure to AI by education attainment.

3 Empirical Strategy

We follow [Giuntella et al. \(2025\)](#) and exploit the longitudinal variation within individuals and across time using Differences-in-Differences (DiD) and event study design to identify the effects of AI exposure.

We start by estimating a DiD with the following specification:

$$Y_{iot} = \alpha + \beta_0 Post_t + \beta_1 AI_{ot} + \beta_3 AI_{ot} \times Post_t + \gamma X_{iot} + \theta_i + \lambda_t + u_{iot}, \quad (1)$$

where Y_{iot} is the outcome of interest for individual i , in occupation o and year t . AI_{ot} is a dummy that takes the value of 1 if the occupation o of individual i at time t is exposed and 0 otherwise. $Post_t$ is a dummy variable equal to 1 after 2016. X_{iot} is a vector of individual controls that includes age, number of kids in the household, a dummy for marriage, and a dummy for tertiary education. θ_i are individual fixed-effects controlling for time-invariant individual heterogeneity, while λ_t are

survey year fixed-effects absorbing common year variation across individuals. u_{iot} is an error term.

Our choice to define 2016–2022 as the post-treatment period and 2011 as the base year in our difference-in-differences and event study specifications deserves more consideration. While the term AI has gained significant traction in recent years, early indicators of rising AI activity can be traced back to the mid-2000s. For example, growth in AI-related academic publications and patents began as early as 2005, and the takeoff in AI-related course enrollments and conference participation occurred between 2010-13 (Shoham et al., 2018).

Our coefficient of interest is β_3 which captures the differential effect between exposed and non-exposed individuals after 2016, using within individual variation across time.

In order for β to recover the causal effect of the AI exposure, we require no differential pre-trends between the treated (exposed to AI) and control (non-exposed) groups. To assess the validity of the parallel trends assumption underlying our DiD design, and to explore the dynamics of the treatment effect over time, we estimate the following event study specification shown in Equation (2).

$$Y_{iot} = \alpha + \sum_{t=2002, t \neq 2011}^{2022} \alpha_t L_t + \sum_{t=2002, t \neq 2011}^{2022} \beta_t AI_{ot} \times L_t + \gamma X_{iot} + \theta_i + \epsilon_{iot}, \quad (2)$$

where L_t is a set of survey year dummies from 2002 to 2022, taking 2011 as a reference period, ϵ_{iot} is an error term, and all the remaining terms are defined as above. In this specification, we interact year dummies L_t with the AI exposure variable AI_{ot} . The coefficients β_t trace out the difference in outcomes between exposed and non-exposed occupations for each year relative to the omitted base year (2011). This approach allows us to examine two key features: (i) whether any differential trends between treated and control groups were already present before the treatment period, and (ii) how the effect of AI exposure evolves after 2016, providing a dynamic view of the treatment effect.

If the identifying assumption holds, the interaction coefficients β_t for years prior to 2016 should be statistically indistinguishable from zero, indicating no pre-treatment divergence between ex-

posed and non-exposed groups.

We then turn to an analysis of heterogeneous effects across different demographic groups including gender, college graduates, age, and casual (non-permanent) workers defined as those who are not entitled to paid holidays or sick leave. To do so, we employ a triple difference estimator according to the following specification:

$$Y_{iot} = \alpha + \beta_0 Post_t + \beta_1 AI_{ot} + \beta_3 AI_{ot} \times Post_t + \beta_4 G_i + \beta_5 Post_t \times G_i + \beta_6 AI_{ot} \times G + \beta_7 AI_{ot} \times Post_t \times G_i + \gamma X_{io} + \theta_i + \lambda_t + e_{iot}, \quad (3)$$

where G_i is a dummy variable that equals 1 if the individual i belongs to a certain demographic group, e_{iot} is an error term, and all the remaining terms are defined as above. The coefficient of interest in this case is β_7 which captures the effect of AI exposure for each demographic group compared to its base category.

