2025
Volume 28 | Number 1

AJLE
AUSTRALIAN JOURNAL OF LABOUR ECONOMICS

ajle.org

A journal of labour economics & labour relations

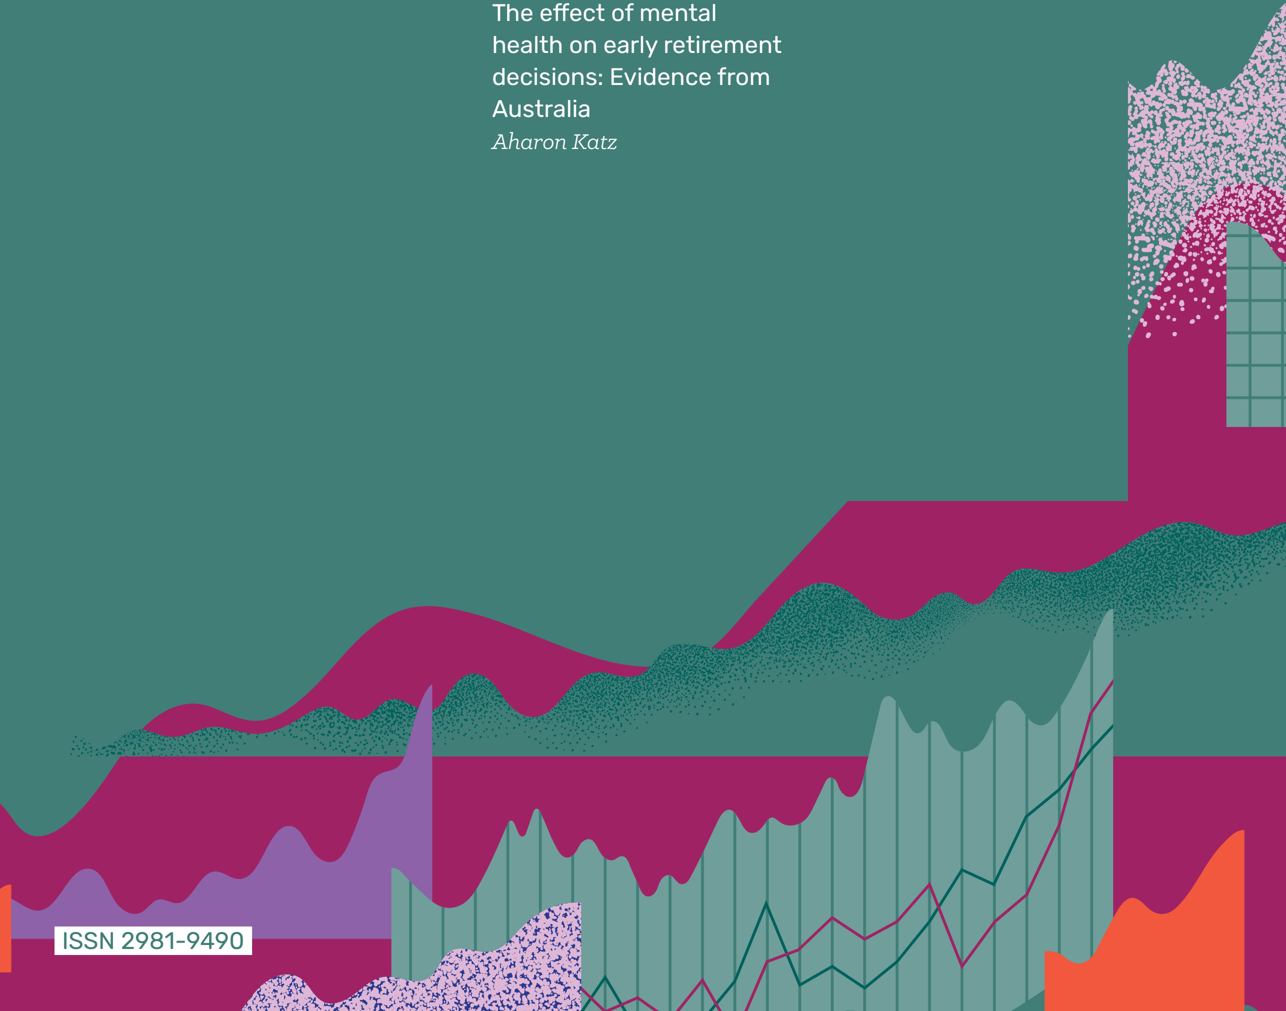

ISSN 2981-9490

AUSTRALIAN JOURNAL OF LABOUR ECONOMICS

A journal of labour economics & labour relations

ISSN 2981-9490
Official journal of the Australian Society of Labour Economists

Graphic Design
Advance Press

Contact details
Telephone: +61 8 9266 1744
Email: ajle@curtin.edu.au
Web: https://ajle.org

Contents

ISSN 2981-9490

Also available from INFORMIT LIBRARY at: http://search.informit.com.au
and PROQUEST LIBRARY at: http://www.proquest.com

From the Managing Editor

Welcome to the latest issue of the *Australian Journal of Labour Economics (AJLE).* In this issue we have, as usual, a range of articles which will be of interest to our readers covering a range of labour market issues and using a variety of approaches to research. The papers also have implications for policy and demonstrate the importance of research in formulating policy.

The first paper examines a subject which is of great interest currently. Claire Mason, Haohui Chen and David Evans of the Commonwealth Scientific and Industrial Organisation (CSIRO) consider firms' Artificial Intelligence (AI) adoption and firm demand for workers and skills.

This study attempts to clarify apparently conflicting findings from the literature. First, the authors specifically test whether trends in demand for workers and skills differ between AI adopting and non-adopting firms (controlling for firm characteristics such as size, industry and geography). Second, focusing on demand for workers and skills in non-AI skilled occupations, they determine how existing occupations are affected by AI adoption. Third, the above effects are examined at the occupational level, which allows a of test whether effects of AI adoption on demand for new workers and skills depends on the occupation's AI exposure. The findings suggest that AI adopting firms experienced slightly stronger growth in job postings than non-adopting firms of the same size, location and industry. In addition, the number of skills required in job postings increased slightly faster in AI adopting firms than in non-adopting firms. The authors conclude that workers who have the skills to use and complement AI remain sought after in an AI-augmented workforce.

The second paper by Aharon Katz, of York University, UK, investigates the relationship between mental illness and early retirement decisions in Australia. It is well-established that health and labour supply are interconnected but research has predominantly focused on the impact of physical rather than mental health. This study seeks to address this deficiency by examining the effect of mental health on early retirement decisions using data from the Household, Income, and Labour Dynamics in Australia (HILDA) Survey using a variety of statistical techniques to help validate and strengthen causal findings. The findings indicate a significant and positive causal impact of poor mental health on early retirement decisions.

The observed differences in the results between males and females support the assumption of gender specific effects. These findings suggest that poor mental health has a significant and potentially causal impact on premature exit from the labour market, particularly among men. The results highlight the importance of effective mental health management in supporting longer working lives.

The final paper, by Karen Mumford, Edith Aguirre, Anna Einarsdóttir, Bridget Lockyer, Melisa Sayli, and Benjamin A. Smith, of various UK universities and health organisations, explores the determinants, and differences, in reported job satisfaction for women, ethnic minority and LGBT+ employees among public sector health employees in the English National Health Service (NHS). A broad range of possible determinants are considered including demographic variables, job characteristics, and supportive workplace measures. The results indicate that women are more likely to be satisfied with their jobs, as are LGBT+ employees from ethnic minorities. There is evidence a higher wage is positively associated with job satisfaction, but relative wages are not consistently related to job satisfaction. In contrast, supportive workplace practices are strongly associated with higher rates of job satisfaction. Of particular importance are effective workplace anti-bullying polices and the presence of relevant minority staff networks, especially for those identifying as LGBT+. These results suggest that organisations can raise job satisfaction by further facilitating these supportive workplace practices.

Again, this issue has been greatly facilitated by my fellow editors and referees but, particularly by the *AJLE* editorial assistant, Sandie Rawnley. Many thanks.

Phil Lewis
Managing Editor

AI adoption and firm demand for workers and skills: New insights from Australia

CLAIRE MASON *Data61, Commonwealth Scientific and Industrial Organisation (CSIRO)*
HAOHUI CHEN *Data61, Commonwealth Scientific and Industrial Organisation (CSIRO)*
DAVID EVANS *Data61, Commonwealth Scientific and Industrial Organisation (CSIRO)*

Abstract

The latest Artificial Intelligence (AI) tools can perform some of the complex tasks that highly skilled and well-paid workers perform. In this study, we analyse an Australian dataset of firm job postings to explore the effects of AI adoption on demand for workers and skills. Our findings suggest that AI adoption increases demand for new workers and skills. Both the number of job postings and the number of skills required in job postings grew slightly faster in AI adopting firms than in non-adopting firms, after controlling for firm size, geography and industry. AI exposed occupations did not experience lower demand than non-exposed occupations, unless they were employed in a non-adopting firm. It is not clear whether the stronger demand for new workers in AI adopting firms is due to the need to replace existing workers with new workers (with new skills) or due to growth in the total number of workers. However, our findings counter fears about AI substituting for or deskilling workers and align with the view that the latest AI tools serve to augment, rather than substitute for, human capability.

Keywords: Employment, Labour Demand, Artificial Intelligence, Skills trends

Introduction

As the adoption of Artificial Intelligence (AI) accelerates, so too do concerns about the labour market impacts of AI. Advances in AI in fields such as natural language processing and image recognition mean that AI tools can now perform a range of non-routine cognitive tasks that are normally performed by highly skilled and well-paid workers (Brynjolfsson *et al.*, 2017; Felten *et al.*, 2019; Tolan *et al.*, 2021a; Webb, 2019). These tools are reigniting concerns about loss of jobs (Arntz *et al.*, 2016; Chen *et al.*, 2022; Priddis *et al.*, 2020; White House, 2022) and deskilling of workers (Shiohira, 2021). In this study, we investigate these concerns by comparing workforce and skills trends in firms that have adopted AI and firms that have not. Using a longitudinal job postings data from 4,395 Australian firms, and controlling for differences in firm size, geographic location and industry, we show that firms adopting AI are exhibiting stronger growth in (a) demand for new workers, and (b) demand for skills, than non-adopting firms. In addition, some formerly non-AI occupations in AI adopting firms are transitioning to become AI skilled roles.

AI and the labour market

AI refers to the capability of a system to perform human-like cognitive functions (learning, understanding, reasoning, and interacting) with the aim of obtaining rational outcomes [1, 2]. According to technology-skill complementarity theory, new technologies increase demand for workers whose skills complement the new technology. In the decades between 1980 and 2010, employment growth was concentrated in high-skill roles (Acemoglu and Autor, 2011). Computers and digital technologies made it easy to automate routine cognitive tasks (Autor *et al.*, 2003). In line with technology-skill complementary theory, it was the low and middle-skilled workers who performed these routine, cognitive tasks (e.g., clerks, bank tellers, cashiers) who were displaced in earlier waves of automation (Brynjolfsson *et al.*, 2018). High-skilled workers had technology-complementary skills (e.g., critical thinking, creativity, manual dexterity and interpersonal skills) that were still needed to perform those tasks that were not yet automatable. For these high-skilled workers, the new technology made them more productive. In addition, the technology created new service delivery opportunities, which in turn required new high-skilled occupations (e.g., software engineers, games developers, digital marketers) to deliver these services.

AI research has progressed to the point where AI tools can perform near to, or better than humans on image recognition, speech recognition, gaming and language translation tasks (Zhang *et al.*, 2023). Unlike previous forms of digital technology, AI tools can perform non-routine cognitive tasks, but they still do not match humans in tasks requiring physical ability (Brynjolfsson *et al.*, 2018), commonsense reasoning (Davis and Marcus, 2015) and metacognition (Eysenck and Eysenck, 2021; Tolan *et al.*, 2021a). Survey data captured by Australia's National AI Centre in 2024 suggests that 35 per cent of large Australian businesses (those with 500 or more employees) have adopted AI (Department of Industry Science and Resources, 2024). The adoption of AI is evident across all industry sectors, although it is more advanced in some sectors (e.g., health, education and manufacturing) than others (e.g., agriculture, forestry and fishing) (Department of Industry Science and Resources, 2024).

To understand how advanced AI tools will affect demand for workers, several groups of researchers have delineated AI capabilities and then investigated which occupations in the labour market traditionally perform tasks that require these capabilities (Brynjolfsson *et al.*, 2018; Felten *et al.*, 2021, 2018a; Webb, 2019). Occupations that perform tasks that are compatible with current AI capabilities are considered to be more 'exposed' to AI than occupations that perform tasks that AI is not yet capable of performing (Acemoglu *et al.*, 2022b). Although their estimates of occupational exposure to AI differ slightly, these researchers all conclude that highly educated and well-paid workers, usually those with high-level cognitive skills (e.g., genetic counsellors, actuaries, teachers, language translators) are most exposed to the latest wave of AI-enabled automation (Brynjolfsson *et al.*, 2017; Eloundou *et al.*, 2023; Felten *et al.*, 2018b; Tolan *et al.*, 2021b; Webb, 2019).

Nevertheless, the workforce impacts of AI adoption are difficult to predict because of the uneven rate of development in AI capabilities and the difficulty of unbundling those aspects of a task that an AI can perform from those that still require a human to perform (Brynjolfsson *et al.*, 2018). As long as a human is required to perform some aspect of the task, there may be no advantage in using AI to perform the other aspects of the task. Alternatively, if the AI allows the human to work with the AI to perform the task more efficiently, more safely or in a better way, it can be a source of competitive advantage for the firms and workers that use it.

This study investigates whether the adoption of AI has more of a substitution effect (wherein the technology substitutes for the human in performing a task) or more of an augmentation effect (wherein using the technology in a task increases productivity in other tasks, thereby enhancing overall labour productivity) (Loaiza and Rigobon, 2024). One way in which we can answer this question is to monitor how demand for high-skilled workers changes when firms adopt AI. If AI substitutes for workers, demand for workers (especially AI exposed workers) should decline in AI adopting firms relative to non-adopting firms. Alternatively, if AI is being used to augment workers, demand for workers should be as strong, or stronger, in AI adopting firms relative to non-adopting firms.

The skills requirements listed in job postings offer another source of insight into AI's augmentation versus substitution effects. If AI is simply being used to automate some of the tasks that were performed in an exposed occupation, we might see a reduction in skills requirements for the occupation because the worker now only needs to perform those aspects of their work that the AI cannot yet perform. However, if AI is augmenting human worker with its computational power and memory (Loaiza and Rigobon, 2024; Schleiger *et al.*, 2024), thereby providing better data or additional insights for humans to draw upon, AI can elevate our human capabilities (Sadiku *et al.*, 2021; Yau *et al.*, 2021). In this context, AI adoption should enlarge the skills profile of an occupation, since workers remain responsible for their AI-augmented work but now need the skills to work with AI to deliver greater value (Boulus-Rødje *et al.*, 2024; Grimberg and Mason, 2025).

Several studies have already been carried out to investigate the relationship between AI adoption and demand for workers. At the occupational level, researchers report that workers in occupations that are more exposed to the latest advances in AI are experiencing increased demand (Albanesi *et al.*, 2023; Alekseeva *et al.*, 2021; Felten *et al.*, 2019; Green and Lamby, 2023) and even increased wages (Felten *et al.*, 2019; Fossen and Sorgner, 2019). A cross-country study found that demand for highly exposed occupations was only enhanced if workers were in an occupation where computer use is high (Georgieff and Hyee, 2022). However, given that firms adopting AI are still in the minority (Acemoglu *et al.*, 2022a; Borgonovi *et al.*, 2023; Nguyen and Hambur, 2023; Zolas *et al.*, 2020), current trends in demand for AI exposed workers are not likely to reveal the effects of AI unless the analyses focus on AI adopting firms.

Other researchers have investigated workforce impacts of AI by monitoring hiring trends in firms that were recruiting workers with AI relevant skills. Recruitment of workers with AI skills is used as an indicator that a firm is adopting AI. Acemoglu *et al*. (2022b) used job postings to identify firms that were adopting AI and then examined trends in demand for non-AI workers within these firms. They found that after adopting AI, firms posted fewer job advertisements for non-AI workers. Furthermore, when AI adopting firms did advertise roles for non-AI workers, the job postings for AI exposed occupations showed more change in the skills and knowledge sought from these workers.

Babina *et al*. (2024) adopted a slightly different approach, characterising firms according to the proportion of the workforce that were AI skilled. Babina *et al*. found that firms with a high proportion of AI skilled workers showed more growth in employment than firms with a low proportion of AI skilled workers. However, Babina *et al*. did not differentiate between AI skilled and non-AI skilled workers. The positive effects on employment reported by Babina *et al*. may be due to growth in demand for AI skilled workers counterbalancing the negative effects of AI on demand for non-AI workers that were reported by Acemoglu *et al*. (2022b).

Contribution of this study

This study clarifies these conflicting findings. First, we specifically test whether trends in demand for workers and skills differ between AI adopting and non-adopting firms (controlling for firm characteristics such as size, industry and geography). Second, we focus on demand for workers and skills in non-AI skilled occupations, to determine how *existing* occupations are affected by AI adoption. Third, we examine the above effects at the occupational level, so that we can test whether effects of AI adoption on demand for new workers and skills depends on the occupation's AI exposure.

We hypothesised that:

> H1: Firms adopting AI will experience more growth in demand for new workers than non-adopting firms.
>
> H2: Occupations that are more exposed to AI will experience more growth in demand for new workers than occupations that are less exposed to AI
>
> H3: Firms that are adopting AI will experience more growth in demand for skills than firms not adopting AI
>
> H4: The effect of occupational exposure to AI on demand for new workers and skills will be moderated by firms' AI adoption status

Next, the study methodology is described.

Method

Datasets

Job postings

A national database of online job postings was obtained from the labour market platform provider, Adzuna Australia. Adzuna Australia aggregate online job ads from more than a thousand sources in Australia. Their sources include job ads listed directly on the Adzuna Australia platform, ads listed in Australia's major newspapers and cross-postings from other recruitment platforms. The representativeness of the Adzuna Australia job ads has been established through comparison of the geographic and occupational composition of the job ads and the geographic and occupational composition of the Australian labour market reported by the Australian Bureau of Statistics (Evans *et al.*, 2023b). The trends over time in Adzuna Australia online job postings also align with the nationally representative Australian Bureau Statistics (ABS) Job Vacancy Survey (JVS) (Duenser and Mason, 2019). Adzuna Australia's job postings have been used in other studies exploring workforce and skills trends in the Australian labour market (Bratanova *et al.*, 2022; Evans *et al.*, 2024, 2023a; Mason *et al.*, 2023; Zhao *et al.*, 2021) and the coverage of this dataset closely matches the coverage of the Litecast (formerly Burning Glass) job postings (Evans *et al.*, 2023b).

With duplicate job ads (Zhao *et al.*, 2021) and advertisements for unpaid (voluntary) or commission-only roles removed, the Adzuna Australia dataset contained 9,550,441 job postings for the period from 1 January 2016 to 31 December 2023. Employers were identified in job ads using a manually compiled dictionary of 7,372 employer names. This dictionary uses unique text patterns to differentiate the employer name from names of other organisations (or locations) mentioned in job postings that could create false classifications. Using this dictionary, employer names were matched to 25 per cent of the job postings.

The analyses investigating trends in demand for workers and skills are carried out on the subset of job postings that could be matched to an employer. To understand the representativeness of this sample, we compared the geographic, industry and occupation profile of these job postings with the geography, industry and occupational characteristics of the Australian labour force at the time of the 2021 Census. The proportion of workers (according to the 2021 Census (Australian Bureau of Statistics, 2021)) and job postings in each major industry division, major occupation group, and greater capital city statistical area is compared in Table 1, Table 2 and Table 3 respectively. These comparisons reveal that the matched job postings over-represented firms in Transport, Postal and Warehousing and Healthcare and Social Assistance, relative to other industry divisions. From an occupational perspective, machinery operators and drivers and labourers were over-represented. Finally, firms in Sydney and Melbourne were over-represented relative to firms in other locations.

Table 1. Industry profile of job postings compared with 2021 Census employment data

	% of employment (2021 Census)	% of job postings
Agriculture, Forestry and Fishing	2	0
Mining	2	5
Manufacturing	6	3
Electricity, Gas, Water and Waste Services	1	1
Construction	9	3
Wholesale Trade	3	2
Retail Trade	10	9
Accommodation and Food Services	7	5
Transport, Postal and Warehousing	5	15
Information Media and Telecommunications	1	2
Financial and Insurance Services	4	7
Rental, Hiring and Real Estate Services	2	3
Professional, Scientific and Technical Services	8	7
Administrative and Support Services	3	0
Public Administration and Safety	7	8
Education and Training	9	4
Health Care and Social Assistance	15	24
Arts and Recreation Services	2	1
Other Services	4	1

Table 2. Occupation profile of job postings compared with 2021 Census employment data

	% of employment (2021 Census)	% of job postings
Labourers	9	14
Machinery Operators and Drivers	6	33
Sales Workers	8	8
Clerical and Administrative Workers	13	10
Community and Personal Service Workers	12	8
Technicians and Trades Workers	13	8
Professionals	24	16
Managers	14	3

Table 3. Geography of job postings compared with 2021 Census employment data

	% of employment (2021 Census)	% of job postings
Australian Capital Territory	2	3
Rest of NT	0	0
Greater Darwin	1	1
Rest of Tas.	1	0
Greater Hobart	1	0
Rest of WA	2	2
Greater Perth	8	9
Rest of SA	2	1
Greater Adelaide	5	4
Rest of Qld	10	4
Greater Brisbane	10	4
Rest of Vic.	6	3
Greater Melbourne	19	28
Rest of NSW	11	8
Greater Sydney	21	32

AI exposure of occupations

Felten *et al*.'s (2018b) estimates of AI exposure for occupations were used in this study. Felten *et al*. used the Electronic Frontier Foundation (EFF) AI Progress Measurement dataset, which assesses progress in AI in different fields (e.g., image recognition) from blog posts, academic literature and websites. To understand what abilities are used in different occupations, they use the Occupational Information Network (O*NET; National Center for O*NET Development, 2022) database, developed by the US Department of Labour, which identifies 52 distinct abilities, matched to occupations in terms of how important the ability is to the relevant occupation. EFF AI domains were mapped to O*Net abilities to assess the relative effect of advances in AI technology on the different abilities and thus, aggregate across all abilities at the occupation level to create an AI exposure score for each occupation. These occupational exposure scores are available on Github at https://github.com/AIOE-Data/AIOE.

We then translated the AI exposure scores for US Occupations to Australian occupations (ANZSCO four-digit and six-digit codes), using the ANZSCO to International Standard Classification of Occupations (ISCO) and the ISCO to Standard Occupational Classification (SOC) cross-walks to match six-digit ANZSCOs to SOC codes. Due to differences in the granularity of ANZSCO and SOC, some SOC codes map to ANZSCO four-digit codes rather than six-digit codes while others map to ANZSCO four-digit codes. To deal with this inconsistency, AIOE scores for each six-digit ANZSCO were averaged and then assigned to the relevant four-digit ANZSCO code. Using this method, AIOE scores were assigned to 336 of the 358 four-digit ANZSCO occupations. The matches were then reviewed and edited manually, using the SOC and ANZSCO look-up functions, to ensure that matches derived from the cross-walks and aggregation process aligned with similarities in occupation descriptions.

IBISWorld

We also captured firm workforce size data from the IBISWorld Australian Enterprise Profiles database (Australia Enterprise Profiles Report 6287). The IBISWorld Enterprise Wizard provides information on leading Australian and New Zealand enterprises, including ASX and NZX listed companies; public, private and foreign-owned companies; local, state and federal government departments (IBISWorld, 2025). Focusing on the largest 1000 enterprises, we searched for matches between firm names in the IBISWorld list and firm names from the job postings. Matches were confirmed by checking industry classifications (available from both the job postings and the IBISWorld database) were aligned. In those instances where IBISWorld firm statistics had not been updated in 2024, they were excluded from the analyses. Using this method, we were able to obtain workforce size and growth data for 373 of the firms identified from the Adzuna Australia job postings. We compared the profile of firms that were matched to the IBISWorld database on study measures. We found that large enterprises were over-represented in the matched dataset but in terms of industry, geography and AI adoption, the profile of matched firms did not differ significantly from the profile of unmatched firms. To reduce the influence of outliers, the four largest firms were assigned a maximum score that was slightly larger than the fifth largest firm.

Measures

Firm status (AI adopter vs non-adopter)

AI skilled job postings have been used by several researchers to differentiate between firms that are adopting AI and those that are not (Acemoglu *et al*., 2022b; Alekseeva *et al*., 2021; Borgonovi *et al*., 2023; Bratanova *et al*., 2022). Since the use of AI technology requires specialised skills, demand for AI skills serves as an indicator of firm adoption of AI (Alekseeva *et al*., 2021). The OECD AI skills dictionary (Borgonovi *et al*., 2023) was used to identify AI skilled job postings. The great majority of the skills words in this list overlap with the dictionaries used by Acemoglu *et al*. (2022b) and other researchers studying AI adoption and AI skilled workers (Alekseeva *et al*., 2021; Green and Lamby, 2023). However, the OECD AI skills list is more stringent because it differentiates between generic and specific AI skills. To be identified as an AI skilled job posting, the job posting must contain at least one specific AI skill word (e.g., 'visual image recognition') or two or more generic AI skill words (e.g., 'autonomous driving' and 'artificial intelligence'). However, checks of the data revealed that four of the specific AI skills ('boosting', 'torch', 'screen reader' and 'caffe') were generating a high rate of false positives when used on their own. Consequently, we chose to treat these as generic AI skill words rather than specific AI skill words. If any of a firm's job postings between 2016 and 2019 (T1) included a specific AI skill word or two or more generic AI skill words, the firm was classified as an AI adopter. Otherwise, the firm was classified as a non-adopter.

Firm characteristics (Industry, geography and size)

The industry, geography and size of each firm was also captured from the job postings so that they could be used as control variables. The geographic location of each firm was determined based on the modal location of the jobs being advertised. Location was classified using the ABS Greater Capital City Statistical Area system, which is designed to represent labour markets and the functional area of Australian capital cities (Australian Bureau of Statistics, 2018). Firm size was classified based on the number of job advertisements each firm posted in T1, with firms grouped into deciles[1].

1 In the analyses predicting job postings, the firm size deciles were not included because numbers of job postings (our proxy measure of firm size) were captured in the first step of the analysis.

Demand for workers and skills

To understand the impact of AI on existing occupations, Acemoglu *et al*. (2022b) examined trends in firms' non-AI job postings. That is, they separated out job postings that required AI skills[2]. We adopted the same approach to derive counts of the number of non-AI job postings for each firm in T1 (2016 – 2019) and T2 (2020 – 2023). Due to the high positive skewness of job postings, numbers of non-AI job postings were transformed using the same inverse hyperbolic sine transformation that Acemoglu *et al*. (2022b) used in their analyses.

To capture the number of skills mentioned in each job posting, we used the ESCO skills taxonomy (European Commission, 2019) which contains a dictionary of preferred and alternative labels for more than 13,000 hierarchically organised skills. Job postings were tagged with the relevant Level 2 ESCO skill if they contained the relevant preferred or alternative label. We then calculated the mean number of Level 2 ESCO skills associated with each firm's (or occupation class') job postings in T1 and T2.

Results

Firm adoption of AI

Figure 1 illustrates how the percentage of Australian job postings mentioning AI skills has been changing over time. Between 2016 to 2023, only 0.18 per cent of job postings mentioned AI skills. The OECD published statistics on the proportion of Australian job postings that mentioned AI skills between 2019 and 2022 (Borgonovi *et al*., 2023). When we focus on job postings for the same time period, we find that 0.26 per cent of job postings mentioned AI skills, which is similar to the OECD estimate of 0.30 per cent.

Figure 1 reveals a steady increase in demand for AI skills, despite some short-term fluctuations. Notably, during the period of COVID-19 shut-downs there was an increase in the proportion of AI skilled job postings. There was an additional spike in AI

2 Acemoglu *et al*. (2022b) also excluded firms from the information technology and professional and business services sector (NAICS 51 and 54) on the grounds that these firms were likely to be selling AI products or in the latter case, supporting the integration of AI in production processes. The major industry division ANZSIC classifications for the Adzuna Australia database are less granular so it was not possible to specifically exclude information technology and professional and business services firms. Instead, the analyses were run twice to check that the findings were consistent when all firms from Professional, Scientific and Technical Services major Industry division were excluded from the analysis. Having determined that the findings were consistent, we report the results for all firms.

skilled job postings when the labour market expanded again post-pandemic. It seems that the accelerated digitisation of product and service delivery driven by the pandemic (Calvino *et al*., 2024) heightened demand for AI skilled workers.

Figure 1. Percentage of job postings requiring AI skills each month (2016 – 2023)

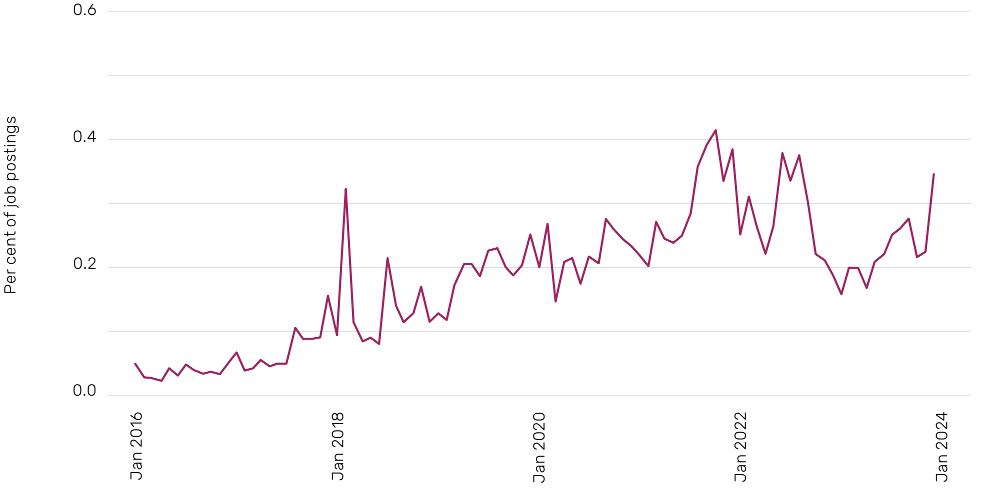

AI skills in job postings were used to differentiate firms that were AI adopters and firms that were not. Firms were categorised as AI adopters if they posted a job ad mentioning one specific or two generic AI skills between 2016 and 2019. Adopting this approach, we identified 322 firms that were adopting AI and 4,073 firms that were not adopting AI at T1.

The observed variability in firm AI adoption across industries and geographies was aligned with previous research. As Figure 2 illustrates, more than 15 per cent of firms in professional, scientific and technical services and financial and insurance services had adopted AI but fewer than 3 per cent of agriculture, forestry and fishing, construction and accommodation and food services firms had adopted AI. These findings are consistent with OCED's cross-country statistics which also found that firms in professional and ICT industries had the most AI skilled job postings and firms in utilities, agriculture, transport, real estate and construction had relatively few AI skilled job postings. The concentration of AI adopting firms in metropolitan regions also aligns with Bratanova *et al*.'s (2022) analysis of Australian AI clusters, which drew upon a wider range of datasets (AI companies, patents and job postings).

Figure 2. Variability in AI adoption by industry and geography at T1 (N denotes the total number of firms represented in the job postings dataset for each industry or geographic location)

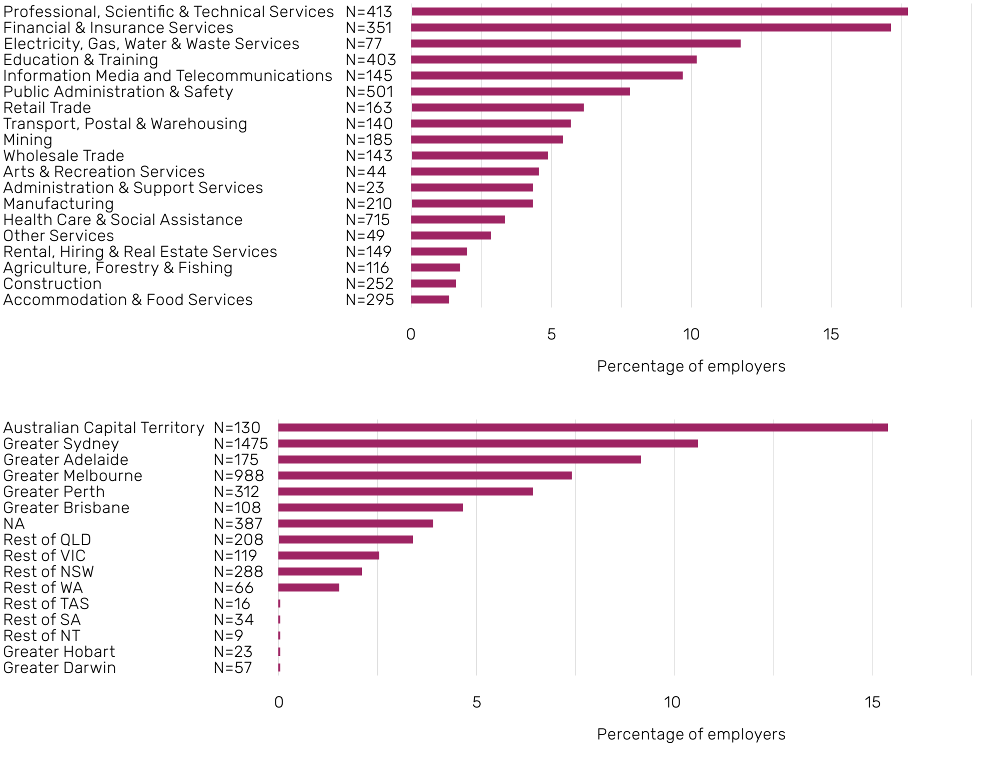

Firm AI adoption and demand for new workers

One of our objectives was to understand whether firms that had adopted AI showed different trends in demand for non-AI workers than firms that had not adopted AI.

To test the effect of AI adoption on demand for non-AI workers, counts of non-AI job postings were captured for firms in both T1 (2016 to 2019) and T2 (2020 to 2023).[3] Figure 3 is based on the counts of job postings for both AI adopting (red) and non-adopting (green) firms at T1 and T2. The steeper regression line for the AI adopting firms reveals that these firms showed a stronger increase in demand for new non-AI workers between T1 and T2.

3 Since the time lag between employer hiring of AI workers and subsequent impacts on demand for non-AI workers is unknown, we tested two alternative timeframes for both T1 (2016-2018 and 2016-2020) and T2 (2019 to 2023 and 2021 to 2023). The findings from these analyses were substantively the same.

Figure 3. Counts of job postings at T1 and T2 within AI adopting and non-adopting firms

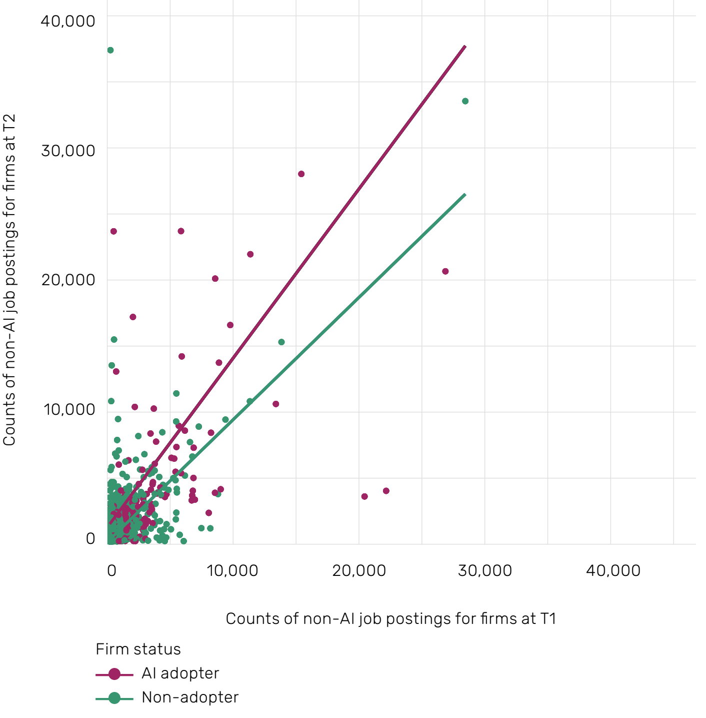

The next step was to determine whether the faster growth in demand for new (non-AI) workers in AI adopting firm was statistically significant after controlling for differences in firm geography and industry. The following model was tested:

$$\text{POSTINGS}_{i,T2} = \beta_0 + \beta_1\text{POSTINGS}_{i,T2} + \beta_2\text{Industry}_i + \beta_3\text{Geography}_i + \beta_4\text{AI}_i + e_i$$

where:

- $\text{POSTINGS}_{i,T2}$ is the (transformed) number of non-AI postings made by firm *i* at T2.
- $\text{POSTINGS}_{i,T1}$ is the (transformed) number of postings made by firm *i* at T1.
- AI_i is a binary variable denoting whether firm *i* was an AI adopter (coded 1) or not (coded 0).
- Industry_i is a series of dummy variables denoting the Industry classification (ANZSIC) of firm *i*.
- Geography_i is a series of dummy variables representing firm *i*'s primary geographic location (according to the ABS Greater Capital City Area classification).

- Observations are weighted by the firm's total job postings in T1, meaning that observations from larger firms were given more weight in estimating regression coefficients than were observations from smaller firms (following Acemoglu *et al.*'s (2022b) approach).
- e_i represents unexplained variance in job postings for firm *i* at T2.

In the first step of the analysis, T1 non-AI job postings were entered into the analysis as a predictor of T2 job postings. Controlling for T1 job postings meant that subsequent predictors added to the model were explaining the change in job postings (Cronbach and Furby, 1970; Edwards, 1994). The effects of firm industry and geography were tested in the second step of the analysis. The third step of the analysis was used to test whether the firm's AI adoption status explained change in numbers of job postings made by firms between T1 and T2.

Table 4 shows how the explanatory power of the model improved as additional variables were entered into the analysis. Supporting hypothesis 1, there was a statistically significant ΔR^2 at step 3 when the firm's AI adoption status was entered into the model. The regression weight for the binary variable representing the firm's AI adoption status was β_4 = 36.07 (LLCI = 25.076, ULCI = 47.07). This indicates that non-AI job postings grew 36 per cent faster for AI adopting firms than for non-adopting firms (holding other factors constant).

Table 4. Predicting change in numbers of non-AI postings at the firm level

Dependent variable: Firm non-AI postings at T_2			
Predictors in the model	ΔR_{adj}^2	df	F value
Step 1: Firm non-AI postings at T_1	0.52	1,3396	3646.00***
Step 2: + Firm industry, geography and size	0.02	33,3364	8.34***
Step 3: + Firm AI adoption status	0.01	34,3363	41.32***

*** p < .001

A limitation of this analysis is the use of job postings as a measure of demand for workers. Job postings reflect demand for *new* workers. If new workers were being recruited to replace existing workers who were being retrenched, the overall size of the workforce might remain unchanged or even decline.

We used the sample of 373 firms that could be matched to an IBISWorld enterprise to investigate whether the faster growth in demand for workers was replicated in the workforce data. This smaller dataset is not representative of Australian employers (it represents the largest Australian enterprises) and the sample was not large enough to allow us to control for the effect of firm industry, geography and size. As Figure 4 reveals, in this sample of large Australian enterprises, AI adopting and non-adopting firms appeared to be growing at the same rate. A regression analysis confirmed that AI adoption did not explain variance in workforce size in 2024 after controlling for workforce size in 2020, β_2 = -200 (LLCI = -866.95, ULCI = 399.27). However, the small and non-representative sample size prevents us from drawing a definitive conclusion from the non-significant result.

Figure 4. Workforce size in 2020 and 2024 for AI adopting and non-adopting firms

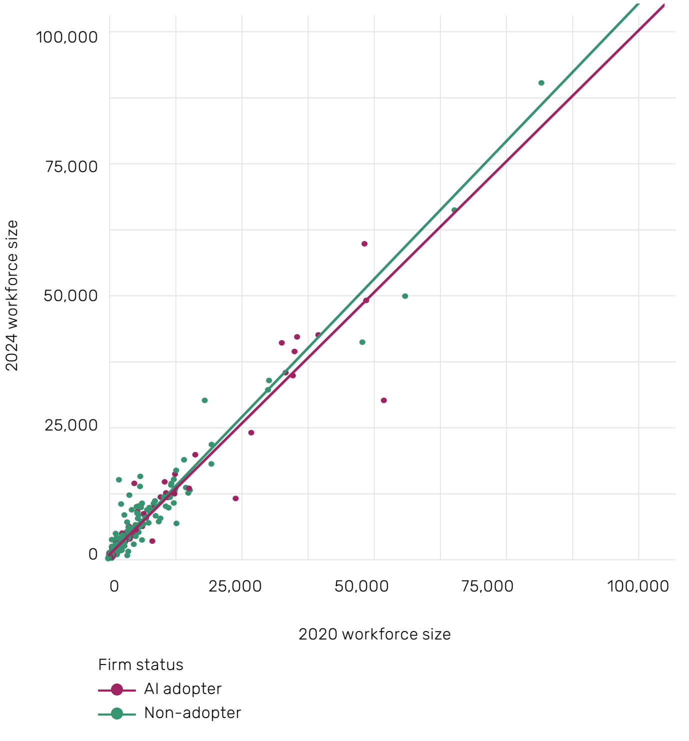

Occupational exposure

The job postings also allow us to examine whether hiring trends vary at the occupational-level, due to some occupations being more exposed to AI (H2). In addition, we wanted to investigate whether demand for skills was growing faster in AI exposed occupations (H3) that were employed in AI adopting firms (H4).

Although the goal was to explore trends in job postings and skills requirements for occupations, it was necessary to control for the effects of firm location, industry and size and to test the effect of firm AI adoption. Therefore, the dependent variable was the number of non-AI postings (or the mean skill count for these non-AI postings) at T2 for each occupation class (i.e., for each unique combination of occupation, firm AI adoption status, firm industry, firm geographic location and firm size).

Occupational AI exposure and demand for workers

The following model was used to test whether firm AI adoption and occupational AI exposure explained change in numbers of job postings for each occupation class:

$$O\_POSTINGS_{i,T2} = \beta_0 + \beta_1 \mathrm{O\_POSTINGS}_{i,T1} + \beta_2 \mathrm{Industry}_i + \beta_3 \mathrm{Geography}_i + \beta_4 \mathrm{Size}_i + \beta_5 \mathrm{AI}_i + \beta_6 \mathrm{AIOE}_i + \beta_7 \mathrm{Product}_i + e_i$$

where:

- $\mathrm{O\_POSTINGS}_{i,T2}$ is the number of non-AI postings for occupation class *i* at T2.
- $\mathrm{O\_POSTINGS}_{i,T1}$ is the number of non-AI postings for occupation class *i* at T1.
- $\mathrm{Industry}_i$ is a series of dummy variables denoting the Industry classification (ANZSIC major industry division) of the firms employing occupation class *i*.
- $\mathrm{Geography}_i$ is a series of dummy variables representing the primary geographic location (Greater Capital City Statistical Area) of the firms employing occupation class *i*.
- Size_i is a series of dummy variables denoting the firm size decile of the firms employing occupation class *i*.
- AI_i is a binary variable denoting the AI adoption status of the firms employing occupation class *i* (coded 1 if the firms were AI adopters and 0 if not).
- AIOE_i is the AI exposure score for occupation class *i*.
- $\mathrm{Product}_i$ represents the moderation effect (the product of AI_i and AIOE_i).
- Observations are weighted by the firm's total job postings in T1, meaning that observations from larger occupation classes were given more weight in estimating regression coefficients than were observations from smaller occupation classes.
- e_i represents unexplained variance in job postings for occupation class *i* at T2.

The hierarchical approach (controlling for T1 job postings in the first step of the analysis) was used again, so that subsequent predictors were explaining change in job postings for each occupation class.

Table 5 shows how the explanatory power of the model changed as effects were added. AI adoption was added to the model in the third step and explained significant incremental variance (supporting hypothesis one). The measure of AI exposure was added in step four but did not add explanatory power (hypothesis two was not supported). However, hypothesis four was supported. At step 5, the moderation effect added significant (albeit small) explanatory power to the model, β_7 = -6.81 (LLCI = -9.77, ULCI = -3.85). A simple slopes analysis revealed that there was a negative relationship between occupational AI exposure and growth in job postings in non-adopting firms, slope = −5.12 (t=−5.96, p<.001) but no relationship between AI exposure and growth in job postings in AI adopting firms, slope −0.49 (t=−0.34, p=.74). In other words, the lowest rate of growth in demand for new workers was experienced by AI exposed occupations in non-adopting firms. Among firms not using AI, a one-unit increase in AI exposure is associated with a 5.12 per cent slower rate of growth in job postings for the relevant occupation. Yet in AI adopting firms, the effect of AI exposure on demand for new workers is not significant.

Table 5. Predicting change in non-AI job postings for occupation classes

Dependent variable: Counts of (non-AI) occupational postings at T_2			
Predictors in the model	ΔR_{adj}^2	df	F value
Step 1: Non-AI occupational postings at T_1	0.77	1, 24759	83,000.00***
Step 2: + Firm industry, geography and size	0.03	41, 24718	1102.61***
Step 3: + Firm AI adoption status	0.01	1, 24717	289.42***
Step 4: + Occupational AI exposure	0.00	1, 24716	0.02
Step 5: + AI adoption * AI exposure	0.00	1, 24715	20.36***

*** p < .001

Occupational AI exposure and demand for skills

Our third hypothesis was that skills requirements would be increasing more in AI adopting firms and in our fourth hypothesis we predicted again that the effect of AI exposure on demand for skills would depend on whether the occupation was employed in an AI adopting firm. A second model was used to test these hypotheses. In this model, the dependent variable was the mean number of skills sought in non-AI job postings for an occupation class at T2. As before, the mean number of skills sought in T1 was entered into the analysis at step 1 so that subsequent predictors were explaining growth in skills requirements. The results of these analyses are reported in Table 6.

First, hypothesis three was supported in that AI adopting firms were exhibiting stronger growth in the number of skills required from occupations than were non-adopting firms, β_5 = 0.31 (LLCI = 0.24, ULCI = 0.37). In addition, the effect of AI exposure was significant and demand for skills was growing faster in occupations that were more exposed to AI, β_6 = 0.30 (LLCI = 0.27, ULCI = 0.33). However, the fourth hypothesis was not supported in this instance; the moderation effect was not significant, β_7 = -0.04 (LLCI = -0.02, ULCI = 0.10).

Table 6. Predicting change in skills per non-AI posting for occupation classes

Dependent variable: Mean skills in non-AI postings at T_2			
Predictors in the model	ΔR_{adj}^2	df	F value
Step 1: Skills per non-AI postings at T_1	0.13	1, 24675	3542.00***
Step 2: + Firm industry, geography and size	0.04	41, 24634	27.67***
Step 3: + Firm AI adoption status	0.00	1, 24633	100.66***
Step 4: + Occupational AI exposure	0.02	1, 24632	513.36***
Step 5: + Product (AI adoption * AI exposure)	0.00	1, 24631	1.66

*** $p < .001$

Occupational AI exposure and demand for AI skills

The stronger growth in skills sought in AI adopting firms and AI exposed occupations raises an interesting possibility. One of the new skills that workers in AI adopting firms might need is the ability to work with AI. However, since our methodology excludes job postings that require AI skills, the transition of formerly non-AI occupations into AI-skilled occupations would not be visible in the analyses. Furthermore, by excluding these job postings from our analysis, we could be underestimating the growth in demand for formerly non-AI workers in AI adopting firms.

To explore whether non-AI occupations were evolving to become AI skilled occupations, we revisited the job postings captured from the AI adopting firms, this time counting the number of AI skilled job postings for each occupation class at T1 and T2. As before, the counts of AI skilled job postings were transformed (inverse hyperbolic sine) to reduce the positive skew in the data.

As Table 7 reveals, after controlling for AI skilled job postings at T1 and the effects of industry, geography and firm size, occupational AI exposure was able to explain significant incremental variance in numbers of AI skilled job postings at T2. The regression coefficient for the effect of AI exposure, β_5 = 12.44 (LLCI = 9.82, ULCI = 15.006), indicates that a difference of one standard deviation in AI exposure was associated with 12.44 per cent more AI skilled postings at T2. That is, occupations were more likely to be transitioning to become AI skilled if they were more exposed to AI. However, although the effect was significant, it was small in magnitude.

Table 7. Predicting change in numbers of AI skilled job postings for occupation classes in AI adopting firms

Dependent variable: Counts of AI skilled postings at T_2			
Predictors in the model	ΔR_{adj}^2	df	F value
Step1: Occupational AI skilled postings at T_1	0.61	1, 6949	1,0790.00***
Step 2: + Firm industry, geography and size	0.11	36, 6913	79.33***
Step 3: + Occupational AI exposure	0.01	1, 6912	86.79***

*** p < .001

For more insight into this AI upskilling trend, the scatterplot in Figure 5 shows how numbers of AI skilled job postings for each occupation class varied between T1 and T2. The steep regression line indicates that numbers of AI skilled job postings were increasing across all occupation classes. Points that land on the y-intercept represent occupations that never required AI skills at T1. Points above the regression line represent occupation classes that were exhibiting faster than average increase in demand for AI skills. Occupation classes that were exhibiting particularly strong AI-upskilling effects have been labelled. Like Alekseeva *et al*. (2021), we find that the trend towards AI upskilling is strongest in IT occupations. Nevertheless, it can be seen in a wide range of occupations, including architects, accountants, sales representatives and security guards.

Figure 5. AI skilled postings at T1 and T2 for occupation classes in AI adopting firms

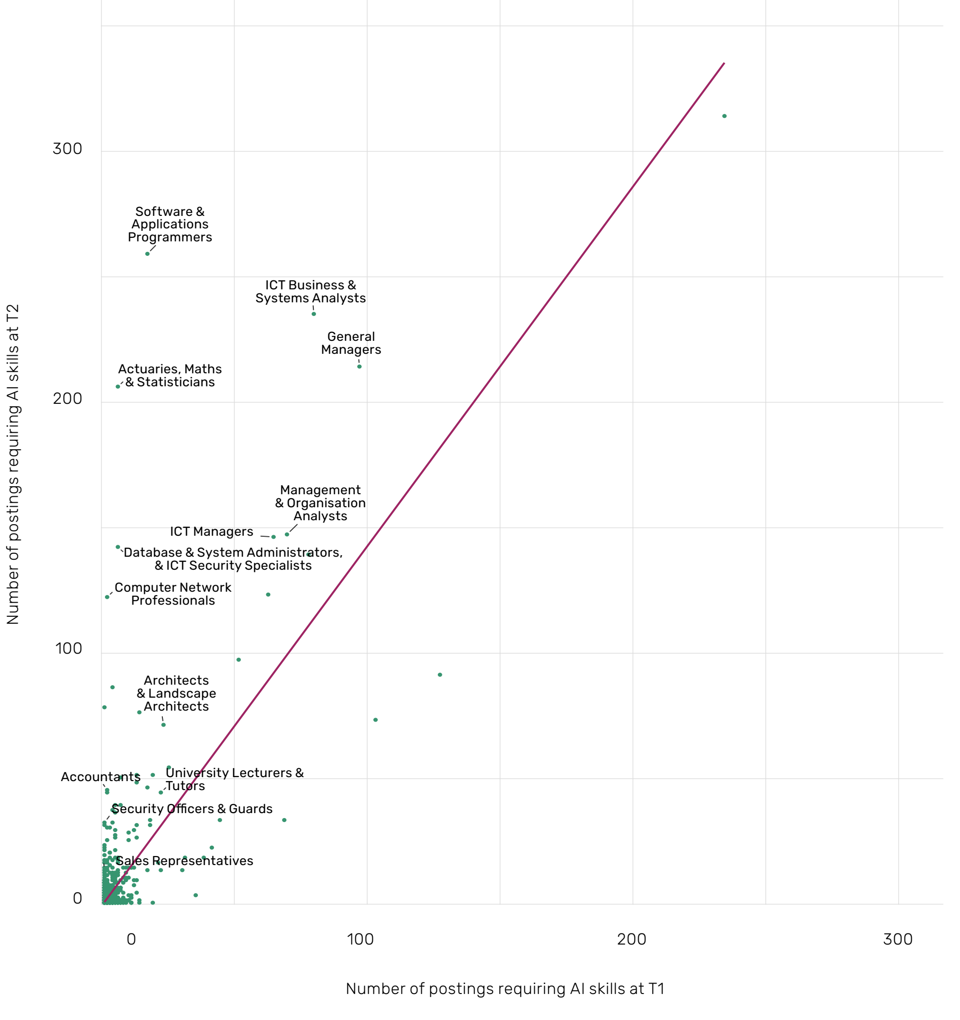

Discussion

The novel combination of methods used in this study provides a cumulative body of evidence countering concerns about AI being used to substitute for workers and/or deskill jobs. Using longitudinal, national data, we find that firms adopting AI (denoted by the firm posting job advertisements that require AI skills) show slightly faster growth in demand for workers and skills, even after controlling for the effects of firm size, geoography and industry. When we investigated these effects at the occupational level, we found that the number of skills sought in job postings was increasing across the board but it was increasing slightly faster in AI adopting firms and occupations that were exposed to AI. Finally, by analysing demand for AI skills at the occupational level, we discovered that some non-AI skilled occupations were transitioning to become AI skilled workers.

More workers or different workers with more skills?

It is important to acknowledge that higher demand for new workers might not lead to workforce growth. Workforce data was only available for 373 large enterprises and in this sample, AI adoption did not explain rates of growth in the workforce. Given that skills requirements were also increasing faster in AI adopting firms, it is possible that AI adopting firms were posting more jobs advertisements because they they were replacing their existing workers with new workers who had the additional skills required in an AI-augmented workplace. Nevertheless, with job postings and skills requirements increasing slightly faster in AI adopting firms and no evidence that AI adoption was associated with a decline in workforce size, our findings provide confidence that AI is not substituting for or deskilling workers.

Although the effects of AI adoption and AI exposure on existing occupations has been studied before (Acemoglu *et al.*, 2022b; Alekseeva *et al.*, 2021; Felten *et al.*, 2019; Green and Lamby, 2023; Webb, 2019), this study is the first to illustrate that the effect of occupational AI exposure on demand for workers depends on whether or not the relevant occupation is employed in an AI adopting firm. Consequently, our study strengthens the grounds for arguing that it is AI adoption (rather than another factor impacting workers who perform tasks that can now be performed by AI) that underlies the observed effects.

The moderation effect (whereby AI exposure was associated with lower rates of growth but only for workers in non-adopting firms) fits with the notion that advanced AI tools allow AI exposed workers augment their productivity or value add in new ways, providing their firms an advantage in the services and products that these workers help to deliver (Saheb and Saheb, 2023). In non-adopting firms, the same AI exposed occupations are then disadvantaged in their ability to compete for market share relative to their AI-augmented peers, with the result that they experience less growth in job postings. We must acknowledge that these effects, while statistically significant, were weak. Nevertheless, they support the theory that advanced AI tools have more of an augmentation than a substitution effect.

AI adopting firms were not just posting more non-AI job ads, they were also posting job advertisements for AI skilled occupations. Furthermore, some of their AI-exposed but formerly non-AI skilled occupations were transitioning to become AI skilled occupations. The latter effect was slightly stronger in occupations that were exposed to AI, providing further evidence that exposed workers are working with, rather than being displaced by, AI. IT occupations were transitioning fastest, a finding that aligns with global research (Alekseeva *et al*., 2021; Borgonovi *et al*., 2023). However, we found several non-IT occupations that were becoming AI skilled. Grinis (2019) argued that the binary classification of occupations into STEM and non-STEM is outdated because STEM skills are now required across a broad range of occupations. The distinction between AI and non-AI workers may also be blurring. Green and Lamby (2023) defined AI workers as those with the skills to develop and maintain AI systems. Our findings reveal that the range of occupations that need to understand and work with AI tools is expanding, encompassing occupations as diverse as security guards, sales representatives and architects.

Limitations

A limitation of this study is the reliance on job postings (and mentions of AI skills in these postings) to determine whether or not a firm is engaging with AI. It may be possible to adopt AI without hiring AI skilled workers if AI development and maintenance is outsourced completely and the use of the AI does not require AI-specific skills. However, the fact that this measure of firm adoption of AI moderated the effect of AI exposure on demand for new workers suggests that this method of differentiating firms that are adopting AI is reasonably effective.