On the labour demand side, U.S. job postings began to increase around 2010, with a marked acceleration from 2015–2016, particularly in roles requiring deep learning skills ([Acemoglu et al., 2022](#); [Shoham et al., 2018](#)). Although data on the incidence of AI-related jobs in the labour market compiled by [Maslej et al. \(2024\)](#) portrays the US as an early adopter with higher AI labour demand penetration compared to Australia (1.62% vs 1% of all job postings in 2023), selecting 2016 as the start of the post-treatment period allows us to capture the early labour market impacts of AI adoption. Moreover, using 2011 as the base year in the event study further informs this choice as it provides a sufficiently long window to detect any earlier shifts and anchors the analysis in a period before substantial AI-related labour market changes began to emerge.

4 Results


4.1 Baseline

We start by presenting the DiD estimation results. Fig. 4 plots the point estimates and confidence intervals of the effect of the coefficient β_3 from Eq. 3. Non-dichotomous outcomes are standardised for ease of interpretation. The figure reports that after 2016, individuals in AI-exposed occupations experienced lower probability of employment, work hours, and wages than non-exposed individuals. At the extensive margin, AI exposure has a negative and statistically significant effect on the probability of employment of about 1.8 percentage points. At the intensive margin, we find negative and statistically significant effects of AI exposure on weekly usual worked hours, with a point estimate of -0.063 standard deviations (approximately 3.44%), and desired worked hours, with a point estimate of -0.041 standard deviations (1.53%). We find a small negative effect on weekly wages of about -0.023 standard deviations (1.74%) and a larger and positive effect on hourly wages of 0.052 standard deviations (2.68%). As hourly wage is constructed as total wages divided by usual hours, a larger increase in hourly wage is driven by a reduction in hours, rather than an increase in total wage.

We find negative and sizable effect of AI exposure on the satisfaction with work hours (0.06 standard deviations, 12.58%) and having a secure future at the job (0.03 standard deviations, 6.38%). This suggests income effect prevailing over substitution effect: as AI exposure increases hourly wage, usual and desired hours decrease but so does the satisfaction with work hours. At the moment it is unclear whether a decrease in hours satisfaction came before or after a decrease in usual hours, which we decompose below using the event study approach.

As for job opinions and perceived health, we observe no significant differences between AI-exposed workers and non-exposed workers after 2016, except for exposed workers reporting higher agreement on having more freedom to decide how to perform the job.

Figure 4: Effects of Exposure to AI - DiD Estimates



Notes: The graph plots the point estimates and 95% confidence intervals of the coefficient β_3 from Eq. 3 for the corresponding outcome in the y-axis. Non-dichotomous outcomes are standardised. Confidence intervals are computed using standard errors clustered at the individual level.


We then present the event study results in order to explore the dynamics of the effects over time. Given our large number of outcomes, we focus our attention on the outcomes for which we find statistically significant DiD estimates. The estimation results reported in Fig. 5 provides evidence supporting the parallel trending assumption as most coefficients are statistically insignificant before 2011. The estimates show the differential effects between AI-exposed and non-exposed individuals after 2011. The results for employment and hourly wages seem to temporarily reverse in 2020, while the positive effects on reported freedom on how to do the perform the job are only statistically significant in that year.

In addition, the negative effect of AI exposure on usual hours worked is statistically significant after 2018, whereas desired hours start to decrease earlier, around 2012. It seems that usual hours worked follow the pattern of desired hours with a time lag for workers exposed to AI. On the

other hand, these workers have a decreased job satisfaction with work hours from 2018 and this dissatisfaction with hours worked matches the usual hours pattern. We do not find matching effects on desired hours worked from 2018 onward.


Finally, we discuss our heterogeneity analysis. [Fig. 6](#) reports the triple-difference estimation results. According to these estimates, AI-exposed casual workers were more affected post-2016 than AI-exposed non-casual workers in terms worked hours and hourly wages. Moreover, older workers have been less affected by the decline in worked hours than younger workers, while employment and hourly wages have increased more for women than men. In contrast, we find little evidence for heterogeneous effects among the studied groups in regards of satisfaction about work hours and job security.