Second, the sample of job postings that could be matched to an employer over-represented some sectors (transport, postal and warehousing, healthcare and social assistance), occupations (machinery operators and drivers and labourers) and locations (Sydney and Melbourne). Since we controlled for the effects of industry, occupation and location in the analyses, this bias in the data should not affect our findings. In addition, firms' geographic location was determined based on where the firm's job postings were located most frequently. In using this approach, we were not able to differentiate between firms that have a national footprint and firms that operating from a single location.

Finally, in adopting Felten *et al*.'s (2018b) AI exposure metrics for the Australian population, we assume that the abilities required in each occupation are the same in the United States and Australia. This assumption is supported by research which found that found that an occupation's exposure to AI varied little, even after taking into account variation in the tasks being performed by workers across different countries (Georgieff and Hyee, 2022). However, our focus on the Australian labour market does limit the generalisability of our findings, since the employment impacts of technology adoption have been found to differ between developed and developing countries (Sharfaei and Bittner, 2024).

Practical implications

Our research suggests that human workers and human skills remain highly sought after in an AI-augmented work environment. Furthermore, the occupations that are most exposed to AI appear to be better off in a firm that adopts AI. The slightly stronger growth in demand for AI exposed workers in AI adopting firms (compared with the same workers in non-adopting firms), along with the significant (albeit small) increase in number of skills required (including AI-related skills) in job postings, align with the argument that more advanced and collaborative AI tools augment (rather than automate) workers (Schleiger *et al*., 2024). The findings validate national policy and investment aimed at educating firms and workers in the effective use of AI (Saheb and Saheb, 2023). Furthermore, our findings reveal key industries (e.g., agriculture, construction, accommodation and food services) and geographic locations (e.g., firms in regional labour markets) that are lagging in terms of AI adoption, offering targets for such efforts.

Directions for further research

Our study focuses on the average effect (across firms) of AI adoption on demand for workers and skills. We note that other researchers have found that the impact of AI adoption in terms of workers' skills depends on the way in which AI is being used (whether to inform a worker's decisions or to direct the worker) and the skill level of the worker (Holm and Lorenz, 2022). Many factors (firm characteristics, type of AI being adopted, type of occupation) could moderate the effects of AI adoption on demand for workers and skills. Elucidating these moderating factors is important, as it may reveal opportunities to strengthen the benefits of AI adoption for both firms and workers.

Finally, we acknowledge that AI adoption is still in its early stages and that workforce impacts may evolve as firms adjust their business processes to the new ways of working that the AI enables (Borgonovi *et al*., 2023). Understanding how ongoing developments in AI systems and business processes are affecting demand for workers and skills requires ongoing research effort.

Conclusion

Although AI tools can now perform some tasks that were traditionally performed by workers in high-skilled and well-paid occupations, our findings counter concerns that they will substitute for workers in AI exposed occupations. We found that AI adopting firms were experiencing slightly stronger growth in job postings than non-adopting firms of the same size, location and industry. In addition, the number of skills required in job postings was increasing slightly faster in AI adopting firms than in non-adopting firms. We conclude that workers who have the skills to use and complement AI remain sought after in an AI-augmented workforce.

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability: Data will be made available on request.
Word count: 9,379

References

Acemoglu, D., Anderson, G., Beede, D., Buffington, C., Chidress, E., Dinlersoz, E., Foster, L., Godschlag, N., Haltiwanger, J., Kroff, Z.,, Zolas, N., 2022a. Automation and the workforce: A firm-level view from the 2019 Annual Business Survey, in: Paper Presented at the NBER/CRIW Conference on Technology, Productivity and Economic Growth. http://pascual.scripts.mit.edu/research/abs/, Washington D.C., pp. 1–45. https://doi.org/10.3386/w30659

Acemoglu, D., Autor, D., 2011. Skills, tasks and technologies: Implications for employment and earnings, in: Handbook of Labor Economics. Elsevier, pp. 1043–1171. https://doi.org/DOI 10.1016/S0169-7218(11)02410-5

Acemoglu, D., Autor, D., Hazell, J., Restrepo, P., 2022b. Artificial intelligence and jobs: evidence from online vacancies. *J Labor Econ* 40, S293–S340. https://doi.org/https://doi.org/10.1086/718327

Albanesi, S., da Silva, A.D., Jimeno, J.F., Lamo, A., Wabitsch, A., 2023. New Technologies and Jobs in Europe. National Bureau of Economic Research. https://doi.org/https://doi.org/10.53479/33414

Alekseeva, L., Azar, J., Gine, M., Samila, S., Taska, B., 2021. The demand for AI skills in the labor market. Labour Econ 102002.

Arntz, M., Gregory, T., Zierahn, U., 2016. The Risk of Automation for Jobs in OECD Countries: A Comparative Analysis (No. 189). Paris. https://doi.org/10.1787/5jlz9h56dvq7-en

Australian Bureau of Statistics, 2021. 2021 Census – employment, income and education [WWW Document]. TableBuilder.

Australian Bureau of Statistics, 2018. Australian Statistical Geography Standard (ASGS) Volume 5 – Remoteness Structure (cat no. 1270.0.55.005) [WWW Document]. URL https://www.abs.gov.au/ausstats/abs@.nsf/mf/1270.0.55.005 (accessed 9.28.22).

Autor, D.H., Levy, F., Murnane, R.J., 2003. The skill content of recent technological change: An empirical exploration. *Q J Econ* 118, 1279–1333.

Babina, T., Fedyk, A., He, A., Hodson, J., 2024. Artificial intelligence, firm growth, and product innovation. *J financ econ* 151, 103745. https://doi.org/https://doi.org/10.1016/j.jfineco.2023.103745

Borgonovi, F., Calvino, F., Criscuolo, C., Nania, J., Nitschke, J., O'Kane, L., Samek, L., Sitz, H., 2023. Emerging trends in AI skill demand across 14 OECD countries. Paris. https://doi.org/https://doi.org/10.1787/7c691b9a-en

Boulus-Rødje, N., Cranefield, J., Doyle, C., Fleron, B., 2024. GenAI and me: the hidden work of building and maintaining an augmentative partnership. *Pers Ubiquitous Comput* 28, 861–874. https://doi.org/10.1007/s00779-024-01810-y

Bratanova, A., Pham, H., Mason, C., Hajkowicz, S., Naughtin, C., Schleiger, E., Sanderson, C., Chen, C., Karimi, S., 2022. Differentiating artificial intelligence activity clusters in Australia. *Technol Soc* 71, 102104. https://doi.org/10.1016/J.TECHSOC.2022.102104

Brynjolfsson, E., Mitchell, T., Rock, D., 2018. What Can Machines Learn, and What Does It Mean for Occupations and the Economy? AEA Papers and Proceedings 108, 43–47. https://doi.org/10.1257/pandp.20181019

Brynjolfsson, E., Rock, D., Syverson, C., 2017. Artificial Intelligence and the Modern Productivity Paradox: a Clash of Expectations and Statistics, NBER Working Paper. University of Chicago Press. https://doi.org/10.3386/w24001

Calvino, F., Criscuolo, C., Ughi, A., 2024. Digital adoption during COVID-19: Cross-country evidence from microdata.

Chen, N., Li, Z., Tang, B., 2022. Can digital skill protect against job displacement risk caused by artificial intelligence? Empirical evidence from 701 detailed occupations. *PLoS One* 17, e0277280. https://doi.org/10.1371/journal.pone.0277280

Cronbach, L.J., Furby, L., 1970. How we should measure "change": Or should we? *Psychol Bull* 74. https://doi.org/10.1037/h0029382

Davis, E., Marcus, G., 2015. Commonsense reasoning and commonsense knowledge in artificial intelligence. Commun ACM 58, 92–103.

Department of Industry Science and Resources, 2024. Exploring AI adoption in Australian businesses [www document].

Duenser, A., Mason, C., 2019. Evaluating online job ads as indicators of demand for new workers: Characterising strengths and weaknesses. Canberra. https://doi.org/10.25919/5e46e02bb21c1

Edwards, J., 1994. Regression analysis as an alternative to difference scores. *J Manage* 20. https://doi.org/10.1016/0149-2063(94)90011-6

Eloundou, T., Manning, S., Mishkin, P., Rock, D., 2023. GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models. https://doi.org/https://doi.org/10.48550/arXiv.2303.10130

European Commission, 2019. ESCO handbook: European skills competences, qualifications and occupations. Brussels, Belgium.

Evans, D., Mason, C., Chen, H., Reeson, A., 2024. Accelerated demand for interpersonal skills in the Australian post-pandemic labour market. *Nat Hum Behav* 8. https://doi.org/10.1038/s41562-023-01788-2

Evans, D., Mason, C., Chen, H., Reeson, A., 2023a. An algorithm for predicting job vacancies using online job postings in Australia. *Humanit Soc Sci Commun* 10, 1–9.

Evans, D., Zhao, Y., Mason, C., Chen, H., Reeson, A., Burns, S., 2023b. An evaluation of Adzuna Australia job postings as a measure of labour demand. Brisbane, Australia. https://doi.org/https://doi.org/10.25919/t2zf-8e10

Eysenck, M.W., Eysenck, C., 2021. AI vs Humans. Routledge.

Felten, E., Raj, M., Seamans, R., 2021. Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal* 42, 2195–2217.

Felten, E.W., Raj, M., Seamans, R., 2019. The effect of artificial intelligence on human labor: An ability-based approach, in: AOM 2019: Understanding the Inclusive Organization - 79th Annual Meeting of the Academy of Management. https://doi.org/10.5465/AMBPP.2019.140

Felten, E.W., Raj, M., Seamans, R., 2018a. A Method to Link Advances in Artificial Intelligence to Occupational Abilities. AEA Papers and Proceedings 108, 54–57. https://doi.org/10.1257/pandp.20181021

Felten, E.W., Raj, M., Seamans, R., 2018b. A Method to Link Advances in Artificial Intelligence to Occupational Abilities. AEA Papers and Proceedings 108, 54–57. https://doi.org/10.1257/pandp.20181021

Fossen, F., Sorgner, A., 2019. Mapping the future of occupations: Transformative and destructive effects of new digital technologies on jobs. *Foresight and STI Governance* 13, 10–18. https://doi.org/10.17323/2500-2597.2019.2.10.18

Georgieff, A., Hyee, R., 2022. Artificial intelligence and employment: New cross-country evidence. *Front Artif Intell* 5, 832736. https://doi.org/https://doi.org/10.3389/frai.2022.832736

Green, A., Lamby, L., 2023. The Supply, Demand, and Characteristics of the AI Workforce across OECD countries. Geneva.

Grimberg, E., Mason, C., 2025. Building Proficiency in GAI: Key Competencies for Success. Qeios. https://doi.org/10.32388/INFMMJ

Grinis, I., 2019. The STEM requirements of “Non-STEM” jobs: Evidence from UK online vacancy postings. *Econ Educ Rev* 70, 144–158. https://doi.org/10.1016/j.econedurev.2019.02.005

Holm, J.R., Lorenz, E., 2022. The impact of artificial intelligence on skills at work in Denmark. *New Technol Work Employ* 37, 79–101.

IBISWorld, 2025. AU and NZ Enterprise Profiles [WWW Document]. URL https://help.ibisworld.com/en/articles/8136631-au-nz-enterprise-profiles (accessed 6.25.25).

Loaiza, I., Rigobon, R., 2024. The EPOCH of AI: Human-Machine Complementarities at Work. Available at SSRN 5028371.

Mason, C.M., Chen, H., Evans, D., Walker, G., 2023. Illustrating the application of a skills taxonomy, machine learning and online data to inform career and training decisions. *International Journal of Information and Learning Technology* 40. https://doi.org/10.1108/IJILT-05-2022-0106

National Center for O*NET Development, 2022. O*NET Questionnaires [WWW Document]. O*NET Resource Center. URL www.onetcenter.org/questionnaires.html (accessed 5.5.22).

Nguyen, K., Hambur, J., 2023. Adoption of Emerging Digital General-purpose Technologies: Determinants and Effects. Reserve Bank of Australia.

Priddis, M., Miller, G., George, R., 2020. Technology impacts on the Australian workforce. ACS, Sydney.

Sadiku, M., Ashaolu, T.J., Ajayi-Majebi, A., Musa, S., 2021. Augmented Intelligence. *International Journal Of Scientific Advances* 2. https://doi.org/10.51542/ijscia.v2i5.17

Saheb, Tahereh, Saheb, Tayebeh, 2023. Topical review of artificial intelligence national policies: A mixed method analysis. *Technol Soc* 74, 102316. https://doi.org/https://doi.org/10.1016/j.techsoc.2023.102316

Schleiger, E., Mason, C., Naughtin, C., Reeson, A., Paris, C., 2024. Collaborative Intelligence: A scoping review of current applications. *Applied Artificial Intelligence* 38, 2327890.

Sharfaei, S., Bittner, J., 2024. Technological employment: Evidence from worldwide robot adoption. *Technol Forecast Soc Change* 209, 123742.

Shiohira, K., 2021. Understanding the Impact of Artificial Intelligence on Skills Development. Education 2030. UNESCO-UNEVOC International Centre for Technical and Vocational Education and Training.

Tolan, S., Pesole, A., Martínez-Plumed, F., Fernández-Macías, E., Hernández-Orallo, J., Gómez, E., 2021a. Measuring the occupational impact of AI: tasks, cognitive abilities and AI benchmarks. *Journal of Artificial Intelligence Research* 71, 191–236. https://doi.org/https://doi.org/10.1613/jair.1.12647

Tolan, S., Pesole, A., Martínez-Plumed, F., Fernández-Macías, E., Hernández-Orallo, J., Gómez, E., 2021b. 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. https://doi.org/http://dx.doi.org/10.2139/ssrn.3482150

White House, 2022. The Impact of Artificial Intelligence on the Future of Workforces in the European Union and the United States of America. Washington D.C.

Yau, K.-L., Lee, H., Chong, Y.-W., Ling, M., Syed, A., Wu, C., Goh, H., 2021. Augmented Intelligence: Surveys of Literature and Expert Opinion to Understand Relations Between Human Intelligence and Artificial Intelligence. IEEE Access PP, 1. https://doi.org/10.1109/ACCESS.2021.3115494

Zhang, Q., Zhang, F., Mai, Q., 2023. Robot adoption and labor demand: A new interpretation from external competition. *Technol Soc* 74, 102310. https://doi.org/DOI: 10.1016/j.techsoc.2023.102310

Zhao, Y., Chen, H., Mason, C.M., 2021. A Framework for Duplicate Detection from Online Job Postings, in: IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology. pp. 249–256. https://doi.org/10.1145/3486622.3493928

Zolas, N., Kroff, Z., Brynjolfsson, E., McElheran, K., Beede, D., Buffington, C., Goldschlag, N., Foster, L., Dinlersoz, E., 2020. Advanced Technologies Adoption and Use by U.S. Firms: Evidence from the Annual Business Survey. Cambridge, MA. https://doi.org/10.3386/w28290

The effect of mental health on early retirement decisions: Evidence from Australia

AHARON KATZ *University of York, United Kingdom*

Abstract

Health and labour supply are interconnected; However, research has predominantly focused on the impact of physical health, leaving a gap in understanding the role of mental health problems. This study addresses this gap by examining the effect of mental health on early retirement decisions using data from the Household, Income, and Labour Dynamics in Australia (HILDA) Survey. We use both linear probability models and a discrete-time hazard approach. While linear models estimate the average effect, the discrete-time hazard model tracks initially employed individuals aged 50 to 64 over time until they retire early or reach retirement age. To mitigate potential bias arising from the timing of reporting of mental health and retirement decisions, lagged measures of mental health are applied, with respect to the temporal sequence of events. To address measurement bias, the association between our derived mental health variable and other objective psychiatric measures is examined. Furthermore, we include the death of a close friend as an instrument for mental health status, helping us validate and strengthen causal findings of our study. Lastly, we examine whether unobserved heterogeneity poses a problem in our analysis by estimating models with and without unobserved heterogeneity. Our findings indicate a significant and positive causal impact of poor mental health on early retirement decisions, which is also supported by the nonlinear analysis. To explore potential gender heterogeneity, separate analyses are conducted for males and females. The observed differences in the results between the two groups support the assumption of gender-specific effects. These findings suggest that poor mental health has a significant and potentially causal impact on premature exit from the labour market, particularly among men. The results highlight the importance of effective mental health management in supporting longer working lives.

JEL Codes: I12, I15, J14, J26

Contact: aharon.katz@york.ac.uk

Introduction

Mental health issues represent a growing concern for labour markets worldwide, significantly influencing employment outcomes across sectors. Poor mental health affects labour market participation through several channels. Affected individuals may experience increased absenteeism due to illness, medical appointments, or treatment regimens (Bloom and Canning 2000). Symptoms, such as depression or mood instability, can impede regular work attendance and reduce motivation to invest in human capital (Fadare *et al*. 2023, Tompa 2002). The resulting decline in labour productivity and economic engagement may encourage earlier retirement, particularly if individuals anticipate a reduced lifespan and wish to maximise leisure time from accumulated wealth. Additionally, presenteeism, or diminished productivity while at work, may further erode performance, ultimately leading to premature labour market exit (Chatterji *et al*. 2007). Mental health problems may also increase reliance on non-wage income sources, such as welfare benefits, thereby potentially reducing the financial incentive to remain employed (Disney *et al*. 2006). However, this relationship is not uniform; some individuals may instead choose to remain in employment longer in order to meet elevated healthcare expenses associated with managing chronic mental health conditions (Bryan *et al*. 2022, Frijters *et al*. 2010, Ngui *et al*. 2010, Hamilton *et al*. 1997).

In Australia, the economic burden of poor mental health is substantial. Estimates by the Australian Bureau of Statistics suggest that mental health-related work absences cost approximately AU$60 billion annually (Australian Bureau of Statistics 2008). Furthermore, Lee *et al*. (2017) estimated that productivity losses stemming from depression, anxiety disorders, and substance use disorders amounted to AU$11.8 billion in 2007, with additional fiscal implications of AU$1.2 billion in lost income tax revenue and AU$12.9 billion in welfare expenditures. For individuals diagnosed with psychosis, Neil *et al*. (2014) calculated productivity losses at AU$40,941 per person, with further indirect costs, such as those associated with supported employment and non-governmental services, totalling an additional AU$14,642 per person.

This paper makes three key contributions to the literature on mental health and labour supply. First, it tries to provide causal evidence on the effect of mental health on early retirement decisions among older individuals in Australia, addressing an underexplored area in comparison to the well-established literature on physical health and labour market exits. Mental health conditions, often less visible than physical ailments, may receive disparate treatment from employers and policymakers (Ngui *et al*. 2010). For example, in a similiar economy, the United Kingdom, one in four people reported a mental disorder in 2016 (Alderwick and Dixon 2019, McManus *et al*. 2016), yet the National Health Service allocated only 12.5 per cent of its total budget to mental health in 2017/18, with a target of 16.2 per cent by 2022/23 (Baker & Kirk-Wade 2023).

As Australia's population ages, with those over 50 projected to increase from 12 per cent to 22 per cent between 2015 and 2050 (World Health Organization 2022), understanding the causes of early retirement becomes critical for economic planning.

An ageing workforce can strain public finances and healthcare systems, increase dependency ratios, and necessitate adjustments to retirement policies (Liddiard 1978). Determining whether early retirement is driven by health or financial considerations has significant implications for public policy (Frijters *et al*. 2 010). Distinguishing mental from physical health conditions helps in targeting interventions and tailoring incentive schemes that can prolong workforce participation.

To investigate the causal relationship between mental health and early retirement, the empirical analysis proceeds in two stages. First, a linear cross-sectional framework is applied, beginning with Ordinary Least Squares (OLS) estimation and followed by a Two-Stage Least Squares (2SLS) specification. The 2 SLS model uses the lagged death o f a close friend as an instrument for mental health. Unlike the more commonly used death of a family member, which may directly affect retirement through caregiving or financial implications, this instrument is assumed to influence mental health without directly impacting retirement decisions, helping to address endogeneity in a transparent and interpretable way.

To establish a benchmark and enable meaningful comparisons, we begin by estimating a linear probability model. This serves as a baseline against which the results of subsequent nonlinear analyses can be assessed. Although linear models are intuitive and straightforward, they impose restrictive assumptions and may yield biased estimates. Nonlinear analysis offers an additional framework for modelling time-to-event outcomes like retirement, as it accounts for the dynamic nature of such decisions and handles censoring more appropriately.

The second contribution to the literature, and the second stage of the analysis, extends the investigation using discrete-time hazard models to examine the timing of retirement while explicitly addressing whether unobserved heterogeneity is an issue in our analysis. This approach compares models with and without unobserved heterogeneity, under various distributional assumptions, to assess the robustness of the results. Specifically, a discrete-time hazard model is paired with a two-stage residual inclusion (2SRI) estimator to control for both endogeneity and unobserved confounding (Terza *et al*. 2008). This more flexible methodology allows for testing whether linear models provide adequate estimates or if failing to account for unobserved heterogeneity and nonlinearity introduces bias, especially when unmeasured factors such as personality traits or job satisfaction may simultaneously influence mental health and retirement decisions.

By leveraging panel data within the discrete-time hazard framework, the longitudinal structure captures individual dynamics over time. While the discrete-time model assumes selection on observables, i.e., that mental health only affects retirement through measured variables, this assumption may be violated if unobserved factors influence both outcomes. To address this, the 2SRI procedure improves causal inference by mitigating bias from such confounding in a panel data setting.

Lastly, the paper explores the role of measurement error and reporting bias in mental health assessment. Unlike many existing studies that rely exclusively on self-reported mental health status, this study compares subjective and objective indicators to account for potential biases in reporting. Social desirability or stigma may lead respondents to underreport psychological conditions, which can result in attenuated estimates of the true effects (Bharadwaj *et al*. 2017, Brohan and Thornicroft 2010, Rüsch *et al*. 2005). By triangulating different data sources and validating mental health measures, the analysis seeks to provide more robust and credible findings.

The results reveal a statistically significant association between poor mental health and an increased likelihood of early retirement. The instrumental variable analysis supports a causal interpretation, with the death of a close friend serving as a valid and relevant instrument, as indicated by the first-stage statistics. When the analysis is extended to a longitudinal framework, the findings remain consistent, with the estimated effect of mental health on early retirement slightly larger than in the non-IV models. This pattern suggests that failing to address endogeneity may lead to an underestimation of the true impact of mental health.

The importance of correcting for endogeneity is further highlighted by the notably larger coefficient, more than six times greater, in the 2SLS model compared to the baseline OLS estimate. Moreover, accounting for unobserved heterogeneity in the longitudinal models does not substantially alter the results, with estimates remaining similar across models with and without frailty.

Importantly, the results indicate gender-specific differences in the mental health and retirement relationship. For males, poor mental health significantly increases the likelihood of early retirement across most specifications, whereas for females, the effect is weaker and not statistically significant in some models. These findings offer important insights for policymakers aiming to extend working lives, improve mental health support, and mitigate the fiscal challenges posed by an ageing population. This study contributes to the existing literature by providing robust evidence from Australia, applying rigorous econometric methods, and underscoring the significance of gender and measurement issues in mental health research.

The outline of the paper is as follows. The next section (section 2) presents a review of the existing literature that examines the impact of mental health on labour market outcomes and early retirement decisions. Section 3 describes the empirical methods, including linear probability model, two-stage least squares, discrete-time hazard models with and without unobserved heterogeneity, validity checks, and the inclusion of the two-stage residual inclusion approach within a discrete-time hazard model. This is followed by the data description in section 4. Section 5 presents the results, while section 6 contains the discussion and conclusion.

Background and Literature Review

Mental health disorders are diverse and vary in their symptoms, treatments, diagnosis, and outcomes (American Psychiatric Association 2013). The severity of mental illness varies greatly from common illnesses, such as general anxiety, mood swings, eating disorders, to more severe mental disorder e.g., schizophrenia. There is a wide range of literature within different fields of research that show mental health disabilities are linked to poorer labour market outcomes, in addition to affecting social life, such as terminating relationships, loneliness, and a greater likelihood of being involved in a crime (Bartel and Taubman 1979, 1986).

There is a substantial amount of previous research on the relationship between health and labour market outcomes, such as economic inactivity or retirement. However, most of the re-search focuses on the effect of physical or general health, without controlling for mental health specifically. Few papers which analyse the effect of mental health and labour market outcomes, such as employment status (Bryan *et al*. 2022, Chatterji *et al*. 2011, Frijters *et al*. 2010, Lu *et al*. 2009, Sainsbury *et al*. 2008, Alexandre and French 2001, Ettner *et al*. 1997), income (Chatterji *et al*. 2011, Ettner *et al*. 1997) and work hours (Ettner *et al*. 1997). Stern (1989) observed the effect of health on retirement decisions, while using longitudinal data and observing health using subjective measures.

The empirical analysis is grounded in the standard life-cycle labour supply framework, in which individuals allocate time between work, leisure, and health investment to maximise life-time utility (Grossman 1972). Within this framework, health, including mental health, affects both the utility derived from work and the productivity of labour. To empirically assess this relationship, we estimate baseline linear models (LPM and 2SLS) where employment status serves as a reduced-form representation of labour supply decisions, and mental health is treated as a potentially endogenous determinant of labour market participation.

Stern (1989) was one of the first to identify a negative and statistically significant effect of health on retirement and labour supply, suggesting that poor health leads to a reduced labour market participation. Similarly, Bryan *et al*. (2022) found that poor mental health decreases the probability of being employed, while Lu *et al*. (2009) reported that a decline in average mental health is associated with a significant reduction in both employment rates and annual income. Their findings also indicate that the mental health index has a positive and significant effect on the likelihood of being employed. Despite these contributions, recent studies have focused mainly on general labour market outcomes, with limited attention to older individuals and specific outcomes such as early retirement decisions, exceptions include the work of Zucchelli *et al*. (2007) and Disney *et al*. (2006), who specifically examine the effect of general health on labour market outcomes across older workers.

Previous studies of the effect of mental health on labour market outcomes, based mainly on cross-sectional data to examine this relationship (Zhang *et al*. 2009, Cai and Kalb 2007, Bound *et al*. 1999, Siddiqui and Ali Shah 1997, Bazzoli 1985). However, the use of cross-sectional data poses limitations in capturing the dynamic effects of health on labour supply and addressing the issue of possible endogeneity through unobserved personal characteristics. To overcome these limitations, longitudinal data has been recommended to mitigate selection bias and gain a deeper understanding of individual behaviour (Chatterji *et al*. 2011, Nerlove 2005).

The discrete-time hazard analysis, widely used in the biomedical field and more recently in health economics research, provides valuable insights into the impact of health on workforce outcomes (Bu¨nnings and Tauchmann 2015, Zucchelli *et al*. 2007, Disney *et al*. 2006). This approach offers several advantages. Firstly, it allows greater flexibility in modelling dynamics and exploring variations in the impact of health on labour market transitions according to current employment status (Disney *et al*. 2006). Secondly, it enables examination of the timing and sequencing of events, offering a more nuanced understanding of the process leading to early retirement. By modelling the transition probabilities over discrete-time intervals, we can capture the dynamic nature of the decision making process and account for changes in mental health status over time. Finally, the discrete-time hazard approach accommodates time-varying covariates, such as changes in mental health status, marital status, and other relevant factors, allowing for a more comprehensive analysis of the determinants of early retirement decisions.

Building on this methodological foundation, previous empirical studies using longitudinal data have shown that poor self-reported general health is strongly associated with early retirement among older adults (Zucchelli *et al*. 2007, Disney *et al*. 2006). For instance, Disney *et al*. (2006) and Zucchelli *et al*. (2007) apply discrete-time hazard models using the HILDA dataset, adopting a non-parametric approach to the hazard function. Our study extends this line of research by employing the same discrete-time hazard framework described by Jenkins (1995), but focusing specifically on mental health rather than general health. Furthermore, consistent with Disney *et al*. (2006), Jenkins (1997, 1995), we account for unobserved heterogeneity to ensure robust inference.

While the discrete-time hazard framework provides a powerful tool to analyse transitions into early retirement, it also brings to the forefront an important econometric concern, namely endogeneity. In studying the relationship between mental health and retirement, endogeneity may arise due to measurement error in self-reported health, the potential for reverse causality between mental health and labour market participation, or omitted variables correlated with both. Addressing this issue is crucial to ensure that estimated effects reflect a causal rather than merely associative relationship.

The issue of endogeneity in our study is particularly relevant due to the potential presence of unobserved characteristics and events associated with both mental health and early retirement decisions (Bryan *et al*. 2022). To mitigate potential endogeneity bias, our study draws on strategies that exploit exogenous variation in health. One common approach involves identifying health shocks that are plausibly independent of unobserved individual characteristics (Disney *et al*. 2006, Bound *et al*. 1999). In this study, we define a health shock by incorporating both lagged and initial-period health. By conditioning on initial health, the coefficient on lagged health can be interpreted as a deviation from an individual's underlying health stock (Disney *et al*. 2006), thereby helping to control for unobserved, time-invariant health-related heterogeneity (Jones 2009).

We also control for the latter to minimise omitted variable bias, as both mental and physical health can influence retirement decisions. Specifically, we construct a variable representing a negative physical health shock, defined as the difference between expected and actual self-reported physical functioning scores. Following Apouey *et al*. (2019), we create a binary indicator that takes the value of one when the observed decline in health is greater than one standard deviation relative to the individual's expected change in health status. This measure helps isolate mental health effects that are not merely reflections of concurrent changes in physical condition.

While exogenous shocks can reduce endogeneity, they may not fully eliminate it. To strengthen causal identification, instrumental variable (IV) techniques are widely employed in health and labour economics. The work by Angrist *et al*. (1996) establishes a comprehensive framework for identifying causal effects using IVs, outlining the necessary assumptions, such as relevance and the exclusion restriction, and the estimation techniques required for valid inference. This framework has since guided a broad range of empirical applications in health and labour economics, highlighting both the potential and the limitations of IV strategies, including issues of weak instruments, overidentification, and bias arising from unobserved confounding variables (Frijters *et al*. 2010, Wooldridge 2010, Angrist and Pischke 2009, Zhang *et al*. 2009, Terza *et al*. 2008, Alexandre and French 2001, Ettner *et al*. 1997, Hamilton *et al*. 1997).

Building on these methodological insights, our study considers instruments proposed in the literature for mental health. One commonly used strategy relies on the death of a family member, under the assumption that such an event affects mental health but is not directly related to labour market outcomes (B¨ockerman *et al*. 2022, Burrell *et al*. 2022). However, applying this strategy in the context of employment decisions introduces additional complexity. The exclusion restriction may be violated if the death of a family member influences labour supply through alternative pathways, such as inheritance or changes in household income. For instance, the death of a spouse may compel the individual to continue working for financial reasons, thereby confounding the relationship between mental health and employment.

To address these challenges, we adopt a two-step strategy that strengthens identification. First, we refine our measurement of mental health by regressing alternative indicators of psychological wellbeing on the main mental health variable as a robustness check, helping assess potential measurement bias. Second, we apply an IV approach that uses the death of a close friend as an instrument for mental health, a shock that is plausibly exogenous to labour market outcomes. The discrete-time hazard model is then estimated under alternative distributional assumptions, both with and without unobserved heterogeneity, and incorporates a two-stage residual inclusion (2SRI) procedure to account for endogeneity within the nonlinear framework.(Bound *et al*. 2001, Butler *et al*. 1987, Kessler *et al*. 2002, Zucchelli *et al*. 2007, Disney *et al*. 2006).

The 2SRI method provides an alternative to the conventional IV approach when applied to nonlinear models. In the context of discrete-time hazard analysis, the 2SRI approach has been applied to estimate the effect of a time-varying exposure or treatment on a time-to-event outcome (Garrido *et al*. 2012, Basu and Rathouz 2005). Previous studies have used the 2SRI approach to estimate the effect of various health-related interventions, such as diabetes treatments, drug treatments, and coverage of health services, on hazard outcomes (Ying *et al*. 2019, Mery *et al*. 2016, Tchetgen *et al*. 2015). The 2SRI approach has the advantage of accounting for the endogeneity of the treatment variable, which can improve the accuracy of the estimated treatment effect, but it may still remain biased.

The use of IVs in discrete-time hazard analysis can be challenging due to the need for large sample sizes and the potential for weak instruments. Weak instruments can result in biased estimates and large standard errors, leading to incorrect inference (Terza *et al*. 2008, Angrist *et al*. 1996). In addition, the use of IVs assumes that the exclusion restriction holds, meaning that the IV only affects the outcome through its effect on the endogenous variable, and not through any other pathway. This assumption can be difficult to verify in practice. Studies have highlighted the effectiveness of the IV approach in addressing endogeneity and providing insights into the effects of variables on survival outcomes (Tchetgen *et al*. 2015, Terza *et al*. 2008, Ettner *et al*. 1997).

Our study builds on the econometric framework established by Terza *et al*. (2008) and Terza (2018), which formalises the use of the 2SRI estimator in nonlinear models. In this context, a valid instrument must satisfy two key criteria: (1) it must be strongly correlated with the endogenous explanatory variable (relevance), and (2) it must be conditionally independent of the outcome, given the endogenous regressor and other covariates (exogeneity). Although, these conditions are similar to those in 2SLS, 2SRI differs in how endogeneity is addressed in nonlinear models: by including the first-stage residual as an additional regressor in the second-stage model, rather than using predicted values as in 2SLS.

Literature on the effect of mental health on labour market outcomes in Australia is infrequent and often limited to a sub-sample analysis of males. Cai and Kalb (2007) use HILDA survey to examine the effect of health and labour participation. Using a simultaneous equation model for working-age participants to control for potential endogeneity of health, their findings show that health has a positive effect on retirement. Wilkins (2004) and Brazenor (2002) use cross-sectional data to analyse the effect of 10 different disabilities on labour market outcomes of the older population, such as income level and employment status. Wilkins (2004) shows that, on average, disability reduces the likelihood of labour market participation with different effects for males and females. However, Brazenor (2002) finds that the impact varies depending on the disability type.

Bubonya *et al*. (2019) focus on depressive symptoms as a proxy for having a mental health disorder and its effect on employment status, using the first 14 waves in HILDA. To measure depressive symptoms, they use the same mental health variable available in HILDA as in our paper. As their primary focus is on anxiety and mood disorders, they assign different weights to the five items used to construct the derived mental health variable. Based on this approach, scores below 60 are coded as moderate symptoms, while scores below 52 are considered indicative of severe symptoms.

Methodology

Existing empirical research has adopted two main strategies to address the structural endogeneity inherent in analysing the relationship between health and labour market outcomes. One common approach involves identifying the model through theoretically motivated exclusion restrictions. This requires finding instruments that affect health but not employment. However, the validity of such instruments is often difficult to defend in practice. An alternative strategy relies on exploiting the temporal ordering of events to mitigate reverse causality, typically by modelling the effects of lagged health and employment outcomes (e.g., Steele *et al*. 2013, Olesen *et al*. 2013).

In line with this, recent work by Bubonya *et al*. (2019) using Australian panel data (HILDA) examines the bidirectional relationship between depressive symptoms and employment. They apply linear fixed-effects and dynamic panel models to control for unobserved heterogeneity and to explore feedback mechanisms over time. Their analysis confirms a robust negative effect of poor mental health on employment probabilities, while also documenting that job loss exacerbates psychological distress, reinforcing concerns about simultaneity bias. These insights inform our empirical strategy. We begin by estimating a linear probability model to establish a baseline association between mental health and early retirement. We then proceed to an instrumental variable framework to address potential endogeneity arising from omitted variables and reverse causality. Finally, to capture the timing of retirement decisions in relation to health, we include a discrete-time hazard model.

Baseline Linear Model: LPM and 2SLS

To gain an initial understanding of the relationship between mental health and early retirement decisions, we first estimate a linear probability model (LPM), in which an outcome variable of reported retirement is regressed on lagged mental health and a set of control variables using a pooled regression framework for individuals aged 50–64. This baseline model offers a straight-forward interpretation of the effects. However, this approach does not account for the timing of retirement events, the longitudinal structure of the data, or the possibility of censoring over time. To more appropriately address these dynamic aspects, the analysis is extended using discrete-time hazard models, which explicitly model the retirement decision as a time-to-event process.

The baseline LPM model captures the relationship between mental health and the probability of retirement. The specification is given by:

$$Y_{it} = \alpha_0 + \alpha_1 m_{it-1} + \alpha_2 X_{it} + \epsilon_{it} \tag{1}$$

where Y_{it} is a binary indicator equal to 1 if individual i reports retirement at time t and 0 otherwise, m_{it-1} denotes the lagged mental health score with lower values indicating poorer mental health and higher values reflecting better mental health, and X_{it} is a vector of socio-demographic controls. The error term ϵ_{it} captures unobserved factors affecting retirement. Time subscripts are included; however, we estimate the model by pooling observations across individuals and time, effectively treating the data as if it consists of $i \times t$ independent units.

Although LPM offers a straightforward interpretation, the coefficient α_1 measures the change in the probability of retirement associated with a one-unit change in lagged mental health score, it has several limitations. Notably, the LPM can produce predicted probabilities outside the [0, 1] interval and assumes linearity in probabilities, which might be restrictive given the binary outcome. More importantly, although OLS provides a straightforward estimate, it is likely to suffer from endogeneity due to reverse causality or omitted variable bias, where poor mental health causes and is caused by retirement decisions. For example, retirement could itself impact mental health, creating reverse causality bias.

To mitigate endogeneity concerns, we implement a Two-Stage Least Squares (2SLS) instrumental variables strategy. We use the death of a close friend in the previous period to mental health (t-2) as an instrument for lagged mental health (t-1). This instrument is relevant because bereavement is empirically linked to worsened mental health outcomes (Frijters *et al*. 2010). Under the key assumption that the death of a close friend affects retirement decisions only through its effect on mental health (exclusion restriction), the 2SLS estimates provide a more reliable causal effect of mental health on retirement probability. The first-stage predicts mental health using the instrument, and the second-stage regresses retirement on the predicted mental health and covariates, thus purging the mental health variable of endogeneity bias.

The first-stage of the 2SLS model is specified as:

$$m_{it} = \pi_0 + \pi_1 Z_{it-1} + \pi_2 X_{it} + \eta_{it} \tag{2}$$

where Z_{it-1} is an indicator for the death of a close friend reported at time t − 1, referring to an event that occurred between $t - 2$ and $t - 1$. This variable therefore instruments $m_{i,t-1}$ for the outcome observed at time t, capturing the delayed effect of bereavement on mental health.

The additional lag allows for the time it can take for grief to translate into measurable mental health deterioration.

The second-stage uses the lagged predicted mental health score from the first stage, $\hat{m}_{it-1}$, to estimate its causal impact on early retirement:

$$Y_{it} = \beta_0 + \beta_1 \hat{m}_{it-1} + \beta_2 X_{it} + \nu_{it} \tag{3}$$

This IV strategy allows us to recover a local average treatment effect (LATE), interpreted as the causal effect of mental health on retirement among individuals whose mental health was affected by the bereavement shock. A statistically significant and negative $\beta 1$ would indicate that better mental health causally decreases the likelihood of retirement.

To assess instrument strength, we report the first-stage F-statistic. A value above the conventional threshold of 10 suggests a strong instrument. Additionally, validity checks are conducted using alternative specifications of bereavement timing and sub-group analyses by gender.

To account for unobserved individual heterogeneity that may bias the estimates in the baseline linear model, we also estimate a fixed-effects panel regression. This approach controls for time-invariant individual characteristics (e.g., personality traits, early-life conditions) that may jointly influence both mental health and retirement decisions. The fixed-effects specification is given by:

$$Y_{it} = \alpha_1 m_{it-1} + \alpha_2 X_{it} + \mu_i + \varepsilon_{it} \tag{4}$$

where μ_i captures unobserved individual-specific effects, and ε_i denotes the idiosyncratic error term after accounting for μ_i. By differencing out time-invariant unobservables, the fixed-effects estimator helps to address bias from omitted variables that are constant over time. However, it does not account for time-varying endogeneity, which is addressed in instrumental variable strategies.

We include the linear approach alongside the discrete-time hazard analysis to establish a baseline and enable meaningful comparison. Although linear probability models offer a straight-forward starting point for examining the relationship between mental health and early retirement, they come with notable limitations. Chief among these are the assumptions of constant marginal effects and the risk that predicted probabilities may fall outside the valid range of 0 to 1. These constraints become particularly problematic when modelling binary or time-to-event outcomes such as retirement.

Furthermore, linear models are limited in their ability to analyse retirement behaviour, as they do not track individuals over time and therefore fail to fully exploit the longitudinal structure of the HILDA dataset. These models also struggle to capture the timing of retirement decisions and the presence of time-related features such as censoring. To address these limitations, we include a discrete-time hazard modelling framework that explicitly follows individuals over time and accounts for the dynamic nature of retirement behaviour. This nonlinear approach provides a more flexible and accurate representation of the timing of retirement, while appropriately handling duration dependence and potential right-censoring in the data.

Discrete-time Hazard Model

We apply the discrete-time hazard model and assume that individuals become at risk of early retirement at the age of 50 years old, which is in line with the previous literature (Butterworth *et al.* 2006, Melzer *et al.* 2004). Setting the pattern of duration dependence is essential prior to analysing the model, and thus we include an additional 12 age dummies to capture the duration dependence by the age of the individuals, starting from

53 to 64 years old. In our sample of HILDA data, which is based on the number of retirees within each age group, the actual risk of retiring early started when participants are 53 years and older, since the number of reported retirements between the ages 50 and 52 is minimal in our sample.

Our research focuses on understanding how mental health influences the decision to retire. In order to observe individuals who are at risk of retirement, we define a stock sample based on the definition of Jenkins (1995). Over the course of the study, only 45 individuals were present for all 15 waves due to sample attrition, retirement, and death. Our retirement models are estimated on complete sequences of observations, using information up to the wave of first exit if an individual leaves the panel and returns in later waves.

We create additional rows of observation per individual based on the time that this individual was at risk of reporting early retirement in our data. This requires that in general, each respondent, i, contributes T_i rows, where T is the number of time periods i was observed at risk of failure. Jenkins (1995) method enables all periods prior to selection to be ignored. This approach relies on only focusing on the relevant time restricted data for observing the effect of interest. Additionally, the data needs to be rearranged and conditioning on stock sampling prior estimation and need to have an unbalanced data format.

We follow the same notation as Jenkins (1995), so time, t, equals to the initial period, τ , ($t = \tau$) when the individual enters the analysis. Each respondent i contributes s_i years of risks while they are still in employment in the interval between the initial period τ and si, so the time when retirement occurs is ($t = \tau + s_i$). Here, we only rely on the age of an individual reported in the first wave in the survey for evaluating the retirement age. Additionally, we follow the age variables on wave 2 onward. The age and our derived time variable may differ as the timing of the surveys varies between the 18 waves and can cause respondents to report the same age in two consecutive waves, if the timing of the later wave was before their birthday. We assume throughout those individuals aged one year between waves of data.

At the end of the sample, each individual will either have retired early (where δ_i = 1) or will be censored and thus will still be working during the last wave in our sample, (δ_i = 0). The age of the individual $t_i = \tau_i + s_i$ will be the retirement age in a complete duration data or the final age of observation if δ_i = 0. As we ignore all the periods prior to selection into the stock sample (Jenkins 1995), we initially selected only those who are participating in the labour market at their first wave, and therefore drop all individuals who reported their occupational status as retired in the wave when joining the survey.

The dependent variable is a binary indicator for whether the participant remained in or left the labour market through retirement in that wave. We do not allow individuals to report multiple retirements by coming back to the labour market and then leaving again. That is retirement is considered as an absorbing (permanent) state, where individuals cannot return to the labour market.

Consider the discrete-time hazard rate, for individual i, with the conditioned probability of retirement at age t, given by:

$$h_{it} = P[T_i = t|T_i > t; X_{it}, m_{it-1}] \tag{5}$$

Where T_i is a discrete random variable, which represents the last derived age observed in the wave at the end of the sample time period (t = $\tau + s_i$). X_{it} is a vector of socio-demographics covariates based on previous literature, and including marital status, gender, number of individuals in the household, education level, unemployment rate in the local region, and household income, and m_{it-1} is mental health score of individual i at time t-1.

The analysis is conditioned a stock sample, implying that all periods prior to the selection period can be ignored, given that individuals are observed to be in employment at the beginning of the sample time period. The conditional probability of observing the event history of an individual without a complete sequence of responses for the whole sample period is:

$$P(T_i > t|T_i > \tau - 1) = \prod_{t=\tau}^{\tau+s_i} (1 - h_{it}) \tag{6}$$

The conditional probability of observing an individual with a complete sequence between τ and the time of the interview is:

$$P(T_i = t|T_i > \tau - 1) = h_{i\tau+s_i} \prod_{t=\tau}^{\tau+s_i-1} (1 - h_{it}) \tag{7}$$

Equation(7) can be simplified by:

$$(\frac{h_{i\tau+s_i}}{1 - h_{i\tau+s_i}}) \prod_{t=\tau}^{\tau+s_i} (1 - h_{it}) \tag{8}$$

Combining Equation(7) and Equation(8), for the corresponding log-likelihood (with and with-out a complete sequence) for the whole sample yields:

$$logL = \sum_{i=1}^{n} \delta_i log(\frac{h_{i\tau+s_i}}{1 - h_{i\tau+s_i}}) + \sum_{i=1}^{n} \sum_{t=\tau}^{\tau+s_i} log(1 - h_{it}) \tag{9}$$

The log likelihood function in Equation (9) depends on the labour market status of the individual i at the end of the sample time period. The individual can retire before the sample time period, $\delta_i = 1$, or to have a complete duration spell if $\delta_i = 0$.

For individuals who stay in the labour market until the last observed wave, denoted by $y_{it} = 0$, the value remains the same for all spell periods. On the other hand, for individuals who exit, denoted by $y_{it} = 0$ for all periods except the exit period, where it becomes $y_{it} = 1$.

The log-likelihood can be simplified by replacing δ_i with y_{it} in Equation 9:

$$logL = \sum_{i=1}^{n} \sum_{t=\tau}^{\tau+s_i} y_{it}\, log(\frac{h_{it}}{(1-h_{it})}) + \sum_{i=1}^{n} \sum_{t=\tau}^{\tau+s_i} log(1-h_{it}) \tag{10}$$

The hazard rate is the rate of retirement at time t and measures how probable an observation is to retire as a function of age of the individual conditioned on surviving to t-1. The hazard rate is defined by applying a complementary log-log hazard rate, as follows:

$$h_{it} = 1 - exp(-exp(\beta_1 m_{it-1} + \beta_2 X_{it} + \theta_t)) \tag{11}$$

Where θ_t is the discrete-time baseline hazard in our model. We include age dummies for every year at risk of retiring early, starting from 53 to 64 years old, as the actual risk of retiring started at 53 years old. The age dummies between 50 and 52 are not included as they are captured in the model's constant. Therefore, we use the semi-parametric form of our hazard model, where $h_0(t)$ specified as a step function, using dummy variables for each year of age. We also estimate sub-samples of males and females separately as it is likely that the response of labour supply to health shocks differs by gender (Stock and Wise 1988).

Unobserved heterogeneity

Unobserved heterogeneity refers to individual-specific characteristics that are not captured by the observed covariates but still influence the outcome variable, such as differences in preferences, motivation, or unmeasured health conditions (Lemeshow *et al.* 2011). Ignoring these latent factors in a hazard model can lead to biased estimates of the baseline hazard and the coefficients of interest, as the model may attribute part of the variation in hazard time to observed covariates rather than to unobserved differences across individuals. This can result in premature censoring and distortions in the estimated duration dependence, producing either an overestimation or underestimation of the true effect of covariates (Balan and Putter 2019, Heckman and Singer 1984, Lancaster 1990, Nicoletti and Rondinelli 2010).

In general, neglecting unobserved heterogeneity or misspecifying its distribution can bias the estimation of the hazard rate and the coefficients associated with explanatory variables (Heckman and Singer 1984). Individuals with higher unobserved risk may exit the labour market earlier, leading to spurious negative duration dependence (Lancaster 1990). Moreover, when the form of heterogeneity is incorrectly modelled, the estimated relationship between regressors and the hazard rate may not reflect the true behavioural mechanisms (Nicoletti and Rondinelli 2010).

These biases are generally more problematic in duration models than in linear regressions, where the consequences of unobserved heterogeneity are less severe if it is uncorrelated with the regressors. In hazard settings, however, the nonlinearity of the model amplifies the effects of omitted heterogeneity, threatening the validity of the estimates (Nicoletti and Rondinelli 2010, Jenkins 1995, Lancaster 1990).

To assess the sensitivity of our estimates, we also follow Nicoletti and Rondinelli (2010) who use Monte Carlo simulations to evaluate the impact of misspecifying unobserved heterogeneity. Their findings suggest that discrete-time hazard models, particularly those using the complementary log-log specification, are relatively robust to certain forms of misspecification, though accounting for frailty can still improve accuracy.

In our empirical analysis, we begin by estimating a complementary log-log model without frailty, followed by specifications that incorporate unobserved heterogeneity using a gamma-mixture distribution (Jenkins 1997, 1995). Model fit is compared using the Akaike Information Criterion (AIC) to evaluate whether including unobserved heterogeneity meaningfully alters the results. This analysis directly informs the following section, where we extend the framework to include a two-stage residual inclusion (2SRI) approach to address potential endogeneity.

Two-stage Residual Inclusion in Discrete-time Hazard Settings

The methodology used in this study is a 2SRI approach in a discrete-time hazard analysis context. The 2SRI estimator is an alternative to the two-stage least squares (2SLS) estimator but applicable to non-linear settings, while in linear models 2SRI is equivalent to 2SLS. The 2SRI incorporates the endogenous variable in the second-stage in addition to the predicted residual term from the first-stage to further account for endogeneity. The 2SRI estimator relies on previous literature, including the control function approach proposed by Heckman and Robb Jr (1985) and the residual inclusion method introduced by Hausman and Taylor (1981) and Terza *et al*. (2008). The 2SRI has been shown to have desirable properties in a range of empirical contexts (Angrist *et al*. 1996).

In the first-stage, the model is estimated with the instrumental variable as the independent variable and the endogenous variable as the dependent variable. The predicted values of the residuals from this model are then used in the second-stage as an additional covariate to estimate the effect of the exposure variable on the hazard outcome. The 2SRI approach can improve the precision of the estimate and reduce bias due to unobserved confounding, but its validity relies on "strong" instrumental variables and the assumption that the residual from the first-stage is uncorrelated with the unobserved confounders.

We consider a discrete-time hazard model specification mentioned in Equation (10). Hence, the first-stage is further specified as:

$$m_{it} = \alpha_0 + \alpha_1 Z_{it-1} + \alpha_2 X_{it} + u_{it} \tag{12}$$

where u_{it} is a mean zero residual error independent of Z, given X_{it}.

In the first-stage, we regress the endogenous covariates on the exogenous instrumental variable Z, and estimate it using an ordinary least squares model (OLS). We obtain the predicted residuals, $\hat{u}_{it}$, of the endogenous covariates from Equation (12).

We assume that the instrumental variable, death of a close friend (Z), used in Equation (12) is exogenous, meaning it is uncorrelated with the error term in the outcome of the discrete-time hazard model in Equation (12). Additionally, given that the effects of a close friend's death on mental health may not be immediately observable and may take some time to materialise, we account for this potential time lag by including a lagged one-period death of a close friend as an instrumental variable in our analysis.

Moreover, we assume that the instrumental variable used in our approach satisfies the exclusion restriction assumption. This indicates that the instrumental variable (lag of death of a close friend), *Z*, used in the first-stage regression in Equation(12) is only affecting the outcome variable (early retirement) through its effect on the endogenous variable, *m* (mental health score at time t-1), and not through any other unobserved factors. If this assumption is violated, the instrumental variable approach may not be valid and the estimated effect of the endogenous variable, m_{it-1}, on the outcome variable, Y_{it}, may be biased.

We then predict the residuals from the first-stage as:

$$\hat{u}_{it-1} = m_{it-1} - (\hat{\alpha}_0 + \hat{\alpha}_1 Z_{it-2} + \hat{\alpha}_2 X_{it-1}) \tag{13}$$

These residuals, $\hat{u}_{it-1}$, represent the unexplained variation in the endogenous variable, mit–1 after controlling for the exogenous variation provided by the instrumental variable of the death of a close friend at the previous period (Z_{it-2}).

In the second-stage, we estimate the effect of the predicted residuals $\hat{u}_{it-1}$ on the likelihood of early retirement, Y_{it}, using a cloglog hazard model. After following Aalen (1989) and Wan *et al*. (2015) where the covariates are allowed to vary over time, we can show that:

$$logit(P(Y_{it} = 1 | \hat{u}_{it-1}, X_{it}, m_{it-1})) = \beta_0 + \beta_1 m_{it-1} + +\beta_2 X_{it} + \beta_3 \hat{u}_{it-1} + v_{it} \tag{14}$$

where $P(Y_{it} = 1| \hat{u}_{it-1}, X_{it})$ is the probability of $(Y_{it}) = 1$ given the predicted residuals $\hat{u}_{it-1}$ and the vector of socio-demographic covariates, X_{it}. The parameter $\beta 2$ represents the effect of the predicted residuals on logit ($P(Y_{it} = 1)$), controlling for Xit and m_{it-1}. v_{it} is a random error term and represents the unobserved factors that affect the outcome variable, Y_{it}. It captures all other factors that might influence Yit but are not included in the model, and it is assumed to be uncorrelated with the instrumental variable, Z_{it-2}, other control variables, X_{it}, and m_{it-1} conditioned on $\hat{u}_{it-1}$.

The discrete-time hazard rate is modelled using a complementary log-log specification, such that:

$$h_{it} = 1 - exp(-exp(\beta_1 m_{it-1} + \beta_2 X_{it} + \beta_3 \hat{u}_{it-1} + \theta_{it})) \tag{15}$$

where $\beta 1$ is the coefficient for mental health status for individual i at time t-1, $\hat{u}_{it-1}$ is the predicated residuals from the first-stage for individual i at time t-1, and other socio-demographic factors for individual i at time t, respectively. θ_{it} is the discrete-time baseline hazard.

We apply the cloglog hazard model that includes a set of exogenous covariates and a discrete-time baseline hazard, θ_{it}, that captures the unobserved factors that affect the hazard of retirement over time. Additionally, we use a semi-parametric approach to model the baseline hazard as a step function, which allows us to estimate the hazard rate at each discrete time t using dummy variables for age to retirement (refer to the next section (section 4) for additional explanation).

Data

This paper uses 18 waves (2001-2018) of the Household, Income and Labour Dynamics in Australia (HILDA)[1] survey. HILDA is a household based panel study which includes information about socioeconomic characteristics, family dynamics and labour market outcomes. The dataset consists of variables related to individuals and household characteristics, including labour market status, wages, and the health status of individuals. Individuals aged 15 and older are eligible for interview. Both personal and self-assessment questionnaires are used in order to obtain information about the participants.

Participants beyond the state retirement age are excluded, as their retirement decisions may not solely be influenced by adverse mental health conditions. During the time of our sample, the Australian government implemented a policy aligning the female state pension age with that of males, resulting in a notable increase in the female retirement age. Our process involves verifying the age of each female participant in relation to the specific wave, excluding those whose age is equal to or greater than the state pension age during that period.

Individuals with missing employment status (i.e. those who refused to respond or did not provide a response) and those with missing mental health scores, the main variable of interest, are excluded from the analysis. In total, 122,808 observations did not report or refused to report their general health status. Furthermore, 63,956 observations are excluded because individuals were above the statutory retirement age for both genders. Finally, 254,468 observations are excluded for being under the age of 50, leaving a total sample of 44,091 observations. When incorporating lagged mental health, we further exclude individuals who reported mental health in only one wave, resulting in a final analytical sample of 32,316 observations.

1 See https://www.melbourneinstitute.com/hilda/

Figure 1. Sample construction and exclusion criteria

Full HILDA sample: 364,427

Exclude respondents outside ages 50–64
254,468 (age < 50), 63,956 (age > 64)

Exclude women reaching pension-eligibility thresholds
1,912

Exclude observations with missing values of their labour market status
4,431

Exclude observations already reported retirement
3,916

Final estimation sample: 32,316

The derived mental health variable in HILDA captures the mental health status of participants as a summary score (1–100), derived from the Short-Form health survey (SF-36) following the method set out in Ware *et al*. (1993). The SF-36 is a self-reported, multidimensional, generic measure of health. Respondents answer 36 short questions regarding their general health status, both physical and mental, relative to other individuals of the same age group, and grade their responses on discrete scales between 1 and 5. The 36 survey items are used to produce an eight-scale profile of functional health and wellbeing as well as a psychometric score. This self-reported questionnaire is based on physical and mental health summary measures, as well as a preference-based health utility index (Ware *et al*. 1993).