Figure 5: Effects of Exposure to AI - Event Study



Notes: The graph plots the point estimates and 95% confidence intervals of the coefficients β_7 from Eq. 2. Confidence intervals are computed using standard errors clustered at the individual level.

Figure 6: Heterogeneity in the Effects of Exposure to AI - Triple-Difference Estimates



Notes: The graph plots the point estimates and 95% confidence intervals of the β coefficients from Eq. 3. Confidence intervals are computed using standard errors clustered at the individual level. Non-dichotomous outcomes are standardised. Casual workers defined as those who are not entitled to paid holidays or sick leave.

4.2 Robustness

As a first robustness check, we estimate placebo event study regressions using two outcomes that are unlikely to be affected by short-term exposure to AI: physical-functioning health scores and the average number of hours worked from home per week. The estimated coefficients, shown in [Fig. A2](#), do not exhibit any discernible patterns or trends following the reference year, providing reassurance that our main findings are not driven by spurious correlations.

Further, it is possible that individuals on the extremes of the AI exposure index distribution are affected differently than those in the middle. To test this, we re-define the AI treatment variable as a dummy that takes the value of 1 for occupations in the top tercile of [Tolan et al. \(2021\)](#) index, and 0 for occupations in the bottom tercile. Estimation results presented in [Fig. A3](#) and [Fig. A4](#) show that our findings are robust to this alternative definition of the treatment variable, with the exception of negative effects on desired work hours which are no longer statistically significant. Compared to our baseline results, individuals in the top tercile of the AI exposure index also report higher agreement with their job requiring them to learn new skills and working very intensively, as well as higher vitality scores than individuals in the bottom tercile after 2016.

An important identification concern is that, even in the absence of pre-trends, individuals with certain unobserved characteristics simultaneously affecting our outcomes of interest could have sorted into exposed occupations in response to the AI shock. Although including individual effects helps to address this concern about unobserved characteristics biasing our estimates, we explore the robustness of our results by using a similar identification strategy as [Giuntella et al. \(2025\)](#) and defining an alternative AI exposure measure that equals 1 if the individual had an exposed occupation before 2010 and restricting the sample to individuals who joined the market before that year. Estimation results presented in [Fig. A5](#) and [Fig. A6](#) imply that our findings on the effects of AI exposure on worked hours, wages, satisfaction about work hours and job security, and perceived role-physical health are robust to using this alternative sample. In contrast, effects on the probability of employment and desired worked hours are not statistically significant. Interestingly, the DiD estimates reported in [Fig. A5](#) reveal that on top of the previous findings, AI-exposed individuals

in this restricted sample report lower overall satisfaction, satisfaction with the work itself, and perceived mental health scores, alongside higher disagreement with the job being complex and often requiring to learn new skills than non-exposed individuals post-2016.

Finally, we test whether the estimated effects reflect COVID-19 induced changes in the workplace rather than changes resulting from the exposure to AI. In order to disentangle both effects, we use self-reported COVID-19 labour incidence to construct a homologous index of COVID exposure. Specifically, we compute the share of individuals reporting that *as a result of the coronavirus, kept working, but with reduced hours* at the 2-digit ANZSCO code level. When looking at employment effects, we construct this index using the share of individuals reporting that *as a result of the coronavirus, employment terminated or made redundant*. We compare both (standardized) indices between AI exposed and non-exposed occupations in [Table A1](#), which shows that AI-exposed occupations had a lower COVID-19 incidence both in terms of reduced hours and redundancy. With these COVID-19 indices at hand, we define a binary variable that takes the value of 1 if the share for that occupation bin is above the median and re-estimate the event study regressions following [Eq. 2](#) setting 2018 as the base year. The estimated effects are significantly different when comparing the AI exposure vs the COVID-19 incidence impacts on the employment probability, usual and desired worked hours, hourly wages, and satisfaction with work hours and wages. At the extensive margin, the trend of the differential employment probability between COVID-19 exposed and non-exposed occupations starts decreasing after 2020 and rebounds thereafter, while it decreases earlier and it is more pronounced when comparing AI exposed vs non-exposed occupations. At the intensive margin, we find no robust evidence of differential effects between COVID-19 exposed and non-exposed occupations in terms of hourly wages and desired worked hours, while the differential effect on satisfaction with usual work hours and job security tends to increase for COVID exposed vs non-exposed occupations, the opposite effect we found for AI exposure.