The mental health variable derived from the SF-36 consists of five questions. The items invite the respondents to grade their mental health by replying on a discrete scale to how much they agree with the following statements: been a nervous person, felt so down in the dumps nothing could cheer you up, felt down, been a happy person, felt calm and peaceful. The last two items are reverse-coded so that all items correspond to the same qualitative effect. Figure 2 displays the distribution of the derived mental health variable. The distribution shows that mental health scores are heavily concentrated in the higher range, with most participants scoring between roughly 70 and 90. The distribution peaks around the mid-80s, indicating that a large share of respondents report relatively good mental health. The mean score of 74.92 on a 0–100 scale (where 100 represents the best mental health) reflects this overall skew toward higher values.

Figure 2. Distribution of the derived mental health variable in HILDA from the SF-36

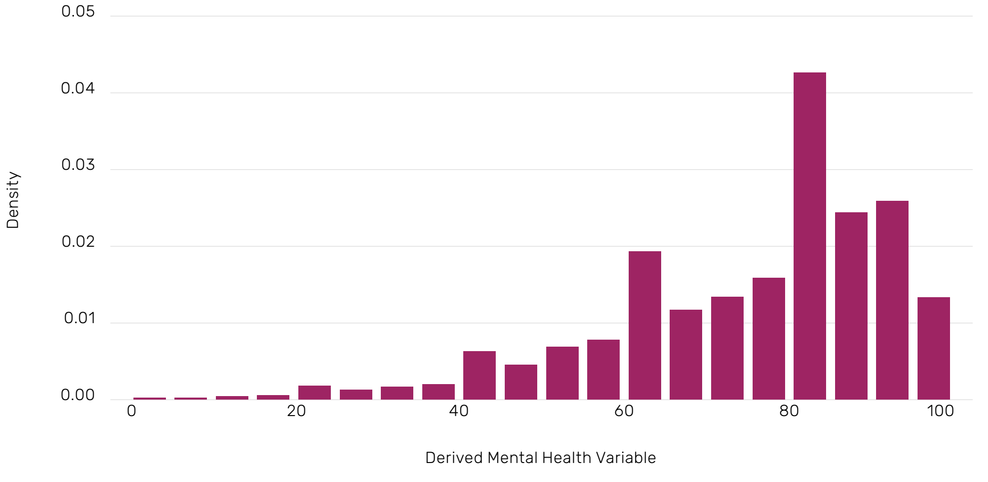

Figure 3 presents the kernel density estimates of the derived mental health scores, stratified by retirement status. The graph reveals a clear distinction in the distribution of mental health between retired and non-retired individuals. Specifically, the density of retired individuals is notably higher at lower mental health scores, indicating poorer mental health within this group. Conversely, the density for non-retired individuals dominates the right tail of the distribution, corresponding to higher (better) mental health scores. This suggests that, on average, retirees report worse general mental health compared to those who remain in the labour force.

Measurement errors in self-reported mental health may occur due to stigma and potential discrimination (Bharadwaj *et al*. 2017, Brohan and Thornicroft 2010, Rüsch *et al*. 2005). While survey-based self-reported questionnaires are commonly used to measure mental health, the potential for inaccuracy due to social desirability bias raises important considerations (Melzer *et al*. 2004). Furthermore, when individuals reported mental health, we expect the information to be influenced by evolving social norms and stigma surrounding mental health disabilities (Bharadwaj *et al*. 2017). However, as societal attitudes towards mental health have evolved in recent years, this shift is reflected in the increased per capita use of public mental health services in Australia (Figure 4).

Figure 3. Density of Mental Health, by Retirement Status

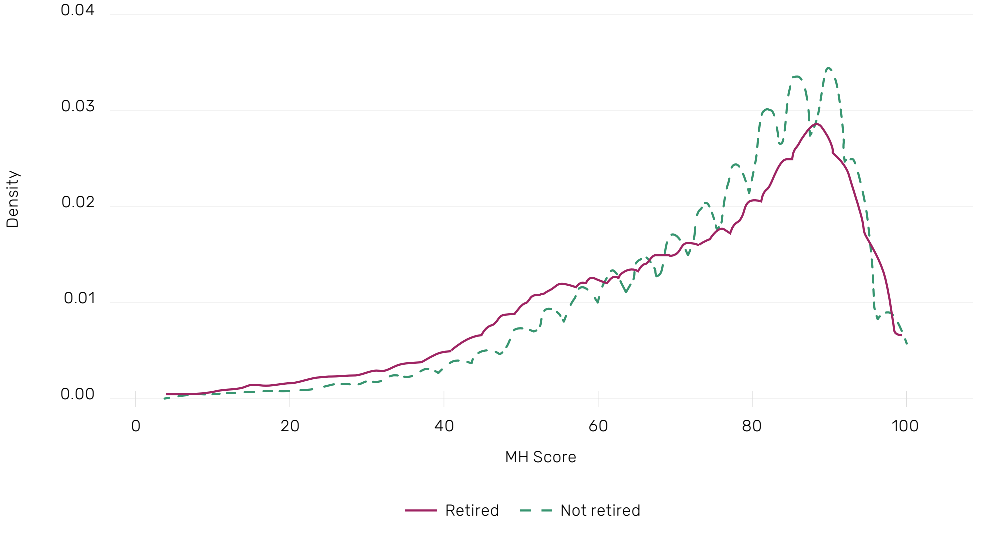

Figure 4. Number of Medicare Subsidised Mental Health Services Consumed by Australians of all providers per year per 1,000 population (Australian Institute of Health and Welfare 2021)

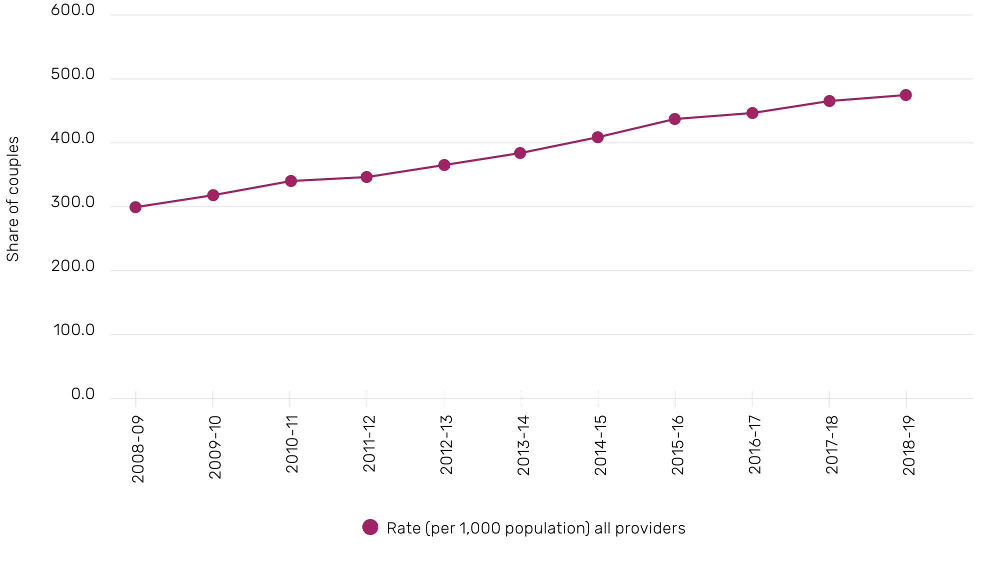

Despite the complex challenge posed by measurement error, we mitigate this concern by analysing the correlation between the Kessler Psychological Distress Scale-10 (K-10) questionnaire (Kessler *et al*. 2002) and our derived mental health variable. The Kessler Psychological Distress Scale includes a 10-item questionnaire, designed specifically for measuring mental health distress caused by anxiety and depressive symptoms following the method of Kessler *et al*. (2002), where higher scores indicate a greater likelihood of experiencing psychological distress. We run a validity check of the self-derived mental health status variable in HILDA, which is available in all 18 waves. However, the measure of Kessler Psychological Distress Scale is not available in every wave and hence cannot be used as the main measurement of mental health in our analysis. The regression method can provide supporting evidence of using the derived mental health variable available in HILDA, while avoiding a potential measurement error by including a more objective in its measurement.

Applying the K-10 as a benchmark, after recoding so that higher values indicate better mental health, we assess the validity of our self-reported mental health variable in the HILDA dataset, thereby enhancing the robustness of our analysis. Our examination reveals a highly statistically significant relationship between the K-10 scores and our derived mental health variable, exhibiting a substantial correlation between the two measures (Table 1). The correlation between these two measure reported to be high with more than 80 per cent. This comprehensive analysis helps alleviate potential concerns of measurement error and underscores the reliability of our findings. By including this approach, we align our investigation with the evolving landscape of attitudes towards mental health and enhance the precision of the derived mental health variable included in our analysis.

Table 1. OLS regression of the derived variable of mental health on K10

VARIABLES	Derived Mental Health
Kessler	2.203*** (0.014)
Constant	108.44*** (0.234)
Observations	13,588
R-squared	0.641
Correlation	-0.801

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

In our paper, an additional regression analysis is applied to examine the association between the derived mental health variable and a binary variable indicating whether individuals have received a formal diagnosis of depression or anxiety by a healthcare professional. This is a formal diagnosis and is more objective and less prone to measurement error. It is important to note that this binary variable is only available in wave 9 and wave 13 and has been queried for a relatively small subset of participants, totalling N=5,436. The regression results, as presented in Table 2, reveal a statistically significant correlation between the binary variable representing a positive diagnosis of either depression or anxiety and the derived mental health score.

Table 2. OLS regression of the derived variable of mental health on depression or anxiety

VARIABLES	Derived Mental Health
Dependanx	-9.470*** (0.326)
Constant	71.96*** (0.244)
Observations	5,436
R-squared	0.134
Correlation	-0.365

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

We include mental health lagged one period to exploit the timing of events between shock to mental health on retirement decisions. This removes bias that would occur if mental health status was measured following the retirement decision during any particular wave. By conditioning on a previous reported mental health status, we can be assured that any health "shock" occurred prior to a retirement decision. Lagged health may also be more informative about the decision to retire than contemporaneous health status, simply because transitions take time. It also may take time for an individual to understand that they suffer from a mental health disorder and to adjust to their new health condition, while observing its effect on labour productivity.

We follow Jones (2009), Zucchelli *et al*. (2007), Disney *et al*. (2006) and include additional covariates to control for other socioeconomics characteristics that might affect early retirement decisions, such as gender, highest educational attainment, and local unemployment rate. Table 3 shows definitions of all the variables included in our analysis. Additionally, we rescaled the household income by the consumer price index in Australia based on 2001 as the base year, in order to capture inflation and ensure values are expressed in constant Australian dollars.

Table 3. Variable Definitions

Variable name	Variable Label	Definition and Measurement
Retired	Reported to be retired	Binary: 1 = retired, 0 = otherwise
Lagged mental health	Lagged self-reported mental health	Continuous scale (0–36), higher = worse mental health
Derived mental health	Mental health status derived from the self-reported GHQ	Continuous GHQ-derived score (0–36), higher = better mental health
Local unemployment rate	Unemployment rate in major statistical region	Continuous percentage (%)
Household size	Number of in-scope persons in household	Count variable
Degree	Holds an academic degree	Binary: 1 = degree, 0 = otherwise
Married	Legally married	Binary: 1 = married, 0 = otherwise
Initial mental health	Initial mental health status	Continuous scale (0–36), baseline measure
Income	Annual income in Australian dollars ($K)	Continuous, measured in thousands of AUD
Death of a friend	Death of a close friend in the last 12 months (instrument)	Binary: 1 = yes, 0 = no
Age	Respondent's age	Continuous, in years
Negative health shock	Negative health shock of physical functioning	Binary: 1 = experienced shock, 0 = no shock
Kessler	Kessler Psychological Distress Scale-10	Continuous scale (10–50), higher = better mental health (rescaled)
Dependanx	Diagnosis of depression and/or anxiety	Binary: 1 = diagnosed, 0 = not diagnosed
D	Early retirement indicator	Binary: 1 = early retired, 0 = censored
T	Analysis time when record ends	Continuous, measured in years
Wave	Survey wave	Integer, indicates survey round

DV: Dummy Variable

An additional measure is incorporated by conditioning mental health on an exogenous indicator of a negative physical health shock. Following Apouey *et al*. (2019), we construct this measure by comparing each individual's expected and actual physical health status across survey waves. Expectations of next period health are derived from a question in the SF-36 survey contained in the HILDA dataset, where respondents indicate their level of agreement (from 1 to 5) with the statement, "I expect my health to get worse." These responses are coded so that higher values reflect a stronger expectation of health deterioration.

Let PH_{it} denote the self-reported physical functioning score for individual i at time t, and Eit represent the expected change in health status for the following wave, derived from responses to the statement "I expect my health to get worse." We first compute the change in actual physical health between two consecutive waves as $\Delta PH_{it} = PH_{it} - PH_{i,t-1}$. We then classify this change into three categories:

- -1 if $\Delta PH_{it} < -1$ standard deviation (decline in health)
- 0 if $|\Delta PH_{it}| \leq 1$ standard deviation (no significant change)
- $+1$ if $\Delta PH_{it} > 1$ standard deviation (improvement in health)

Next, we construct a binary indicator of an unexpected negative health shock, denoted by S_{it}, defined as follows:

$$S_{it} = \begin{cases} 1 & \text{if the individual expected their health to stay the same or improve } (E_{it} \geq 0), \\ & \text{but experienced an actual decline } (\Delta PH_{it} < -1), \\ 0 & \text{otherwise.} \end{cases}$$

In this definition, $S_{it} = 1$ captures cases where the participant's physical health worsened substantially (by at least one standard deviation) despite expecting it to remain stable or improve. Conversely, if an individual anticipated a deterioration in their health and it indeed worsened, Sit = 0, since the change was anticipated. This formulation enables us to isolate unexpected negative health shocks, allowing for a more accurate assessment of how unforeseen declines in physical health influence early retirement decisions.

We also include a discrete-time hazard analysis method that helps us observe the influence of health on the timing of the decision to retire early. To achieve this goal, individuals who are at risk of an early retirement at the start of the HILDA survey are examined. This sample is referred to as a stock sample (Jenkins 1995). We define the stock sample to be comprised of individuals who are aged 50 years or above, have completed a full interview, and are observed to be employed or self-employed in the first wave of the survey. Our main discrete-time hazard analysis sample consists of 9,455 individuals aged between 50 and 64. When additional sociodemographic factors are included, the sample comprises 9,054 individuals, 4,928 males and 4,126 females.

To identify individuals at risk of early retirement in our study, we set an upper age limit of 64, representing the age until which individuals are considered susceptible to early retirement. This choice is consistent with the state retirement age in Australia, which is 65 during the covered period. Individuals aged 64 and below are included in our stock sample, recognising that those over 64 may have different retirement motivations, such as eligibility for social security benefits. Notably, for female participants, we account for variations in the state retirement age until July 2013 when the retirement age was similar for both males and females. This approach determines the observation period for each female participant based on her birth year and the specific state retirement age during the relevant wave.

We follow individuals in every wave until they reach state retirement age or exit the labour market. Individuals cannot be included in all the 18 waves, even if they joined the sample in the first wave of the analysis, since we only estimate the model with workers aged between 50 and 64. Thus, the maximum number of waves per individual is 15. Individuals are followed from the first wave of the survey until they retire, which is assumed to be an absorbing state, are lost to follow-up, or stay in the labour market after state retirement age (65 years old).

We use the death of a close friend as an IV for mental health in both our linear (2SLS) and nonlinear (2SRI) approaches. The relevant question in the HILDA survey asks: "Did any of these happen to you in the past 12 months? Death of a close friend?". We include the IV lagged by one period to align with the lagged structure of our mental health variable. The rationale is that the impact of a close friend's death on an individual's mental health may take time to manifest. Moreover, because the question refers to events in the past 12 months, the death could have occurred at any point within that timeframe, from one month to nearly a year ago, and may already have influenced the previously measured mental health score, depending on the timing of previous wave. By lagging the IV, we aim to allow sufficient time for the event to affect mental health and better capture its causal impact.

Results

Summary Statistics

Table 4 presents summary statistics for the key variables used in the analysis. Approximately 17 per cent of the observations report being retired, indicating a relatively small but meaningful of early retirement in the sample. The average lagged subjective mental health score is around 74.9 (out of 100), with a comparable mean of 74.4 for the initial mental health measure included in the longitudinal analysis, suggesting general stability in reported mental health over time. However, both variables exhibit considerable variation (standard deviation of roughly 17.4), highlighting notable heterogeneity in mental health within the sample.

The average annual income is AU$86,190, with a substantial standard deviation of AU$81,480, indicating a wide dispersion in earnings. Negative health shocks are relatively rare, affecting only 7.5 per cent of the sample, while 13 per cent experienced the death of a close friend, a variable used to capture exogenous emotional distress. Women constitute approximately half the sample (49 per cent), and 62 per cent of individuals are married. About 23 per cent hold a university degree, and the average household size is 2.5 persons. The local unemployment rate, used as a proxy for labour market conditions, averages around 5.3 per cent but ranges from 1.9 per cent to 8.8 per cent, suggesting some regional variation in economic environments.

Table 4. Descriptive Statistics

Variable	Obs	Mean	Std Dev	Min	Max
Retired	44,091	0.169	0.375	0	1
Lagged Mental Health	32,316	74.92	17.42	4	100
Initial Mental Health	38,676	74.41	17.40	0	100
Income	44,091	86.19	81.48	0	1030.59
Negative Health Shock	44,091	0.075	0.26	0	1
Female	44,091	0.49	0.50	0	1
Married	44,091	0.62	0.49	0	1
Degree	41,428	0.23	0.42	0	1
Household size	44,091	2.54	1.23	1	14
Local Unemployment Rate	44,090	5.31	1.11	1.9	8.8
Death of a Close Friend	35,784	0.13	0.34	0	1

DV: Dummy Variable

Table 5 presents descriptive statistics for the variables of interest, disaggregated by individuals' early retirement status and compared with the characteristics of the restricted sample, which includes individuals outside the eligible age range excluded from the main analysis (e.g. pension-aged individuals or younger than 50 years old. On average, individuals report good or very good lagged self-assessed mental health, with a mean of around 75, though this declines among those who have retired early. A notable difference is also observed in the gender composition: approximately 48 per cent of individuals who remain in the labour force are female, compared with around 54 per cent among those who retire early.

Table 5. Descriptive Statistics, by Retirement Status and Restricted Sample

Variable	All	Pre-Retirement	Retirement	Restricted Sample
Retirement Status	0.169	0	1	0.112
Lagged Mental Health	74.924	75.600	71.913	73.859
Academic Degree	0.233	0.248	0.161	0.228
Income	86.193	92.512	55.214	83.441
Female	0.492	0.483	0.538	0.490
Household Size	2.539	2.623	2.127	3.465
Married	0.619	0.619	0.615	0.297
Local Unemployment Rate	5.314	5.309	5.337	5.292
Negative Health Shock	0.075	0.077	0.067	0.087
Death of a Friend (IV)	0.131	0.128	0.143	0.099

Other variables in the analysis show that most individuals in the sample do not have an academic degree. Table 5 also provides a comparison of the information on the local unemployment rate and the number of individuals in the household. The average local unemployment rate is highest for the post-retirement group at around 5.33 per cent, which is slightly similiar to the pre-retirement group (5.31 per cent). The average household size is 2.54 individuals, with 2.62 individuals for the pre-retirement and slightly lower numbers for post-retirement (2.13). This may indicate that individuals tend to have fewer household members after retirement, possibly due to adult children leaving the household or the passing of a spouse.

The restricted sample provides a useful benchmark to assess whether individuals excluded from the core analysis differ systematically from the main study population. Compared with the full sample, the restricted group displays lower average mental health scores and a higher proportion of negative health shocks, suggesting that younger and older excluded individuals may experience different health dynamics. The restricted group also shows lower marriage rates and substantially larger household sizes, reflecting life-cycle differences in living arrangements. These patterns confirm that the main analytical sample is more homogeneous in terms of age-related socioeconomic characteristics, reinforcing the internal validity of the empirical strategy. The individuals included and excluded in the sample share unsystematic characteristics that we expect to not bias the results of the main analysis.

Interestingly, the number of people experiencing negative health shock decreasing post-retirement, suggesting a potential association between an exogenous physical health shock and early retirement. These findings raise intriguing questions about the role of health-related events in influencing retirement decisions and warrant further investigation.

Linear analysis

Table 6 presents the LPM estimates of the effect of lagged mental health on the probability of early retirement. Across all three specifications, lagged mental health is statistically significantly and negatively associated with retirement, suggesting that poorer mental health in the previous wave increases the likelihood of early exiting the labour market. In the most basic model (column 1), which includes only the mental health variable, the coefficient is −0.18 per cent and highly statistically significant at the 1 per cent level. This result is robust to the inclusion of sociodemographic controls (column 2), where the coefficient remains negative and significant, though slightly attenuated to −0.15 per cent. When age dummies are also included to account for nonlinear age-related retirement patterns (column 3), the magnitude of the effect increases slightly to −0.17 per cent, reinforcing the conclusion that deteriorating mental health is a strong predictor of early retirement.

Table 6. Effect of Lagged Mental Health on Retirement (LPM Models)

	(1)	(2)	(3)
Dependent variable:		*Retired*	
Lagged mental health	−0.0018*** (0.0001)	−0.0015*** (0.0001)	−0.0017*** (0.0001)
Sociodemographic controls[a]	–	Y	Y
Age dummies (53–64)	–	–	Y
Observations	32,316	32,057	32,057
R^2	0.0067	0.0495	0.1067

[a]Includes degree, gender, household size, marital status, and regional unemployment rate.
Note: Standard errors in parentheses. ***p < 0.01

Since mental health is likely endogenous, potentially influenced by unobserved factors that also affect retirement decisions, Table 7 presents estimates from instrumental variable (2SLS) models using the death of a close friend in the previous period as an instrument for mental health. In Column 1, where no control variables are included, the effect of lagged mental health on retirement is substantially larger in magnitude, more than six times greater, compared to the corresponding LPM estimate −0.18 per cent, and remains statistically significant at the 1 per cent level. This suggests that the LPM model may fail to adequately capture the causal effect of mental health on retirement, potentially due to omitted variable bias or reverse causality.

Table 7. Effect of Lagged Mental Health on Retirement (2SLS Estimates)

	(1)	(2)	(3)
Dependent variable:		*Retired*	
Lagged mental health	−0.0118*** (0.0032)	−0.0067* (0.0040)	0.0036 (0.0032)
Sociodemographic controls[a]	–	Y	Y
Age dummies (53–64)	–	–	Y
Observations	29,603	29,432	29,432
First-stage F-stat	59.5	33.55	38.8

[a]Includes degree, gender, household size, marital status, and regional unemployment rate.
Note: Standard errors in parentheses. **p < 0.05, ***p < 0.01. All models use L.ledfr as instrument.

When sociodemographic controls are added (column 2), the coefficient decreases to −0.7 per cent and is statistically significant at the 10 per cent level, indicating the robustness of the relationship after adjusting for confounding variables. However, once age dummies are included in column 3, the coefficient becomes minimal and statistically insignificant.

Across all 2SLS models, the first-stage Cragg-Donald F-statistics (ranging from 33.6 to 59.5) exceed the Stock-Yogo critical threshold of 16.38 for the 10 per cent maximal IV size, indicating that the instrument is sufficiently strong and the estimates are not likely to suffer from weak instrument bias. This exceeds the commonly used rule-of-thumb threshold for identifying a strong instrument, with an F-statistic above 10 and hence indicating that 2SLS may capture the causal effect of mental health on retirement.

Together, these findings support the hypothesis that poor mental health also causally increases the probability of early retirement. However, the attenuation and eventual insignificance of the coefficient once richer controls are included, particularly age, highlight the importance of accounting for life-course factors when modelling retirement behaviour. This highlights the need to apply a longitudinal framework to better capture how individuals' retirement decisions evolve over time. In particular, we expect substantial heterogeneity in retirement behaviour across the age distribution, with older individuals likely exhibiting different patterns and sensitivities compared to their younger counterparts.

The fixed-effects estimation results presented in Table 8 show that the coefficient on lagged mental health is negative, consistent in direction with previous models, suggesting that better mental health is associated with a lower probability of retirement. However, unlike earlier findings, this effect is not statistically significant within the fixed-effects framework. This implies that, after accounting for time-invariant individual characteristics, changes in mental health do not have a statistically meaningful impact on the likelihood of retirement within individuals over time. The lack of significance may be due to limited within-individual variation in mental health or retirement status.

Table 8. Effect of Lagged Mental Health on Retirement (Fixed-Effects Model)

Dependent variable:	*Retired*
Lagged mental health	−0.0001 (0.0002)
Income	−0.0002*** (0.00004)
Degree	−0.0757 (0.0408)
Sociodemographic controls[a]	Y
Age dummies (53–64)	Y
Observations	32,057
Within R^2	0.0688

Note: Robust standard errors clustered at the individual level in parentheses.
***p < 0.01
[a] Includes degree, gender, household size, marital status, and regional unemployment rate.

As shown in Table 9, retirement status remains constant for the majority of the sample across panel waves; only a small proportion of individuals report a change in retirement status, and those who do typically maintain that status over time. This limited within-person variation in retirement partly explains why fixed-effects estimates tend to be less precise or insignificant for retirement outcomes, there is simply insufficient "within" variation to identify the effect reliably.

Table 9. Within-Individual Variation in Retirement Status

Retired Variation	Frequency
0 (No variation)	27,310
1 (Variation)	16,781
Total	44,091

At the same time, other variables such as household income and household size show statistically significant effects, with higher income and larger household size reducing the probability of retirement. This indicates that economic resources and family context remain important predictors of retirement decisions.

Discrete-time hazard model

We apply discrete-time hazard analysis to examine how the results of the baseline model change when tracking individuals longitudinally, wave by wave, until they report early retirement or are excluded from the sample due to attrition or other circumstances. This approach allows us to capture the timing of retirement decisions and better account for dynamic factors influencing exit from the labour market.

Figure 5 presents Kaplan-Meier survival estimates of the probability of remaining in the labour market across successive time intervals leading up to retirement. Mental health is measured at time $t - 1 = t0$ (baseline interview) and treated as a fixed covariate in the survival analysis. The probability of early retirement increases as individuals approach state retirement age. Mental health status is measured at time $t - 1$, representing the most recent observation prior to the retirement outcome, to ensure consistency with the econometric specifications presented in the subsequent tables. Individuals reporting excellent or good mental health at $t - 1$ are more likely to remain in the labour market at time t compared to those with poorer mental health. Conversely, workers who report very poor mental health exhibit the lowest probability of remaining in the labour market relative to individuals who describe their mental health as excellent, good, fair, or poor.

Figure 5. Kaplan–Meier survival estimates, by baseline MH status.

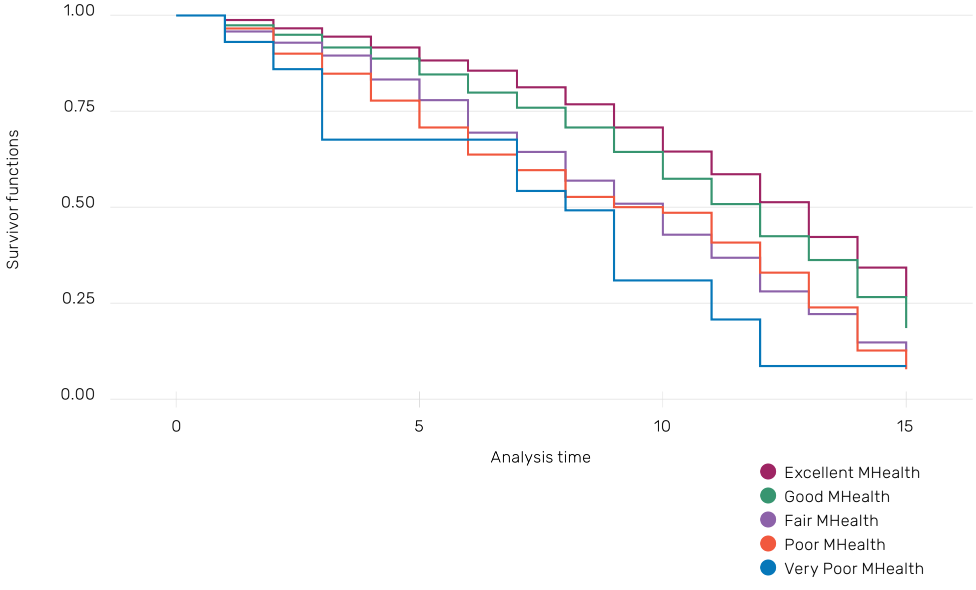

We estimate a series of discrete-time hazard models using a complementary log-log specification to assess how lagged mental health influences the probability of early retirement across waves. Table 10 presents six model specifications with increasing levels of control. In Model (1), we estimate the baseline relationship between lagged mental health and retirement. The results reveal a strong and statistically significant negative association, with a coefficient of −0.0139 ($p < 0.01$). This implies that poorer mental health in the previous wave is associated with a higher hazard of early retirement.

Model (2) introduces a control for initial mental health status, allowing us to net out individuals' initial condition. Both the lagged and baseline mental health variables remain statistically significant, though the magnitude of the lagged effect decreases to −0.09 per cent. This suggests that changes in mental health, rather than static poor levels alone, contribute to early exit from the labour market. Additionally, in Model (3), we add a dummy for experiencing a recent negative health shock. Although the mental health coefficients remain stable and significant, the health shock itself is not statistically significant ($p = 0.178$), implying that acute shocks do not independently predict retirement once mental health is controlled.

Model (4) incorporates household income. The coefficient of lagged mental health remains significant and relatively unchanged −0.086 per cent, while income is also negatively and significantly associated with retirement ($p = 0.017$). This suggests that financial security may partially buffer the impact on retirement decisions. In Model

(5), we further control for sociodemographic characteristics including education, gender, household size, and regional unemployment rate. The effect of lagged mental health remains robust and significant (−0.076 per cent, $p < 0.01$), and household size and having an academic degree are also statistically significant.

Finally, Model (6) adds a full set of age dummies to capture non-linear age-specific retirement risks. The coefficient for lagged mental health remains significant and stable (−0.083 per cent, $p < 0.01$), confirming the robustness of the relationship. Notably, the inclusion of age controls substantially improves model fit and reduces the size and significance of several other covariates. Many age dummies, particularly those in the early 50s, are large and negative with reduction of the effect size for the older age groups, as expected. This pattern aligns with expectations: the negative association between age and early retirement diminishes as individuals approach the statutory retirement age, reflecting prevailing social norms around the timing of retirement.

We also observe a gradient across educational attainment compared to the baseline category of not having any educational qualifications, with higher level of education are positively linked with a decreasing hazard of retiring. Possessing an academic degree, regardless of other variables, exhibits a substantial positive relationship in the probability of retirement. Moreover, the number of individuals in the household has a negative association with the likelihood not to retire early. According to our results, households with a larger number of living individuals would be less likely to leave the labour market earlier than state retirement age, independently of their health difficulties. The results of the full model are reported in Table 16.

Overall, the discrete-time hazard analysis reinforces earlier findings from the LPM and IV models, demonstrating stronger statistical significance and confirming that lagged mental health is a consistent and significant predictor of early retirement. This effect remains robust even after controlling for initial mental health, physical health shocks, income, sociodemographic characteristics, and age-dummies.

Discrete-time hazard model (including unobserved heterogeneity)

In this section, we include a model with unobserved heterogeneity and test the validation of this model compared to a model without frailty. Table 11 presents the results of a gamma-distributed frailty following the method of Jenkins (1997) for observing frailty. The results show that the expected effect of mental health on early retirement decision is -0.014 for the model without frailty and follows the same pattern of the model with frailty. The effect is statistically significant for both models.

Moreover, we estimate the model fit using information criteria (Akaike and Bayesian information criterions) which were greater with the complementary log-log estimator without frailty.

Table 10. Effect of Lagged Mental Health on Early Retirement (discrete-time hazard Models)

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable:			*Early retirement (hazard)*			
Lagged mental health	−0.0139*** (0.0021)	−0.0091*** (0.0029)	−0.0094*** (0.0029)	−0.0086*** (0.0030)	−0.0076*** (0.0029)	−0.0083*** (0.0029)
Observations	29,603		29,432		29,432	
Initial mental health	–	Y	Y	Y	Y	Y
Health shock	–	–	Y	Y	Y	Y
Income	–	–	–	Y	Y	Y
Sociodemographics[a]	–	–	–	–	Y	Y
Age dummies (53–64)	–	–	–	–	–	Y
Observations	9,455	9,139	9,139	9,139	9,054	9,054

[a] Includes: degree, gender, household size, marital status, and regional unemployment rate.
Note: Standard errors in parentheses. ***p < 0.01. Estimated using complementary log-log models.

Table 11. Complementary Log-Log Models of Early Retirement: With and Without Gamma-Distributed Heterogeneity

	No Unobserved Heterogeneity	Unobserved Heterogeneity
Lagged mental health	−0.0139*** (0.0021)	−0.0138*** (0.0036)
Constant	−1.596*** (0.162)	−1.597*** (0.273)
ln(var(ν))	– –	−14.852 (64.841)
Gamma variance	–	3.55×10^{-7}
Observations	9,455	9,455

Note: Standard errors in parentheses. ***p < 0.01. Unobserved Heterogeneity is estimated assuming Gamma frailty.

This shows the similar effect of the independent variable and indicates that our analysis does not set to include frailty. As frailty is not much a concern in this paper and the results do not vary between estimators with or without frailty, we have chosen to apply the complementary log-log regression for the discrete-time hazard analysis assuming no prevalence of frailty in our paper.

Two-stage residual inclusion

To assess the validity of our IV approach, we first estimated the first-stage regression applying LPM, where we regressed mental health score (t-1) on the death of a close friend at time t-2 (Table 17).

The results from the two-stage residual inclusion (2SRI) model presented in Table 12 reinforce the negative and statistically significant impact of lagged mental health on the likelihood of early retirement. Specifically, the coefficient of −1.1 per cent (significant at the 1 per cent level) indicates that a deterioration in mental health in the previous period is associated with a higher hazard of transitioning into early retirement.

Table 12. Two-Stage Residual Inclusion

	Two-Stage Residual Inclusion (2SRI)	
Variable	Coefficient	(SE)
Second-stage of 2SRI: Early Retirement		
Lagged mental health	-0.011***	(0.003)
Initial mental health	-0.008***	(0.003)
Income	0.001	(0.001)
Degree	-0.199*	(0.105)
Negative health shock	-0.175	(0.194)
Observations	7,335	
First-stage of 2SRI: OLS Regression on Mental Health Status		
Lagged death of a friend	-1.954***	(0.518)

Note: Standard errors in parentheses.
***p < 0.01, **p < 0.05, *p < 0.1
F-statistic (First-stage): 14.18

Initial mental health status, representing time-invariant baseline psychological conditions, also emerges as a significant predictor, with a coefficient of −0.08 per cent (significant at the 1 per cent level). This highlights that both dynamic (lagged) and baseline (initial) mental health exert meaningful influence on retirement decisions. However, income does not appear to have a statistically significant effect on early retirement in this model, with a small positive coefficient (0.001) and a p-value above conventional thresholds. This suggests that after accounting for health, sociodemographics, and age effects, income levels may not directly be associated with early retirement decisions in this sample.

Educational attainment, proxied by holding an academic degree, shows a marginally significant negative effect (–1.99 per cent, p = 0.057), implying that more educated individuals are less likely to retire early, possibly reflecting greater access to less physically or mentally demanding jobs, stronger labour market attachment, or different preferences regarding continued work.

Other sociodemographic controls, such as household size and local unemployment rate, also exhibit significant associations. A larger household size is associated with a smaller risk of early retirement, while higher local unemployment is positively related to retirement risk, potentially capturing economic discouragement effects. Age dummies behave as expected, with older individuals facing higher retirement hazards. The F-statistic for the instrument in the first-stage regression is 14.2, which exceeds the conventional threshold of 10, suggesting that the instrument is not weak.

These findings highlight the critical role of mental health, both dynamic and baseline, in shaping early retirement behaviour and underline the importance of integrating psychological well-being into labour market and retirement policy frameworks.

Sub-sample: Gender-specific

In this section, we present the results of the sub-sample analysis of estimating males and females separately, as previous research traditionally shows that the results differ when estimating the responses for labour market within the different genders. Figure 6 demonstrate the Kaplan-Meier survival estimates by sex. Based on the gender difference alone, independent of other factors, Figure 6 reveals that male workers in our sample exhibit a higher likelihood of remaining in the labour market compared to females. Additionally, mental health may affect males differently compared to females, with different proportions of individuals with a severe or mild mental health disorder in this sub-groups (Australian Bureau of Statistics 2022). The traditional approach when estimating mental health is to observe the effect of mental health on males, however, in this paper we include the results of the two sub-samples (Table 13).

Figure 6. Kaplan-Meier Survival Estimates, by sex

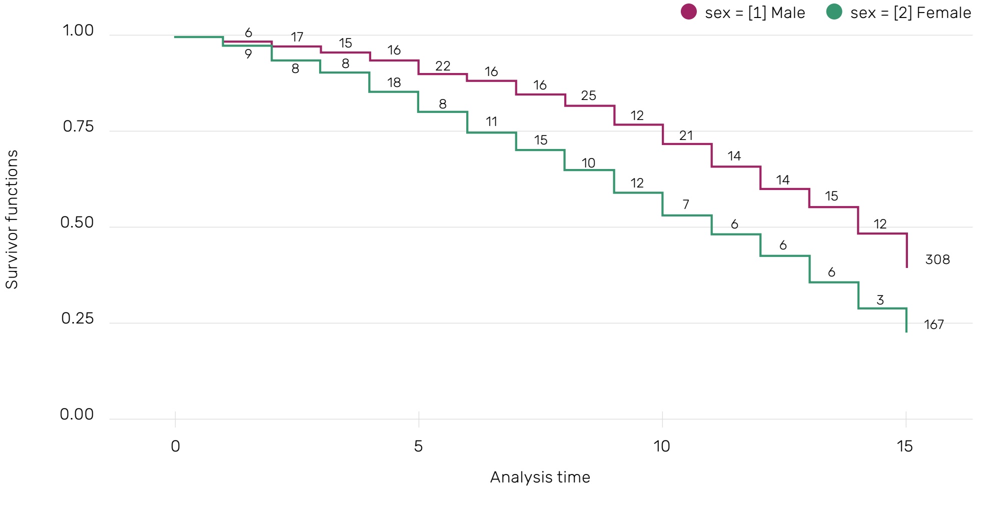

The results in Table 13 present estimates from four different empirical approaches: linear probability models using LPM and 2SLS, and nonlinear discrete-time hazard models using complementary log-log and 2SRI methods across genders. These models aim to evaluate the impact of lagged mental health on early retirement decisions.

Across most models, for both males and females, better lagged mental health is associated with a lower likelihood of early retirement, confirming that mental health issues increase the probability of exiting the labour force prematurely. This effect is consistently negative and statistically significant in nearly all models, except for the 2SLS specification for females, where the sign is positive and statistically insignificant, possibly due to weak instrument bias or sample-specific variation (first-stage F-stats below 10, at 7.907). Notably, the estimated effects tend to be larger in magnitude in the nonlinear models (complementary log-log and 2SRI) for males, particularly after addressing endogeneity in the discrete-time hazard framework.

Comparing across genders reveals notable heterogeneity. The effects of poor mental health on early retirement appear stronger for males than for females in most specifications, underlining the importance of gender-specific analyses when studying retirement behaviour.

Other covariates also exhibit gendered patterns. Education is positively associated with delaying retirement for females, with statistically significant effects across all models. For males, while the direction of the relationship is less consistent and mostly insignificant, the point estimates suggest that having an academic degree may in some cases be linked with earlier retirement, possibly reflecting differences in financial security, job satisfaction, or pension eligibility between educational groups.

Income exhibits consistent effects across linear and nonlinear models and across genders. The association is statistically significant in the linear models for both men and women but not in the nonlinear specifications. This discrepancy may reflect differences in the way time-to-event outcomes are captured or reduced statistical power in the discrete-time hazard models.

Table 13. Comparison of Coefficient Estimates Across Models and Gender

	LPM	2SLS	Cloglog	2SRI
Panel A: Female				
Mental health	−0.0014*** (0.0002)	0.0116 (0.0090)	−0.0057 (0.0041)	−0.0063 (0.0038)
Degree	−0.0307*** (0.0071)	−0.0635*** (0.0210)	−0.3473** (0.1651)	−0.3685** (0.1651)
Income	−0.0006*** (0.0001)	−0.0009*** (0.0002)	−0.0006 (0.0008)	−0.0006 (0.0008)
Negative health shock	−0.0316** (0.0140)	0.0384 (0.0488)	−0.5775* (0.3400)	−0.4792 (0.3403)
Observations	16,192	14,926	4,126	3,320
F-stat first-stage (2SLS)	–	7.907	–	–
Panel B: Male				
Mental health	−0.0022*** (0.0002)	−0.0004 (0.0034)	−0.0147*** (0.0041)	−0.0164*** (0.0038)
Degree	−0.0123* (0.0065)	−0.0185** (0.0073)	0.0174 (0.1380)	−0.0101 (0.1383)
Income	−0.0005*** (0.0001)	−0.0006*** (0.0001)	−0.0010 (0.0008)	−0.0012 (0.0009)
Negative health shock	−0.0078 (0.0125)	0.0069 (0.0214)	−0.0095 (0.2378)	0.0398 (0.2381)
Observations	15,865	14,506	4,928	4,015
F-stat first-stage (2SLS)	–	39.923	–	–

Notes: Each cell reports the coefficient (standard error in parentheses) for the corresponding variable and model.
Significance levels: *p < 0.10, **p < 0.05, ***p < 0.01.

The effect of health shocks is uniformly insignificant across all specifications. This suggests that physical health shocks may have a stronger influence on early retirement among women, possibly due to differing occupational exposures or caregiving responsibilities.

Discussion and conclusion

The present study examines the impact of mental health on labour market participation, specifically the decision to remain in the workforce or retire early. Exiting the labour market before reaching the state pension age can have negative implications, including reduced social interaction and financial insecurity. Policymakers therefore need a clearer understanding of how mental health influences early retirement decisions.

Previous literature has explored the relationship between health and labour market outcomes (Porru *et al*. 2019, Frijters *et al*. 2010, Disney *et al*. 2006, Ettner *et al*. 1997, Hamilton *et al*. 1997), yet the specific role of mental health among older workers remains underexplored. Empirical challenges such as reverse causality, endogeneity, and measurement error further complicate causal inference. To address these, our study applies a range of econometric methods designed to isolate the causal effect of mental health on early retirement.

First, we validate our derived mental health score using alternative subjective measures, in-cluding the Kessler scale (Kessler *et al*. 1999), and medical diagnoses. The high correlation between these measures suggests that our constructed mental health variable provides a reliable representation of individuals' mental health status. Second, we estimate linear probability models, including two-stage least squares (2SLS) regressions using the death of a close friend as an instrument for mental health. This event is assumed to affect labour market participation only through its impact on mental health (Frijters *et al*. 2010).

Although the instrument performs reasonably well, some limitations remain. Only about 13 per cent of respondents reported the death of a close friend, which may restrict variation. Furthermore, the meaning of a "close friend" is subjective and may differ across individuals. Nonetheless, the first-stage F-statistic exceeds 10 for the overall sample, confirming instrument strength. For females, the statistic falls slightly below this threshold at 7.9, but increases to around 40 under the 2SRI specification, reinforcing the validity of our identification strategy.

A further limitation relates to the absence of detailed occupational or industry controls. Job type can plausibly affect physical demands, psychosocial stressors, and exposure to health risks, and may therefore influence both mental health trajectories and early retirement decisions. In our setting, inclusion of occupation may introduce post-treatment bias given that job choice is itself shaped by earlier health and socioeconomic factors. While our models adjust for a broad set of socioeconomic characteristics, including education, income, and household circumstances, that capture some of the systematic differences associated with occupational sorting, we recognise that residual confounding may remain. Future research with richer longitudinal occupational histories would help to further disentangle these channels.

Third, we test for unobserved heterogeneity by estimating discrete-time hazard models both with and without frailty, assuming a gamma mixture distribution (Jenkins 1997, 1995). The results remain consistent across specifications, suggesting that unobserved frailty is not a major concern. Finally, to account for the timing of early retirement, we use discrete-time hazard models that handle right-censoring and exploit the longitudinal structure of the data. These models show a robust and statistically significant relationship between poorer mental health and higher likelihood of early retirement, even after controlling for physical health and socioeconomic characteristics.

Across all specifications, the results consistently show that individuals with better mental health are significantly less likely to retire early. The estimated effects, although modest in size, are economically meaningful. A one-standard-deviation improvement in mental health reduces the probability of early retirement by roughly 2–4 percentage points. However, an immediate deterioration in mental health would require several unit-level changes to produce a large effect on the probability of exiting the labour market prior retirement age.

The gender analysis further reveals that the effect of mental health is stronger and statistically significant for men, while smaller and insignificant for women. This suggests that men's retirement decisions are more sensitive to changes in mental health, possibly reflecting gendered social norms that link male identity more strongly with employment. Women, by contrast, appear more responsive to financial and family-related factors. Lastly, older individuals are more likely to retire early, reflecting both social norms and reduced financial or social penalties for doing so as they approach the state pension age.

The results of our study are broadly consistent with evidence on the effect of general health on labour market outcomes (Disney *et al*. 2006, Bryan *et al*. 2022, García-Gómez *et al*. 2010). However, given the complexity of the relationship between mental health and retirement be-haviour, not all studies reach the same conclusion. For example, Andersen *et al*. (2024) find no causal effect of mental health on early retirement when exploiting a fireworks disaster as an instrumental variable. These contrasting findings highlight the challenges inherent in isolating causal effects in this domain. Our paper contributes meaningfully to the existing literature by providing new evidence on this relationship and by applying a combination of econometric techniques to strengthen identification and robustness.

From a policy perspective, these findings highlight the importance of integrating mental health into retirement and employment policy. Early screening and workplace mental health support can help reduce premature labour market exits. Gender-sensitive interventions, such as promoting open discussions about men's mental health and reducing stigma, could further enhance labour force retention among older workers.

While the estimated effects are relatively modest, they remain economically significant, particularly in the context of cumulative public costs associated with early retirement. Improving access to mental health services and promoting workplace well-being could help delay retirement, thereby extending productive working lives and reducing fiscal pressures linked to ageing populations.

In conclusion, this study contributes new evidence on the causal effect of mental health on early retirement decisions using robust longitudinal and instrumental-variable methods. Our findings underscore that mental health, though not the only driver of early retirement, is a critical factor shaping labour market participation among older workers. Strengthening mental health support, particularly in the workplace, can generate both social and economic benefits.

References

Aalen, O. O. (1989), 'A linear regression model for the analysis of life times', *Statistics in Medicine* 8(8), 907–925.

Alderwick, H. and Dixon, J. (2019), 'The NHA long term plan', *BMJ (Clinical research ed.)* 364, l84.

Alexandre, P. K. and French, M. T. (2001), 'Labor supply of poor residents in metropolitan Miami, Florida: The role of depression and the co-morbid effects of substance use', *Journal of Mental Health Policy and Economics* 4(4), 161–174.

American Psychiatric Association (2013), *Diagnostic and Statistical Manual of Mental Disorders: DSM-5*, 5th edn, American Psychiatric Association, Washington, D.C.; London.

Andersen, S. H., Richmond-Rakerd, L. S., Moffitt, T. E. and Caspi, A. (2024), 'The causal effect of mental health on labor market outcomes: The case of stress-related mental disorders following a human-made disaster', *Proceedings of the National Academy of Sciences* 121(27), e2316423121. URL: https://www.pnas.org/doi/abs/10.1073/pnas.2316423121

Angrist, J. D., Imbens, G. W. and Rubin, D. B. (1996), 'Identification of causal effects using instrumental variables', *Journal of the American Statistical Association* 91(434), 444–455.

Angrist, J. D. and Pischke, J.-S. (2009), Mostly harmless econometrics: An empiricist's companion, Princeton University Press.

Apouey, B. H., Guven, C. and Senik, C. (2019), 'Retirement and unexpected health shocks', *Economics & Human Biology* 33, 116–123.

Australian Bureau of Statistics (2008), 'National Survey of Mental Health and Wellbeing: Summary of Results, 2007', Canberra: ABS, Cat. no. (4326.0).

Australian Bureau of Statistics 2020–2022, National Study of Mental Health and Wellbeing, ABS, viewed 13 July 2023, https://www.abs.gov.au/statistics/health/mental-health/national-study-mental-health-and-wellbeing/latest-release.

Australian Institute of Health and Welfare (2021), 'Mental health services in Australia', Retrieved from https://www.aihw.gov.au/reports/mental-health-services/mental-health-services-in-australia.

Baker, C. and Kirk-Wade, E. (2023), 'Mental health statistics: prevalence, services and funding in England', House of Commons Library .

Balan, T. A. and Putter, H. (2019), 'Nonproportional hazards and unobserved heterogeneity in clustered survival data: When can we tell the difference?', *Statistics in Medicine* 38(18), 3405–3420.

Bartel, A. and Taubman, P. (1979), 'Health and labor market success: The role of various diseases', *The Review of Economics and Statistics* pp. 1–8.

Bartel, A. and Taubman, P. (1986), 'Some economic and demographic consequences of mental illness', *Journal of Labor Economics* 4(2), 243–256.

Basu, A. and Rathouz, P. J. (2005), 'Estimating marginal and incremental effects on health out-comes using flexible link and variance function models', *Biostatistics* 6(1), 93–109.

Bazzoli, G. J. (1985), 'The early retirement decision: new empirical evidence on the influence of health', *Journal of Human Resources* pp. 214–234.

Bharadwaj, P., Pai, M. M. and Suziedelyte, A. (2017), 'Mental health stigma', *Economics Letters* 159, 57–60.

Bloom, D. E. and Canning, D. (2000), 'The health and wealth of nations', Science 287(5456), 1207–1209.

Böckerman, P., Haapanen, M. and Jepsen, C. (2022), 'Dark passage: Mental health consequences of parental death', *Journal of Human Resources* 58(1), 196–229.

Bound, J., Brown, C. and Mathiowetz, N. (2001), Measurement error in survey data, in 'Handbook of Econometrics', Vol. 5, Elsevier, pp. 3705–3843.

Bound, J., Schoenbaum, M., Stinebrickner, T. R. and Waidmann, T. (1999), 'The dynamic effects of health on the labor force transitions of older workers', *Labour Economics* 6(2), 179–202.

Brazenor, R. (2002), 'Disabilities and labour market earnings in Australia', *Australian Journal of Labour Economics* 5(3), 319–334.

Brohan, E. and Thornicroft, G. (2010), 'Stigma and discrimination of mental health problems: workplace implications'.

Bryan, M. L., Rice, N., Roberts, J. and Sechel, C. (2022), 'Mental health and employment: a bounding approach using panel data', *Oxford Bulletin of Economics and Statistics* 84(5), 1018–1051.

Bubonya, M., Cobb-Clark, D. A. and Ribar, D. C. (2019), 'The reciprocal relationship between depressive symptoms and employment status', *Economics & Human Biology* 35, 96–106.

Bönnings, C. and Tauchmann, H. (2015), 'Who opts out of the statutory health insurance? A discrete time hazard model for Germany', *Health Economics* 24(10), 1331–1347.

Burrell, L. V., Mehlum, L. and Qin, P. (2022), 'Parental death by external causes during childhood and risk of psychiatric disorders in bereaved offspring', *Child and Adolescent Mental Health* 27(2), 122–130.

Butler, J. S., Burkhauser, R. V., Mitchell, J. M. and Pincus, T. P. (1987), 'Measurement error in self-reported health variables', *The Review of Economics and Statistics* pp. 644–650.

Butterworth, P., Gill, S. C., Rodgers, B., Anstey, K. J., Villamil, E. and Melzer, D. (2006), 'Retirement and mental health: Analysis of the Australian National Survey of Mental Health and Wellbeing', *Social Science & Medicine* 62(5), 1179–1191.

Cai, L. and Kalb, G. (2007), 'Health status and labour force status of older working-age Australian men', *Australian Journal of Labour Economics* 10(4), 227–252.

Chatterji, P., Alegría, M., Lu, M. and Takeuchi, D. (2007), 'Psychiatric disorders and labor market outcomes: Evidence from the National Latino and Asian American study', *Health Economics* 16(10), 1069–1090.

Chatterji, P., Alegria, M. and Takeuchi, D. (2011), 'Psychiatric disorders and labor market out-comes: Evidence from the national comorbidity survey-replication', *Journal of Health Economics* 30(5), 858–868.

Disney, R., Emmerson, C. and Wakefield, M. (2006), 'Ill health and retirement in Britain: A panel data-based analysis', *Journal of Health Economics* 25(4), 621–649.

Ettner, S. L., Frank, R. G. and Kessler, R. C. (1997), 'The impact of psychiatric disorders on labor market outcomes', *ILR Review* 51(1), 64–81.

Fadare, O., Zanello, G. and Srinivasan, C. (2023), 'Stressor or succour? Examining the association between conflict, livestock assets, and farmers' mental health in Nigeria', *Economics & Human Biology* 49, 101234. URL: https://www.sciencedirect.com/science/article/pii/S1570677X23000151

Frijters, P., Johnston, D. and Shields, M. (2010), 'Mental health and labour market participation: Evidence from IV panel data models', Institute for the Study of Labor (IZA), IZA Discussion Papers.

García-Gómez, P., Jones, A. M. and Rice, N. (2010), 'Health effects on labour market exits and entries', *Labour Economics* 17(1), 62–76. URL: https://www.sciencedirect.com/science/article/pii/S0927537109000463

Garrido, M. M., Deb, P., Burgess Jr, J. F. and Penrod, J. D. (2012), 'Choosing models for health care cost analyses: issues of nonlinearity and endogeneity', *Health Services Research* 47(6), 2377–2397.

Grossman, M. (1972), 'On the concept of health capital and the demand for health', *Journal of Political Economy* 80(2), 223–255.

Hamilton, V. H., Merrigan, P. and Dufresne, É. (1997), 'Down and out: Estimating the relationship between mental health and unemployment', *Health Economics* 6(4), 397–406.

Hausman, J. A. and Taylor, W. E. (1981), 'Panel data and unobservable individual effects', *Econometrica: Journal of the Econometric Society* pp. 1377–1398.

Heckman, J. J. and Robb Jr, R. (1985), 'Alternative methods for evaluating the impact of interventions: An overview', *Journal of Econometrics* 30(1-2), 239–267.

Heckman, J. and Singer, B. (1984), 'A method for minimizing the impact of distributional assumptions in econometric models for duration data', *Econometrica: Journal of the Econometric Society* pp. 271–320.

Jenkins, S. P. (1995), 'Easy estimation methods for discrete-time duration models', *Oxford Bulletin of Economics and Statistics* 57(1), 129–138.

Jenkins, S. P. (1997), 'Estimation of discrete time (grouped duration data) proportional hazards models: pgmhaz', *Stata Technical Bulletin Reprints*, STB 17(7), 109–121.

Jones, A. M. (2009), 'Panel data methods and applications to health economics', *Palgrave Hand-book of Econometrics: Volume 2: Applied Econometrics* pp. 557–631.

Kessler, R. C., Andrews, G., Colpe, L. J., Hiripi, E., Mroczek, D. K., Normand, S.-L., Walters, E. E. and Zaslavsky, A. M. (2002), 'Short screening scales to monitor population prevalences and trends in non-specific psychological distress', *Psychological Medicine* 32(6), 959–976.