5 Conclusion

This paper examined the effects of AI exposure on employment, work hours, wages, job satisfaction in Australia, using longitudinal microdata from the Household, Income and Labour Dynamics in Australia (HILDA) survey spanning 2002-2022.

Our difference-in-differences estimates show that AI exposure is associated with a modest but statistically significant decline in employment probability and a reduction in weekly hours worked, both usual and desired, accompanied by a large increase in hourly wages.

Beyond labor market outcomes, we find that AI exposure negatively affects satisfaction with work hours and perceived job security. Interestingly, we find no significant effects on broader job satisfaction or perceived health, although AI-exposed workers report greater autonomy in how they perform their jobs.

To assess the identifying assumption and study effect dynamics, we conduct an event study analysis. We find no significant pre-trends before 2011, supporting the parallel trends assumption. Post-2016, the negative employment and hours effects become more pronounced, especially after 2018. Desired hours begin declining as early as 2012, while actual hours and satisfaction follow with a lag.

Finally, heterogeneity analyses reveal that casual workers experience stronger negative effects of AI exposure on hours and hourly wages, while older workers are somewhat shielded from reductions in hours compared to younger individuals. Women benefit more than men in terms of employment and hourly wage growth, consistent with prior findings on gendered complementarities with AI [Cazzaniga \(2024\)](#). In contrast, college-educated individuals do not appear to experience significantly more favorable outcomes than non-college individuals.

References

- Abeliansky, A. L., Beulmann, M., & Prettnner, K. (2024). Are they coming for us? industrial robots and the mental health of workers. *Research Policy*, 53(3), 104956. Retrieved from <https://www.sciencedirect.com/science/article/pii/S0048733324000052> doi: <https://doi.org/10.1016/j.respol.2024.104956>
- Acemoglu, D., Autor, D., Hazell, J., & Restrepo, P. (2022). Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics*, 40(S1), S293–S340.
- Albanesi, S., Dias da Silva, A., Jimeno, J. F., Lamo, A., & Wabitsch, A. (2024, 12). New technologies and jobs in europe*. *Economic Policy*, 40(121), 71-139.
- Brynjolfsson, E., Mitchell, T., & Rock, D. (2018). What can machines learn, and what does it mean for occupations and the economy? *AER Papers and Proceedings*, 108, 43–47.
- Cazzaniga, M. (2024, 06). Exposure to artificial intelligence and occupational mobility: A cross-country analysis. *IMF Working Papers*, 2024, 1. doi: 10.5089/9798400278631.001
- Chakravorti, B., Bhalla, A., & Chaturvedi, R. S. (2023). The emerging geography of ai. *Digital Planet, The Fletcher School at Tufts University*. Retrieved from <http://data.worldbank.org/indicator/SP.DYN.LE00.FE.IN>
- Felten, E., Raj, M., & Seamans, R. (2018, May). A method to link advances in artificial intelligence to occupational abilities. *AEA Papers and Proceedings*, 108, 54–57.
- Felten, E., Raj, M., & Seamans, R. (2019). The occupational impact of artificial intelligence: Labor, skills, and polarization. *NYU Stern School of Business*.
- Felten, E., Ray, M., & Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42, 2195–2217.
- Georgieff, A., & Hyee, R. (2021). Artificial intelligence and employment. *OECD Social, Employment and Migration Working Papers*(265). Retrieved from <https://www.oecd-ilibrary.org/content/paper/c2c1d276-en> doi: <https://doi.org/https://doi.org/10.1787/c2c1d276-en>
- Giuntella, O., Konig, J., & Stella, L. (2025). Artificial intelligence and the wellbeing of workers. *Scientific Reports*, 15(1), 20087.
- Maslej, N., Fattorini, L., Perrault, R., Parli, V., Reuel, A., Brynjolfsson, E., ... Clark, J. (2024). *The AI Index 2024 Annual Report* (Tech. Rep.). AI Index Steering Committee, Institute for Human-Centered AI, Stanford University, Stanford.
- Patel, P. C., Devaraj, S., Hicks, M. J., & Wornell, E. J. (2018). County-level job automation risk and health: Evidence from the united states. *Social Science & Medicine*, 202, 54–60.