Kessler, R. C., Mickelson, K. D. and Williams, D. R. (1999), 'The prevalence, distribution, and mental health correlates of perceived discrimination in the united states', *Journal of Health and Social Behavior* pp. 208–230.
Lancaster, T. (1990), *The econometric analysis of transition data*, 17, Cambridge University Press.
Lee, Y.-C., Chatterton, M. L., Magnus, A., Mohebbi, M., Le, L. K.-D. and Mihalopoulos, C. (2017), 'Cost of high prevalence mental disorders: Findings from the 2007 Australian National Survey of Mental Health and Wellbeing', *Australian & New Zealand Journal of Psychiatry* 51(12), 1198–1211.
Lemeshow, S., May, S. and Hosmer Jr, D. W. (2011), *Applied survival analysis: Regression modeling of time-to-event data*, John Wiley & Sons.
Liddiard, P. (1978), *An Ageing Population: A Reader and Sourcebook*, Hodder & Stoughton.
Lu, C., Frank, R. G., Liu, Y., Shen, J. *et al*. (2009), 'The impact of mental health on labour market outcomes in China', *Journal of Mental Health Policy and Economics* 12(3), 157.
McManus, S., Bebbington, P., Jenkins, R. and Brugha, T. (2016), *Mental health and wellbeing in England: Adult Psychiatric Morbidity Survey 2014*, NHS Digital, Leeds.
Melzer, D., Buxton, J. and Villamil, E. (2004), 'Decline in common mental disorder prevalence in men during the sixth decade of life: Evidence from the national psychiatric morbidity survey', *Social Psychiatry and Psychiatric Epidemiology* 39, 33–38.
Mery, G., Wodchis, W. P. and Laporte, A. (2016), 'The determinants of the propensity to receive publicly funded home care services for the elderly in Canada: A panel two-stage residual inclusion approach', *Health Economics Review* 6(1), 1–18.
Neil, A. L., Carr, V. J., Mihalopoulos, C., Mackinnon, A. and Morgan, V. A. (2014), 'Costs of psychosis in 2010: Findings from the second Australian National Survey of Psychosis', *Australian & New Zealand Journal of Psychiatry* 48(2), 169–182.
Nerlove, M. (2005), *Essays in panel data econometrics*, Cambridge University Press.
Ngui, E. M., Khasakhala, L., Ndetei, D. and Roberts, L. W. (2010), 'Mental disorders, health inequalities and ethics: A global perspective', *International Review of Psychiatry* 22(3), 235–244.
Nicoletti, C. and Rondinelli, C. (2010), 'The (mis) specification of discrete duration models with unobserved heterogeneity: a Monte Carlo study', *Journal of Econometrics* 159(1), 1–13.
Olesen, S. C., Butterworth, P., Leach, L. S., Kelaher, M. and Pirkis, J. (2013), 'Mental health affects future employment as job loss affects mental health: Findings from a longitudinal population study', *BMC Psychiatry* 13, 1–9.
Porru, F., Burdorf, A. and Robroek, S. J. (2019), 'The impact of depressive symptoms on exit from paid employment in Europe: A longitudinal study with 4 years follow-up', *European Journal of Public Health* 29(1), 134–139.
Rüsch, N., Angermeyer, M. C. and Corrigan, P. W. (2005), 'Mental illness stigma: Concepts, consequences, and initiatives to reduce stigma', *European Psychiatry* 20(8), 529–539.

Sainsbury, R., Irvine, A., Aston, J., Wilson, S., Williams, C. and Sinclair, A. (2008), 'Mental health and employment', *White Rose*.

Siddiqui, S. and Ali Shah, S. A. (1997), 'Siddiqui-shah depression scale (ssds): Development and validation', *Psychology and Developing Societies* 9(2), 245–262.

Steele, F., French, R. and Bartley, M. (2013), 'Adjusting for selection bias in longitudinal analyses using simultaneous equations modeling: the relationship between employment transitions and mental health', *Epidemiology* 24(5), 703–711.

Stern, S. (1989), 'Measuring the effect of disability on labor force participation', *Journal of Human Resources* pp. 361–395.

Stock, J. H. and Wise, D. A. (1988), 'Pensions, the option value of work, and retirement'. Tchetgen, E. J. T., Walter, S., Vansteelandt, S., Martinussen, T. and Glymour, M. (2015), 'Instrumental variable estimation in a survival context', *Epidemiology* (Cambridge, Mass.) 26(3), 402.

Terza, J. V. (2018), 'Two-stage residual inclusion estimation in health services research and health economics', *Health Services Research* 53(3), 1890–1899.

Terza, J. V., Basu, A. and Rathouz, P. J. (2008), 'Two-stage residual inclusion estimation: Addressing endogeneity in health econometric modeling', *Journal of Health Economics* 27(3), 531–543.

Tompa, E. (2002), *The impact of health on productivity: empirical evidence and policy implications*, CIGI.

Wan, F., Small, D., Bekelman, J. E. and Mitra, N. (2015), 'Bias in estimating the causal hazard ratio when using two-stage instrumental variable methods', *Statistics in Medicine* 34(14), 2235–2265.

Ware, J. E., Snow, K. K., Kosinski, M. and Gandek, B. (1993), 'Sf-36 Health Survey', *Manual and Interpretation Guide* 2.

Wilkins, R. (2004), 'The effects of disability on labour force status in Australia', *Australian Economic Review* 37(4), 359–382.

Wooldridge, J. M. (2010), *Econometric analysis of cross section and panel data*, MIT press.

World Health Organization (2022), 'Ageing and health', https://www.who.int/news-room/fact-sheets/detail/ageing-and-health.

Ying, A., Xu, R. and Murphy, J. (2019), 'Two-stage residual inclusion for survival data and competing risks—An instrumental variable approach with application to seer-medicare linked data', *Statistics in Medicine* 38(10), 1775–1801.

Zhang, X., Zhao, X. and Harris, A. (2009), 'Chronic diseases and labour force participation in Australia', *Journal of Health Economics* 28(1), 91–108.

Zucchelli, E., Harris, A., Jones, A. M. and Rice, N. (2007), 'Health and retirement among older workers', *Health, Econometrics and Data Group Working Paper* 19.

A: Full-model: OLS Regression

Table 14. OLS Regression: Effect of Mental Health on Early Retirement

VARIABLES	Retired
Lagged mental health	-0.0017*** (0.0001)
Income	-0.0005*** (0.0001)
Female	0.0360*** (0.0041)
Married	0.0398*** (0.0048)
Degree	-0.0281*** (0.0049)
Household size	-0.0258*** (0.0020)
Local unemployment rate	-0.0080*** (0.0020)
Age 53	-0.1055*** (0.0088)
Age 54	-0.1035*** (0.0089)
Age 55	-0.0914*** (0.0090)
Age 56	-0.0567*** (0.0092)
Age 57	-0.0333*** (0.0093)
Age 58	-0.0184* (0.0094)
Age 59	0.0119 (0.0095)
Age 60	0.0583*** (0.0097)
Age 61	0.1052*** (0.0098)
Age 62	0.1388*** (0.0099)
Age 63	0.1757*** (0.0106)
Age 64	0.2228*** (0.0114)
Constant	0.4176*** (0.0159)
Observations	32,057
R-squared	0.1067

Standard errors in parentheses
*** p<0.01, **p<0.05, * p<0.1

B: Full-model: 2SLS Regression

Table 15. 2SLS Regression: Effect of Mental Health on Early Retirement

VARIABLES	Retired
Lagged mental health	0.0036 (0.0032)
Degree	-0.0394*** (0.0073)
Income	-0.0007*** (0.0001)
Household size	-0.0172*** (0.0030)
Married	0.0088 (0.0169)
Local unemployment rate	0.0008 (0.0024)
Age 52	-0.4115*** (0.0167)
Age 53	-0.3960*** (0.0173)
Age 54	-0.3924*** (0.0175)
Age 55	-0.3820*** (0.0162)
Age 56	-0.3422*** (0.0172)
Age 57	-0.3213*** (0.0159)
Age 58	-0.3061*** (0.0151)
Age 59	-0.2735*** (0.0155)
Age 60	-0.2255*** (0.0152)
Age 61	-0.1737*** (0.0152)
Age 62	-0.1354*** (0.0146)
Age 63	-0.1178*** (0.0148)
Age 64	-0.0687*** (0.0157)
Constant	0.3077 (0.2662)
Observations	29,432
R-squared (Centered)	0.0870
R-squared (Uncentered)	0.2659
Root MSE	0.3792
Underidentification test (LM stat.)	38.794
Underidentification p-value	0.0000
Weak identification test (CD F stat.)	38.819

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

C: Full-model: Complementary log-log model

Table 16. Complementary Log-Log Regression: Effect of Mental Health on Early Retirement

VARIABLES	Early Retirement (Hazard)
Lagged mental health	-0.0083*** (0.0028)
Initial mental health	-0.0076*** (0.0029)
Income	0.0006 (0.0005)
Female	0.0620 (0.0836)
Married	0.0613 (0.0989)
Degree	0.5969*** (0.1109)
Professional/Managerial Jobs	-1.9837*** (0.1484)
Negative Health Shock	-0.2812 (0.1941)
Household size	-0.1893*** (0.0544)
Local unemployment rate	-0.0112 (0.0390)
Age 53	-1.5787*** (0.4287)
Age 54	-1.7052*** (0.3763)
Age 55	-1.4337*** (0.2962)
Age 56	-0.4093** (0.1926)
Age 57	-0.6121*** (0.1992)
Age 58	-0.6215*** (0.1981)
Age 59	-0.5594*** (0.1898)
Age 60	-0.3353* (0.1769)
Age 61	-0.0721 (0.1680)
Age 62	-0.1441 (0.1715)
Age 63	-0.1332 (0.1805)
Age 64	-0.0783 (0.1857)
Constant	-0.3713 (0.3381)
Observations	9,054

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

D: Full-model: Second-Stage of 2SRI

Table 17. The second-stage of the Two-Stage Residual Inclusion

VARIABLES	Early Retirement
Lagged mental health	-0.011*** (0.003)
Initial mental health	-0.008*** (0.003)
Income	-0.001 (0.001)
Married	0.117 (0.100)
Degree	-0.199* (0.105)
Negative Health Shock	-0.175 (0.194)
Household size	-0.138** (0.055)
Local unemployment rate	0.149*** (0.041)
Age 53	-1.331*** (0.429)
Age 54	-1.589*** (0.376)
Age 55	-1.471*** (0.296)
Age 56	-0.480** (0.194)
Age 57	-0.732*** (0.201)
Age 58	-0.776*** (0.200)
Age 59	-0.724*** (0.191)
Age 60	-0.505*** (0.177)
Age 61	-0.234 (0.168)
Age 62	-0.304* (0.172)
Age 63	-0.417** (0.181)
Age 64	-0.319* (0.185)
Xuhat	0.052 (0.059)
Constant	-4.847 (4.558)
Observations	7,335

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

E: Full-model: Sub-sample of gender specific

Table 18. OLS estimates of early retirement on mental health and covariates, by gender

VARIABLES	Female	Male
Lagged mental health	-0.001*** (0.001)	-0.002*** (0.001)
Degree	-0.031*** (0.007)	-0.199* (0.105)
Income	-0.001*** (0.001)	-0.001*** (0.001)
Household size	-0.012*** (0.003)	-0.016*** (0.003)
Married	0.067*** (0.007)	-0.023*** (0.007)
Local unemployment rate	-0.012*** (0.003)	-0.007** (0.003)
Negative Health Shock	-0.032** (0.014)	-0.008 (0.013)
Age 53 to 64 Controls	Included	Included
Observations	16,192	15,865
R-squared	0.131	0.152

Notes: Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 19. 2SLS Estimates of Mental Health on Retirement, Instrumented by Death of Close Friend

	Dependent variable: Retired	
VARIABLES	Female (N=14,926)	Male (N=14,506)
Lagged mental health	0.0116 (0.0090)	-0.0004 (0.0034)
Degree	-0.0635** (0.0210)	-0.0185** (0.0073)
Income	-0.0009*** (0.0002)	-0.0006*** (0.0001)
Household size	-0.0040 (0.0082)	-0.0174*** (0.0033)
Married	0.0100 (0.0469)	-0.0241* (0.0140)
Local unemployment rate	-0.0013 (0.0041)	0.0024 (0.0030)
Negative Health Shock	0.0384 (0.0488)	0.0069 (0.0214)
Age 53 to 64 Controls	Included	Included
Constant	480	480
Observations	7,335	7,335
R-squared	335	335

Standard errors in parentheses
*p < 0.05, **p < 0.01, ***p < 0.001, p < 0.10

Table 20. Complementary Log-Log Regression on Retirement by Sex

	Dependent variable: Early retirement (d)	
VARIABLES	Female (N=14,926)	Male (N=14,506)
Lagged mental health	-0.0057 (0.0041)	-0.0147*** (0.0041)
Initial mental health	-0.0068 (0.0042)	-0.0080* (0.0041)
Income	-0.0006 (0.0008)	-0.0010 (0.0008)
Married	0.2223 (0.1428)	-0.1042 (0.1390)
Degree	-0.3473* (0.1651)	0.0174 (0.1380)
Negative health shock	-0.5775* (0.3400)	-0.0095 (0.2378)
Household size	-0.0259 (0.0843)	-0.2362*** (0.0734)
Local unemployment rate	0.0622 (0.0578)	-0.0477 (0.0537)
Age 53	-0.5912	-2.8383**
Age 54	-0.9307*	-2.2630***
Age 55	-0.7205	-1.8665***
Age 56	0.1301	-0.5563**
Age 57	0.2913	-1.1537***
Age 58	-0.0018	-0.8705***
Age 59	0.1140	-0.7936***
Age 60	0.5169	-0.7007***
Age 61	0.7893*	-0.4382**
Age 62	0.7452*	-0.5104**
Age 63	0.4823	-0.2896
Age 64	0.7171*	-0.2862
Constant	-2.2369*** (0.5493)	0.6060 (0.4526)
Number of obs	4,126	4,928
Log likelihood	-1020.27	-1131.84
LR chi2(20)	79.18	136.95
Prob > chi2	0.0000	0.0000

Standard errors in parentheses
*p < 0.05, **p < 0.01, ***p < 0.001, p < 0.10

Table 21. 2SRI Complementary Log-Log Regression on Retirement by Sex

VARIABLES	Female (N=3,320)	Male (N=4,015)
Lagged mental health	-0.0063 (0.0038)	-0.0164*** (0.0038)
Residual from first stage (Xuhat)	0.0072 (0.1205)	0.0688 (0.0670)
Initial mental health	-0.0080** (0.0040)	-0.0094** (0.0038)
Income	-0.0006 (0.0008)	-0.0012 (0.0009)
Married	0.2147 (0.1438)	-0.1149 (0.1395)
Degree	-0.3685** (0.1651)	-0.0101 (0.1383)
Negative health shock	-0.4792 (0.3403)	0.0398 (0.2381)
Household size	0.0090 (0.0857)	-0.2078*** (0.0733)
Local unemployment rate	0.2065*** (0.0603)	0.1027* (0.0558)
Age 53	-0.4963	-2.4146**
Age 54	-1.0005*	-1.9276***
Age 55	-0.9464**	-1.7801***
Age 56	-0.0708	-0.5921**
Age 57	-0.0556	-1.1802***
Age 58	-0.3461	-0.8813***
Age 59	-0.2607	-0.8656***
Age 60	0.1000	-0.7887***
Age 61	0.3809	-0.5138**
Age 62	0.3692	-0.6467***
Age 63	-0.0481	-0.4738**
Age 64	0.1268	-0.4338**
Constant	-2.9029 (9.1336)	0.6060 (0.4526)
Log likelihood	-947.359	-1059.165
Observations	3,320	4,015

Robust standard errors in parentheses
*p < 0.1, **p < 0.05, ***p < 0.01

Job satisfaction among public sector health employees: Gender, sexual identity, and ethnicity

KAREN MUMFORD *School for Business and Society, University of York, UK; IZA, Institute for the Study of Labour, Germany*
EDITH AGUIRRE *Institute for Social and Economic Research, University of Essex, UK*
ANNA EINARSDÓTTIR *School for Business and Society, University of York, UK*
BRIDGET LOCKYER *Bradford Institute for Health Research, UK*
MELISA SAYLI *School of Economics, University of Surrey, UK*
BENJAMIN A. SMITH *Townsville University Hospital and Health Service, Australia*

Abstract

We explore the determinants, and differences, in reported job satisfaction for women, ethnic minority and LGB+ employees among public sector health employees in the English National Health Service (NHS). A broad range of possible determinants are considered including demographic variables, job characteristics, and supportive workplace measures. Women are found to be more likely to be satisfied with their jobs, as are LGB+ employees from ethnic minorities. There is evidence that higher wage is positively associated with job satisfaction, but relative wages are not consistently related to job satisfaction. In contrast, supportive workplace practices are strongly associated with higher rates of job satisfaction. Of particular importance are effective workplace anti-bullying polices and the presence of relevant minority staff networks, especially for those identifying as LGB+. These results suggest that organisations can raise job satisfaction by further facilitating these supportive workplace practices.

Keywords: job satisfaction, LGB+, gender; ethnicity, networks, NHS, public health

Acknowledgements
The results in this article were first presented at the 33rd Australian Labour Market Research Workshop 2025, we thank the organisers and participants for this opportunity and their helpful comments. We are especially grateful to our Discussant Leonora Risse for her thoughtful and constructive insights. We also thank Yvonne Birks and Peter Smith; members of our Research Advisory Board for their support and guidance, particularly the Chair, Tracy Myhill, for her leadership and for paving the way for the project; and NHS Employers; Stonewall; and Employers Networks for Equality and Inclusion (enei); for supporting the development, progress and dissemination of the research. We gratefully acknowledge funding from the Economic and Social Research Council (ESRC) [grant number ES/N019334/1].

Corresponding author: Karen Mumford, karen.mumford@york.ac.uk, School of Business and Society, Church Lane Building, Church Lane, University of York, Heslington York YO10 5DG, UK.

Introduction

Job satisfaction amongst workers is strongly related to productive work behaviour, lower absenteeism and fewer quits (Yean *et al.*, 2023). There is a large and ongoing literature describing, and seeking to explain, job satisfaction (Freeman, 1978; Clark, 1997; Dolan *et al.*, 2008; Falk *et al.*, 2008; Clark *et al.*, 2009; Mumford and Smith, 2015; Green *et al.*, 2018). The main contribution we make to this research area is to provide a more detailed (and comparative) exploration of the determinants of job satisfaction for women, LGB+, and ethnic minorities.

Despite women facing extensive historical discrimination in the labour market and encountering sizeable gender pay gaps (Blau and Kahn, 2017; Goldin *et al.*, 2022), early studies typically found that women were more likely to be satisfied with their jobs than men (Clark, 1997). More recent studies suggest the gendered difference in job satisfaction is diminishing as women's expectations of work are changing to be closer to those of men (Green *et al.*, 2018).

There is considerable literature suggesting LGB+ employees are discriminated against in the labour market (Badgett, 1996; Badgett *et al.*, 2021) with lower rates of promotion and fewer management roles (Frank, 2006). The LGB+ are also more likely to seek employment in female dominated occupations associated with less productivity and lower pay (Plug *et al.*, 2014; Bridges and Mann, 2019). Nevertheless, empirical evidence of the relationship between sexual identity and job satisfaction is rare. Drydakis (2019a) presents survey results arguing that gay men experience lower levels of job satisfaction than male heterosexuals in Canada, Greece, Sweden and the US; and that lesbians have lower levels of job satisfaction relative to female heterosexuals in Canada, Greece and Sweden. Analogously, Fletcher *et al.* (2021) find LGBT+ workers in the UK are less likely to be satisfied with their job than heterosexuals (19 per cent compared to 15 per cent).

Sexual identity may be a hidden characteristic in the workplace if the employee chooses not to reveal (disclose) their preferences. LGB+ workers may choose to conceal their sexual identity if they fear stigmatism and discrimination (Goffman, 1963; Tajfel and Turner, 1979, page 281; Myer, 2013). Disclosure of sexual identity at work may, however, itself also impact on job satisfaction for LGB+ workers. Swann (2011) posits that people seek verification between their self-perception and the perceptions of others towards them. Disclosure can enable feelings of authenticity and improve social adhesion in the workplace; increasing identity pride, self-esteem, resilience, and mental health for LGB+ employees (Perrin *et al.*, 2020). It would appear that the relationship between disclosure of sexual identity and job satisfaction is ambiguous, and likely contingent on workplace conditions. Empirical evidence is again scarce, but evidence from Greece suggests gay men and lesbians who have disclosed their sexual identity in the workplace have higher job satisfaction than those who have not disclosed, although the size of these differences is small (Drydakis, 2015).

The pernicious nature of racial discrimination in workplaces is well documented in the literature (Heath and Di Stasio, 2019) but the direct empirical relationship between ethnicity and job satisfaction is likewise still under researched. In an early UK study, Shields and Wheatley Price (2002) found non-white nurses who frequently faced racial harassment from work colleagues were seven times more likely to report job dissatisfaction, and those facing racial harassment from patients were four times more likely. Doede (2017) found that black and hispanic nurses in the US were considerably less likely to be satisfied with their jobs than were white nurses. Across broader occupation comparison, however, Campbell (2011) found that within the United States, neither race nor ethnicity is a reliable predictor of workers' satisfaction with any specific element of a job.

Having multiple stigmatised characteristics may compound the relationship between minority status and job satisfaction. Early work on the impact of the intersection of characteristics (Crenshaw, 1989) focussed on the implications of dual stigmatised identities for black women. Subsequent work has broadened awareness across a range of minority characteristics including LGB+ status (Frost and Myer, 2023). Empirical implications of intersectionality between gender, race and sexual identity have not been established for job satisfaction. Related studies, for example of job insecurity (Lavaysse *et al*., 2018), suggest such intersections have complex impacts on employees and there may not necessarily be predictable associations with labour market outcomes (Raver and Nishi, 2010).

The risk (and experience) of negative outcomes in the work environment, including the occurrence and impact of discrimination, may be reduced by the presence of supportive workplace practices. Indeed, there is a very broad literature finding supportive work environments are positively associated with job satisfaction (Freeman, 1978; Shields and Ward, 2001; Huffman *et al*., 2008; Hebl *et al*., 2012; McFadden, 2015; Mumford and Sechel, 2019; Perales, 2022). Whilst the measures of workplace support vary across these studies, they generally include a range of indicators capturing social capital.

Established employee networks are an important component of workplace social capital. They facilitate a supportive work environment and enable transfer of social capital within minority groups in the workplace (Scrivens and Smith, 2013). Recent evidence shows that networks of members from the relevant minority group and their allies, are associated with greater appreciation and understanding of the minority group; a decrease in minority stressors; and a gain in wellbeing amongst the minority group (Perales, 2022). Network presence is expected to be positively associated with job satisfaction.

In contrast, the presence of bullying in the workplace can be considered as a failure of supportive workplace practices. There is considerable evidence showing that employees from minority groups – such as women (Salin and Hoel, 2013), the LGB (Hoel *et al*. 2022), and ethnic minorities (Lewis and Gunn, 2007) – are more likely to suffer from workplace bullying, and to encounter wellbeing losses from working in environments where bullying occurs (McFadden, 2015, page 142). There is also evidence that workplace bullying lowers job satisfaction for LGB+ employees (Drydakis, 2019b).

We seek to bring these strands in the literature together to provide a more detailed, and critically comparative, exploration of the determinants of job satisfaction for women, LGB+, and ethnic minorities The remainder of the paper is structured as follows: the data are described in section 2; methodology and the estimation of the determinants of job satisfaction are considered in section 3; results and discussions in section 4; and conclusions are presented in section 5.

Data and variables of interest

Data

To the best of our knowledge, the only data set that includes information on job satisfaction, pay, sexual orientation and disclosure, bullying and discrimination, and staff networks is the National Health Service Employee Engagement Survey (EES-NHS) in England. These employees are all covered by Agenda for Change contracts and the NHS Pay Review Board (NHSPRB)[1], which means doctors and dentists are excluded. The NHS is a particularly relevant workforce to survey as it is large enough to generate a suitably sized LGB+ sample for statistically meaningful analysis. Furthermore, NHS employees are all working in the public health sector where they share a common employer, with well recognised pay and working conditions. These commonalities help us to focus the empirical analysis presented below. Although, they may also limit extrapolation of the findings outside of the NHS to other more diverse workforces. This is a caveat that will be returned to in the discussion and interpretation of the results below.

The EES-NHS was launched in January 2019 and closed in May 2019; it is a fully pre-pandemic survey. The full sample taken from the EES-NHS includes 3,724 NHS employees. The NHS Digital's headcount data from September 2018[2] suggests that the potential sample frame was 1.19 million, implying a response rate of less than 1 per cent for the EES-NHS. In absolute terms, a sample of more than 3,700 employees is large enough for the development and testing of meaningful hypotheses. Nevertheless, such a low response rate raises obvious concerns that the sample may not reflect the population of NHS employees. Compared to the 2018 NHS-Staff Survey (NHS-SS), the EES-NHS sample has similar aggregate descriptive statistics including gender breakdown (with around 77 per cent women employees) and age distribution (Einarsdóttir *et al.*, 2020; Table 8). In terms of sexual orientation, however, the EES-NHS sample has a larger proportion of respondents declaring as LGB+ (12 per cent compared to 3.5 per cent) and fewer respondents opting for 'prefer not to say' (2.3 per cent relative to 6.5 per cent). This greater engagement is not surprising as LGBT+ labelling was included in the advertising for the EES-NHS survey. Comparing the EES-NHS with the National LGBT Survey suggests the former is broadly

1 https://www.gov.uk/government/organisations/nhs-pay-review-body
2 https://digital.nhs.uk/data-and-information/publications/statistical/nhs-workforce-statistics/september-2018

representative of the LGB+, given the focus of the former on employees in the NHS. For example, 78.2 per cent of respondents in the National LGBT Survey subsample were satisfied with their lives, and 70.9 per cent of the EES-NHS respondents are. Whilst 62.5 per cent of the National LGBT respondents said they were open about their sexuality with all or most co-workers, and 60.3 per cent did in the EES-NHS sample (see Einarsdóttir *et al.*, 2020, for more detailed comparison across these data sets).

Missing observations for variables used in the analysis below limits the usable sample from the EES-NHS to 3497 observations. As discussed above, one compensation for the overrepresentation of LGB+ employees in the EES-NHS sample is the inclusion of a reasonable number of observations in the analysis (435 LGB+).

Variable descriptions are provided in Table 1, further information regarding variable definitions and descriptive statistics is provided in Table A1 of the Appendix. Variable means are provided for the full sample (in column 1 of Table 1); for women (column 2); the LGB+ (column 3); and ethnic minorities (column 4). Statistically significant subgroup mean differences are indicated in Table 1 by bold font and a '+' (or '-') to show the minority value is higher (or lower) than the comparator group. More detailed information on the subgroup mean differences is provided in Table A2 of the Appendix.

Variables used in the analysis

Job satisfaction

The focus of this article is job satisfaction; it is the dependent variable throughout the analysis. The EES-NHS survey respondents were asked: "Overall how satisfied are you with your job these days?". There were five potential responses: extremely dissatisfied, somewhat dissatisfied, neither satisfied or dissatisfied, somewhat satisfied, or extremely satisfied. The responses are not evenly distanced from each other; even for a single individual the ranking is not necessarily cardinal. In aggregate, at best, the rankings may be considered as ordinal. Furthermore, the responses are not symmetric around the neutral mid-value (see Table 1), revealing that job satisfaction is not simply the inverse of job dissatisfaction and assuming that it could lead to incorrect interpretation of the results.

To avoid these possible measurement issues, a binary measure of job satisfaction is created and used as the dependent variable. It is set equal to one if the respondent is extremely or somewhat satisfied, and zero otherwise. Using this measure, some 53 per cent of the full sample report that they are satisfied with their job[3].

---

3 This average value is similar to the level of job satisfaction found amongst academic economists in the UK in 2016 (Mumford and Sechel, 2019). Focussing on sexual identity, Fletcher *et al.* (2021) found 19 per cent of the LGBT+ in Britain were dissatisfied, we found 25 per cent for this group in our sample.

Demographic variables

The gender of choice from the survey respondents is used in the EES-NHS, this is clearly relevant for transgender respondents.[4] There are 745 men and 2752 women in the full sample, with 2752 women in the woman only sample. Within group subsample differences (reported in the Appendix Table A2) reveals that 49 per cent of the men and 54 per cent of the women are satisfied with their job in our sample. These subgroup mean differences are statistically significant at standard confidence levels, as indicated by the bold formatting of the relevant means in Table A2. Survey respondents are also categorised as LGB+ in the EES-NHS according to their own choices. There are 435 LGB+ included in the sample, 296 gay or lesbian, 100 bisexual, and 39 in categories grouped together as plus. In our sample (on average) LGB+ employees are more likely to report job satisfaction (58 per cent) than heterosexuals (53 per cent). Those selecting as LGB+ were asked, 'What best describes how open you are about your sexuality/sexual orientation in your current job'. The possible responses are: give the heterosexual impression; not open at all; only reveal if asked; avoid drawing attention to it; make no secret about it; totally open. Those choosing that they make no secret about their sexual identity, or that they are totally open about it, are classified as disclosed. Just over half of the LGB+ employee sample have disclosed (51 per cent), with this disclosure more common amongst men (60 per cent) than women (42 per cent).