Shoham, Y., Perrault, R., Brynjolfsson, E., Clark, J., Manyika, J., Niebles, J. C., ... Baue, Z. (2018). *The AI Index 2018 Annual Report* (Tech. Rep.). AI Index Steering Committee, Institute for Human-Centered AI, Stanford University, Stanford.


Tolan, S., Pesole, A., Martínez-Plumed, F., Fernández-Macías, E., Hernández-Orallo, J., & Gómez, E. (2021). Measuring the occupational impact of ai: tasks, cognitive abilities and ai benchmarks. *Journal of Artificial Intelligence Research*, 71, 191–236.

Webb, M. (2019). The impact of artificial intelligence on the labor market. *Available at SSRN 3482150*.

Appendix

A Data

Figure A1: Hourly Wage Histogram




Notes: The sample are currently employed individuals who we observe as employed in at least two waves. We calculate the hourly wage rate from usual hours worked and their weekly wage. We drop individuals with hourly wage rate higher than the 99th percentile, producing this histogram over waves 2001 – 2022.

B Robustness

B.1 Placebo Test


Figure A2: Event Study - Placebo Tests



Notes: The graph reproduces the estimates reported in [Fig. 5](#) by running placebo regressions checking for effects on physical-functioning health scores and hours work on average each week at home.


B.2 Top Tercile Treatment

Figure A3: Effects of Exposure to AI - DiD Estimates - Top Tercile Treatment



Notes: The graph reproduces the estimates reported in [Fig. 4](#) by defining AI exposure for occupations in the top tercile of [Tolan et al. \(2021\)](#) index, and non-exposure for occupations in the bottom tercile.


Figure A4: Effects of Exposure to AI - Event Study - Top Tercile Treatment



Notes: The graph reproduces the estimates reported in Fig. 5 by defining AI exposure for occupations in the top tercile of Tolan et al. (2021) index, and non-exposure for occupations in the bottom tercile.


B.3 Pre-2010 Restricted Sample

Figure A5: Effects of Exposure to AI - DiD Estimates - Restricted Sample



Notes: The graph reproduces the estimates reported in [Fig. 4](#) by defining AI exposure if the individual had an exposed occupation before 2010 and restricting the sample to individuals who joined the labour market before 2010.

Figure A6: Effects of Exposure to AI - Event Study - Restricted Sample




Notes: The graph reproduces the estimates reported in Fig. 5 by defining AI exposure if the individual had an exposed occupation before 2010 and restricting the sample to individuals who joined the labour market before 2010.

B.4 COVID-19

Table A1: AI Exposure vs COVID-19 Incidence Indices

	(1)					(2)				
	Mean	AI exposed				Mean	Non-exposed			
		SD	P25	P50	P75		SD	P25	P50	P75
covid_	0.301	0.16	0.23	0.27	0.23	0.525	0.19	0.42	0.52	0.42
covid2_	0.137	0.10	0.06	0.09	0.06	0.364	0.25	0.18	0.30	0.18
Observations	117433					82046				

Figure A7: Effects of Exposure to AI vs COVID - Event Study



Notes: The graph plots the event studies estimates contrasting the AI exposure to a COVID exposure binary variable constructed at the 2-digit occupation level using self-reported reduced working hours or redundancy as a result of the pandemic in 2020. Base years are 2016 and 2018 for AI and COVID exposure treatments, respectively.