An ethnicity measure is also included and set equal to one if survey respondents selected Asian; Black; Arab; other Non-White; or Mixed to the question 'What is your ethnic background?' On average, ethnicity is associated with equal lowest job satisfaction (along with men) at 49 per cent compared to 54 per cent for the non-ethnic. This subgroup mean difference is again statistically significant at standard confidence levels (as indicated by bold font in Table A2).

The remaining demographic measures (age, having dependent children, living with partner, having a disability, and highest acquired education) may be considered primarily as control variables in the analysis, although they are also of interest in their own right. Mild associations have been found with age and job satisfaction, with younger woman workers being less satisfied with their jobs and older workers being more satisfied (O'Brien and Dowling, 2011). Clark (1997) finds that being married is positively associated with job satisfaction for women but not for men, and Guler (2024) finds that having dependent children, especially preschool children, increases job dissatisfaction for women but not for men. The majority of the work investigating the relationship between disability and job satisfaction has focused on alternative modes of employment. Recent findings, however, suggest the relationship between disability and job satisfaction for employees (as compared to the self-employed) is negative but small (Keating *et al.*,

4 There are 17 self-identifying transgender individuals included in the respondents, removing the transgender individuals from the sample does not change the findings in any substantial or statistically significant manner. Results without the transgender respondents are available from the authors upon request.

2022). The relationship between education and job satisfaction has been found to be weakly negative but becoming insignificant when additional indications of job stress are included amongst the regressors (Solomon *et al.*, 2022).

Job characteristics

Higher wages are commonly believed to help compensate for negative aspects related to working thereby increasing job satisfaction (Lavetti, 2023). The NHS employees are paid in twelve bands (or sub-bands) set by the government with advice from the NHS Pay Review Board (NHSPRB)[5] and other parties. The average hourly own-wage measure is used below and is constructed from the mid-point of the employee's salary band, allowing for their usual working hours, adjusting for overtime hours (paid and unpaid), and expressed in natural logarithms. This hourly own wage (W) for individual i is denoted as W_i.

Workers may also care about their own wage relative to the wage of other comparable employees. They may gain utility (and be more likely to be satisfied with their job) from having an own wage that is relatively higher, or disutility (and greater likelihood of job dissatisfaction) from being relatively lower paid (Card *et al.*, 2012). An alternative argument is that workers view their co-workers wage as a signal of their own future wage growth (Clark *et al.*, 2009); seeing comparable workers earning higher wages would encourage them to think that their own wage will also increase in the future. These positive expectations would increase job satisfaction. Gender may also be important, Mumford and Smith (2015) find that higher own wage, and higher relative wage, is associated with higher job satisfaction for British men. In contrast, British women appear to care only about their own wage, with higher own wage increasing job satisfaction. In order to address the possible relationship between relative wage and job satisfaction, we include a relative wage measure (RW_i) equal to the difference in hourly own-wage (W_i see above) and the mean average hourly wage in the particular workers own occupation and age band (AW_{oa}), all divided by the mean hourly wage in the particular workers own occupation and age band (AW_{oa}) : $RW_i = (W_i - AW_{oa}) / AW_{oa}$.

The remaining job characteristics include working part-time, having a permanent contract, having training opportunities, recent promotion, being a trade union member, often feeling under pressure, feeling overwhelmed at work, wanting to work less hours, and being able to maintain work-life balance. Liu and Zhang (2014) found part-time employees had lower levels of job satisfaction, especially if they worked more hours than they wanted to. We might expect that having a permanent contract, training opportunities and recent promotion all reflect employer approval, stronger job matching, and a higher likelihood of job satisfaction. The relationship between trade union membership and job satisfaction is not obvious. Being prepared to pay trade union membership fees may reflect poor working conditions and job dissatisfaction or a fear

5 https://www.gov.uk/government/organisations/nhs-pay-review-body

of future job dissatisfaction (Bessa *et al.*, 2021). Trade unions may, however, provide a range of services for their members, including a voice mechanism for the individual, that increasing job satisfaction (Freeman, 1978). Blanchflower and Bryson (2022) recently find trade union membership is associated with lower job satisfaction, although this relationship is small and not found for newer union members. Of the remaining job characteristics, feeling under pressure, feeling overwhelmed at work, and not being able to maintain work-life balance are all indications of employee stress; and are predicted to be negatively related to job satisfaction.

Measures of workplace support

The supportive workplace practice measures included in the analysis are: having supportive coworkers; being in a cooperative workplace; having a responsive supervisor; having an effective mentor or coach; feeling part of the family at their organisation; having influence over their job; and satisfaction with the quality of care they provide. (Additional definition of these variables can be found in Table A1 of the Appendix.) As discussed above, supportive workplace practices help to diminish the incidence, and the extent, of negative outcomes for employees in the work environment. We expect supportive work practices to be positively associated with job satisfaction (Freeman, 1978; Shields and Ward, 2001; Huffman *et al.*, 2008; Hebl *et al.*, 2012; McFadden, 2015; Mumford and Sechel, 2019; Perales, 2022). Notably, ethnic minorities in our sample report the lowest average values for all of these supportive workplace practices (see Table 1).

NHS workers are located in organisational units called Trusts, there are 203 individual Trusts included in the analysis. Trusts are grouped into several types by the NHS according to primary function and medical specialty (see Appendix Table A1), they are also categorised by region. We include a measure of relevant staff network existence in the Trust: the existence of a women's network in the Trust when considering the women only sample; LGBT+ network existence for the LGB+ sample; and BAME network existence for the ethnic minority sample. When considering the total sample, the existence of any of these minority group staff networks in the worker's Trust is included. As discussed above, employee networks can facilitate a supportive work environment and ease the transfer of social capital within minority groups in the workplace. Recent evidence shows that networks of members from the relevant minority group, and their allies, are associated with greater appreciation and understanding of the minority group; a decrease in minority stressors; and a gain in wellbeing amongst the minority group (Perales, 2022). Network presence is expected to be positively associated with job satisfaction.

The final measures of supportive workplace practices address bullying and discrimination. There is considerable evidence showing that employees from minority groups are more likely to experience bullying and suffer hardship from working in environments where bullying occurs (McFadden, 2015, page 142). We expect workplace bullying to be associated with lower levels of job satisfaction. We include two measures to help reflect this possibility at (i) the individual level and (ii) the broader environment context. In recognition that the understanding of what constitutes bullying may not be uniform, the EES-NHS questions on bullying were preempted with the statement: “Bullying at work involves repeated negative actions and practices that are directed at one or more workers/employees. The behaviours are unwelcome to the victim and undertaken in circumstances where the victim has difficulty defending themselves. We do not think of one-off incidents as bullying”.

At the individual level we create a bullying index ranging from 0 to 3: coded with one point for each of: having been “bullied at work in the last six months” (some 22 per cent have); “witnessed bullying at work” (51 per cent have); or “been subject to discrimination in the last 12 months” (17 per cent have). On average, less than a third of workers have had encountered one type of these negative experiences recently, 37 per cent have had no experience and 35 per cent have encountered two or more. Reflecting the broader work environment, a categorical variable capturing “do you think the measures your organisation takes to prevent bullying or discrimination are effective” is included, ranging from not effective at all (22 per cent), slightly effective (17 per cent), moderately effective (37 per cent), to very or extremely effective (21 per cent). Ethnic minorities are more likely to experience bullying or discrimination personally, and more likely to rate their organisation’s measures to prevent this behaviour as ineffective. We expect job satisfaction will be negatively related to higher bullying index outcomes and to less effective bullying reduction measures.

Trust characteristics

The Trust controls included in the analysis are variables that are common to all workers in that Trust that can vary across Trusts; they are measures of regional location and Trust type. The Trust measures are primarily included as control variables in the analysis. There are five regions: the north of England; Midlands and the east of England; London; the south-west; and the south-east. Compared to heterosexuals (see Appendix Table A2), LGB+ employees are substantially more likely to be located in London (24 per cent compared to 14 per cent) or, to a lesser extent, the south-east; whereas they are less likely to work in the Midlands or east of England, or the south-west. The concentration in London is even more extreme for ethnic minorities, 46 per cent work in London compared to 11 per cent of the white respondents; whilst they are less likely to be located in the north of England, the Midlands or the east of England, or the south-west. The trends for women compared to men are quite different. Women are slightly less likely to work in London (15 per cent for women, 18 per cent for men) and the south-east, and they are more likely to be located in the Midlands or the east of England. There are seven Trust types included, of particular note is the concentration of LGB+ working in the Ambulance Trusts and the prevalence of ethnic minorities employed in the Mental Health Trusts.

Summary of cross group differences

Considering the information presented in Table 1 in more detail. The NHS is a female dominated work environment: some 79 per cent of the workforce are women (see Table

1). Compared to the men, the women in our sample are on average less likely to identify as LGB+ or ethnic, they are older, earn less, are considerably more likely to work part-time, be satisfied with their training, belong to a trade union, and feel overwhelmed at work. The women are also more likely to have a mentor, feel part of the work family, have influence over their work, and be satisfied with the quality of care provided. However, the women are less likely to have a staff network, and they have (perhaps surprising) experienced less workplace bullying than the men (see Table 1, column 2).

LGB+ employees in our sample are on average younger than heterosexuals, they are less likely to be a woman, be living with their partner, or have dependent children. They are less likely to work part-time, and are more likely to have been recently promoted, work in cooperative workplaces and to have a staff network. On average, ethnic minorities are younger than whites, they have more dependent children, are less likely to live with their partner and are more likely to be graduates. Ethnic minorities tend to earn more, are less likely to work part-time and more likely to feel overwhelmed at work. They are, on average, less likely to report having supportive workplace practices in their work environment, except for having a staff network. Ethnic minorities also consistently report experiencing more bullying and that their workplace measures to prevent bullying are not effective. We next consider the formal estimation of the determinants of job satisfaction.

Methodology and estimation

We estimate the probability that individual *i* is satisfied with their job (S_i) conditional on a range of observable characteristics expected to predict that probability. Probit regressions are estimated for the total sample, and for the subsamples of interest, with the latent dependent variable (job satisfaction, S_i) set equal to 1 if the individual responds they are somewhat satisfied or extremely satisfied with their job, and zero otherwise. The probit models the relationship between the probability of satisfaction and its determinants as

$$Pr\,(S_i = 1) = \theta(\beta X_i) \qquad (1)$$

where X_i is a vector of explanatory variables and θ is the standard normal distribution function (Maddala 1992; 327).

We begin with a parsimonious model which includes the demographic and job characteristics but excludes the workplace support measures (all discussed in Section 2 above). Selected estimation results are presented in Table 2. We next include the workplace support measures in our preferred 'full' model, with selected results presented in Table 3 (complete results are provided in Table A3 of the Appendix). The more intuitive marginal effects at the means of the explanatory variables are reported in Tables 2 and 3 with differential effects for the binary variables. Results for the total sample are reported in column 1 of Tables 2 and 3. While, columns 2 to 4 of Tables 2 and 3 provide results for the subsamples of interest: women (column 2), the LGB+ (column 3), and ethnic minorities (column 4). With a single cross section of data, such as the EES-NHS, the results should not be interpreted as causal; they indicate the direction and strength of the relationship between each determinant variable and job satisfaction.

Results and discussion

Considering the results for the parsimonious model (Table 2) in more detail, amongst the demographic variables, women are 6.4 per cent more likely to be satisfied with their job in the total sample results (see column 1 of Table 2); the size of this gap is arguably not large compared to historical studies (Green *et al.*, 2018), but it is clearly statistically significant. This result is consistent with findings from previous studies that women are more satisfied with their jobs than men, but this gender gap has become smaller over time as women's expectations of work move closer to those of men (Clark, 1997; Mumford and Smith, 2015; Green *et al.*, 2018). Comparing columns 1 and 2 of Table 2 reveals that the other results are typically not statistically significantly different between the total sample (column 1) and the woman only sample (column 2); this finding may further support the claim that the gender gap in job satisfaction is related to differences in expectations rather than characteristics.

Being either an ethnic minority or LGB+ employee is not found to be statistically related to job satisfaction in the total sample estimations (column 1). However, ethnic LGB+ employees are 20.1 per cent more likely to be satisfied than heterosexual ethnic employees (column 4), and ethnic LGB+ employees are 13 per cent more likely to be satisfied than white LGB+ employees (column 3); suggesting that this intersection effect may be important and supporting separate estimation for these major demographic groups. This point will be returned to in discussions of the full model below. The relationship between disclosure of sexual identity and job satisfaction is small, negative and not statistically significant. In net terms, this finding does not support the argument that positive aspects related to disclosure outweigh the stressors associated with stigmatism for this minority group in our sample. Age appears to have a negative but very small relationship with job satisfaction. There is also evidence that women and ethnic minorities who are living with their partner are more likely to report job satisfaction.

Amongst the job characteristics, we find some evidence supporting the standard utility wage model. In particular, employees are more likely to be satisfied with their jobs if they have higher own earnings. This relationship is not, however, statistically significant for either LGB+ or ethnic minorities. The relationship with relative wage is smaller, inconsistent, and is only statistically significantly different to zero for women. Taken as a whole, this may suggest that the minority groups are not competing with their comparators over relative wage (Mumford and Smith, 2015). Having a permanent job is associated with less job satisfaction for ethnic minorities, as is working part-time. Furthermore, and consistent with the priors, being satisfied with training, being recently promoted, maintaining work-life balance or not feeling overwhelmed with work are positively associated with job satisfaction for all groups.

Introducing the measures of supportive workplace practices, discussed in section 2 above, into the model (see the full model results in Table 3) generally leads to (statistically insignificant) decreases in the estimated marginal effects suggesting, unsurprisingly, some correlation between some demographic groups (especially women), job characteristics (such as having training, recent promotion, having work-life balance and feeling overwhelmed) and the measures of supportive workplace practices. Nevertheless, including measures of supportive workplace practices increases the overall model fit considerably (for example, comparing the pseudo R-square of 0.2726 for the parsimonious model in column 1 of Table 2, with the pseudo R-square of 0.4407 for the full model in column 1 of Table 3) revealing that, in aggregate, the supportive workplace measures are contributing additionally and substantially to the explanation of job satisfaction.

Considering the additional results in the full model in Table 3 in more detail, the measures of supportive workplace practices are all associated with women being more likely to report job satisfaction (see Table 3), especially strong is the effect of feeling part of the family at work; working in a cooperative workplace; and having influence over job. Having a network staff network in the Trust is also important. Women in Trusts with a staff gender network are 6.2 per cent more likely to report job satisfaction. The results for LGB+ employees are similar although having supportive colleagues is not as important whereas working in a cooperative workplace, having a responsive supervisor, having a mentor and not feeling overwhelmed are all more strongly related to higher job satisfaction. Furthermore, LGB+ workers in Trusts with a LGBT+ staff network are 13.6 per cent more likely to report satisfaction, twice the likelihood found for women of 6.2 per cent. This is the strongest effect found for any job or workplace determinant of job satisfaction for the LGB+ sample, revealing a strong policy role for organisations to facilitate the provision of minority staff networks.

In contrast, there is an insubstantial and insignificant relationship found between the presence of a BAME staff network in the worker's Trust and job satisfaction for the ethnic minority; despite this group being the most likely to say that a staff network was available in their Trust (see Table 1). This result may be, at least partially, due to heterogeneity amongst the ethnic minorities not being reflected in the BAME staff network that is available at their Trust.[6] The only supportive workplace practice measures found to be significantly related to job satisfaction for ethnic minorities are having a mentor,

6 This is also true if relevant staff network membership, rather than network existence, is considered. Ethnic minorities who are members of an ethnic minority staff network at their Trust are 2 per cent more likely to report job satisfaction, however, this result is not statistically significant at standard confidence levels. Full results are available upon request from the authors.

feeling part of the family, and having job influence. These relationships are all strong but, taken as a whole, these results suggest that ethnic minorities are not as included, nor as supported, in their work environments as are the other minority employees.

Finally, believing the organisation has effective measures to prevent bullying and or discrimination are broadly associated with more reported job satisfaction for women. For ethnic minorities and LGB+ employees this relationship is only statistically significant when the organisation's prevention policies are considered to be very or extremely effective. While experiencing bullying and/or discrimination is associated with all worker groups being less likely to report job satisfaction (with a similar sized, but not statistically significant, relationship for ethnic minorities).[7]

There is some evidence in both the parsimonious (Table 2) and the full model (Table 3) results that intersection between LGB+ and ethnic minority status is associated with higher job satisfaction. To further explore this, as well as address possible intersection effects with gender, three two-way interactive terms (gender*LGB+, gender*ethnicity, ethnicity*LGB+) were included in the full model (see column 2 of Appendix Table A4); and the three-way interactive term (gender*LGB+*ethnicity) also additionally included (see column 3 of Appendix Table A4). While further evidence of a statistically significant intersection effect between LGB+ status and being an ethnic minority is found (see coefficient estimates in Appendix Table A5 and also the table notes for Appendix Table A4), including this term and/or the other interactive terms is not found to change the results either qualitatively or quantitatively (as shown by reading across the columns in Table A4).

Using the EES-NHS data allows us to include a broad range of explanatory variables in the analysis, this helps to control for omitted variable bias and possible endogeneity. However, there are also limitations with the data which may influence interpretation of the findings. For example, all of the employees in the sample have accepted jobs in the public health sector (the NHS), this is a possible self-selection effect in the analysis for which we do not have a suitable control measure. It may be that (at least some) individuals who select into employment in this sector have a vocation or an intrinsic motivation to provide care (Heyes, 2005). These individuals may have stronger preferences (and higher job satisfaction) associated with inclusion, equity and collegial support than those found in other sectors. In aggregate, these employees might also be less motivated by pay, especially relative pay. Not being able to control for this possible self-selection effect limits extrapolation of the findings to other sectors. Furthermore, while the EES-NHS sample size is large enough for statistically meaningful analysis, the survey response rate is low and the sampling process is not fully random, both of which may limit the extrapolation of the findings across the full NHS workforce. The goodness of fit measures presented in Table 3 suggest the full model estimations are reasonable

7 We also include results for the subsample of nurses for all descriptive statistics and estimations in the Appendix. Insubstantial, and generally statistically insignificant, differences were found comparing nurses with the woman subsample; the analysis presented in this article focuses on women accordingly.

but are only capturing some 44 per cent of the total variation in job satisfaction. There are clearly other factors affecting job satisfaction that are not included in our modelling.

In an ideal scenario, equivalent data sets to the EES-NHS would exist for other countries, including Australia, enabling critical international modelling of job satisfaction. To the best of our knowledge, however, no comparable international data exists. We might expect our results from the NHS to have some relevance to the Australian public health sector. Both sectors share similar occupation definitions and training requirements, work task allocation, patient health requirements, and a reliance on immigrant labour. However, the private health sector is relatively larger in Australia, and the public sector is more highly funded (per population head) across a range of measures, resulting in substantially better patient outcomes in Australia than England (Anandaciva, 2023) and higher associated job satisfaction levels (Dattoli and Cohidon, 2025).

To reiterate, the EES-NHS survey provides only a single cross section of data for public health sector workers in England. It allows for an exceptionally broad range of explanatory variables, as in our analysis, which could reasonably be expected to capture some potential omitted variable bias. Nevertheless, low survey response rates and the focus of the data exclusively on the English public health sector limits the generality of the results. Furthermore, it is not possible to identify causality with cross-sectional data such as that used in this article, however, the results do show clear observational relationships.

Conclusion

This study explores job satisfaction amongst a sample of National Health Service (NHS) employees in England. This is a particularly relevant workforce to consider as it is large enough to generate suitable LGB+ and ethnic samples for statistically meaningful analysis. Furthermore, the respondents are all working in the public health sector where they share a common employer, with well recognised pay and working conditions.

Women are found to be more likely to be satisfied with their jobs, however, this gender gap is not large. Taken as a whole, the results are in accordance with other recent findings suggesting women's job satisfaction levels are becoming similar to men's as women's job expectations are moving towards those of men. LGB+ employees who are an ethnic minority (or ethnic minorities who identify as LGB+) are also found to be more likely to be satisfied with their current job. In contrast to our predictions, workers who have disclosed minority sexual identity to their workplace colleagues are not found to have higher levels of job satisfaction. The determinants of the decision to disclose sexual identity to work colleagues is under researched; our results suggest that this is clearly a potentially fruitful area for future research.

Unsurprisingly, higher own-wage is associated with higher job satisfaction for the NHS workers in the sample, consistent with standard labour economic theories of wage determination. The finding of no clear pattern between relative wages and job satisfaction indicates that these workers are not competing about their relative wages. This muted response to the pay of comparator colleagues may be influenced by the cross-occupation job-team tasking commonly used in public health sector workplaces.

Supportive workplace practices are strongly related to higher rates of job satisfaction, less so for the ethnic minorities. In common across all the groups, experiencing bullying or discrimination in the workplace is associated with decreased likelihood of job satisfaction. Furthermore, simply having measures to prevent bullying and/or discrimination is not enough to improve job satisfaction. It is only when these measures are considered to be very or extremely effective that ethnic minority or LGB+ workers report significantly more job satisfaction. These results have obvious implications for policy implementation for health sector managers. Of particular importance amongst the supportive workplace measures is the existence of relevant staff networks. Women are 6.2 per cent more likely to report job satisfaction when they have a staff network; LGB+ employees are 13 per cent more likely to. This is the strongest effect found for any determinant of job satisfaction for those identifying as LGB+; revealing a further policy role for organisations to facilitate the provision of minority staff networks.

A cautionary finding is that the results reveal low job satisfaction, low inclusion, and low levels of workplace support for the ethnic minorities. This is a concerning finding for a sizable proportion of the workers in the sample and suggests that the provision of supportive workplace practices is not extended fully across all minority groups.

The conclusions are valid for the sample of NHS employees addressed in this study. The sampling process was not random and, while the sample sizes are reasonable, they are small relative to the sampling population. Both factors limit extrapolation of the findings across the NHS or to a broader social context and suggest a need for additional studies. It is also not possible to identify causality between job satisfaction and the explanatory variables used in this paper, however, the results do show clear observational relationships.

Nevertheless, this article provides important new evidence, revealing limited job satisfaction amongst women, ethnic minority, and LGB+ employees. By encouraging supportive workplace practices across all employees, having effective anti-bullying and discrimination policies, and (especially) facilitating minority staff networks, managers can further improve these levels of job satisfaction.

References

Anandaciva, S. (2023) How does the NHS compare to the health care systems of other countries? Kings Fund Report. https://www.kingsfund.org.uk/insight-and-analysis/reports/nhs-compare-health-care-systems-other-countries Accessed on 12/11/2025.

Badgett, M. V. L. (1996) Employment and sexual orientation: Disclosure and discrimination in the workplace, in A. L. Ellis & E. D. B. Riggle (Eds.), *Sexual Identity on the Job: Issues and Services*. Harrington Park Press/Haworth Press, pages 29–52.

Badgett, M.V.L., Carpenter, C.S. and Sasone, D. (2021) LGBTQ Economics. *Journal of Economic Perspectives* 35(2): 141-171.

Bessa, I., Charlwood, A. and Valizade, D. (2021) Do unions cause job dissatisfaction? Evidence from a quasi-experiment in the United Kingdom. *British Journal of Industrial Relations* 59(2): 251-278.

Blanchflower, D. and Bryson, A. (2022) Union membership and job satisfaction over the life course. *Industrial Relations Journal* 53(5): 411-429.

Blau, F. and Kahn, L. (2017) The gender wage gap: Extent, trends, and sources. *Journal of Economic Literature* 55(3): 789-865.

Bridges, S. and Mann, S. (2019) Sexual orientation, legal partnerships and wages in Britain. *Work, Employment and Society* 33 (6): 1020–1038.

Campbell, D.G. (2011) Diversity and job satisfaction: Reconciling conflicting theories and findings. *International Journal of Applied Management and Technology* 10(1): 1–15.

Card, D., Mas, A., Moretti, E. and Saez, E. (2012) Inequality at work: The effect of peer salaries on job satisfaction. *American Economic Review* 102(6): 2981-3003.

Clark, A. (1997) Job satisfaction and gender: Why are women so happy at work? *Labour Economics* 4: 341-372.

Clark, A., Kristensen, N. and Westergard-Nielsen, N. (2009) Job satisfaction and co-workers wages: Status or signal. *Economic Journal* 119: 430-447.

Crenshaw, K.W. (1989) Demarginalizing the intersection of race and sex: A black feminist critique of antidiscrimination doctrine, feminist theory, and antiracist politics. *University of Chicago Legal Forum* 1989: 139–167.

Dattoli, E. and Cohidon, C. (2025) Primary care physicians' job satisfaction in eleven Western countries: A cross-sectional study. *BMC Primary Care* 26, 241. https://doi.org/10.1186/s12875-025-02949-0

Doede, M. (2017) Race as a predictor of job satisfaction and turnover in US nurses. *Journal of Nursing Management* 25: 207–214.

Dolan, P., Peasgood, T. and White, M. (2008) Do we really know what makes us happy: A review of the economic literature on the factors associated with subjective well-being. *Journal of Economic Psychology* 29(1): 94-122.

Drydakis N (2015) Effect of sexual orientation on job satisfaction: Evidence from Greece. *Industrial Relations* 54(1): 162-187.

Drydakis, N. (2019a) Sexual orientation and labour market outcomes. *IZA World of Labor*, July 2019.

Drydakis, N. (2019b) School-age bullying, workplace bullying and job satisfaction: Experiences of LGB people in Britain. *The Manchester School* 87 (4): 455-488.

Einarsdóttir, A., Mumford, K., Birks, Y., Aguirre, E., Lockyer, B. and Sayli, M. (2020) LGBT employee networks within the NHS: Technical report and data addendum. York: University of York.

Falk, A., Fehr, E. and Fischbacher, U. (2008) Testing theories of fairness: intentions matter. *Games and Economic Behavior* 62: 287-303.

Frank, J. (2006) Gay glass ceilings. *Economica* 73: 485-508.

Fletcher, L., Miller, J. and Green, M. (2021) Inclusion at work: Perspectives on LGBT+ working lives. *CIPD Research Report*: London.

Freeman, R. (1978) Job satisfaction as an economic variable. *American.Economic. Review* 68(2): 135-141.

Frost, D. and Meyer, I. (2023) Minority stress theory: Application, critique, and continued relevance. *Current Opinion in Psychology* 51: 101579.

Goldin, C., Kerr, S. and Olivetti, C. (2022) When the kids grow up: Women's employment and earnings across the family lifecycle. *VOX Eu CEPR*. 2022; Oct. 7.

Government Equalities Office GEO. (2018) National LGBT Survey Research Report, GEO-RR001 (ISBN: 978-1-78655-661-5), Manchester.

Goffman, E. (1963) *Stigma: Notes on the Management of Spoiled Identity*. Simon & Shuster: New York.

Green, C., Heywood, J., Kier, P. and Leeves, G. (2018) Paradox lost: The disappearing female job satisfaction premium. *British Journal of Industrial Relations* 56(3): 484-502.

Guler, A. (2024) Gender, children, and employment: A study of the effects of children on job satisfaction. *Academy of Management Annual Meeting Proceedings* 2024(1), forthcoming.

Heath, A. and Di Stasio, V. (2019) Racial discrimination in Britain, 1969–2017: A meta-analysis of field experiments on racial discrimination in the British labour market. *British Journal of Sociology* 70(5): 1774-1798.

Hebl, M., Tonidandel, S. and Ruggs, E. (2012) The impact of like-mentors for gay/lesbian employees. *Human Resources* 25: 52-71.

Heyes, A. (2005) The economics of vocation or why is a badly paid nurse a good nurse? *Journal of Health Economics* 24: 561-569.

Hoel, H., Lewis, D., Einarsdóttir,A. and Notelaers, G. (2022) Openness about sexual orientation and exposure to workplace bullying. *Scandinavian Journal of Work and Organizational Psychology* 7(1): 9, 1–17. DOI: https://doi.org/10.16993/sjwop.164

Huffman, A., Watrous-Rodriguez, K. and King, E. B. (2008) Supporting a diverse workforce: What type of support is most meaningful for lesbian and gay employees? *Human Resource Management* 47 (2): 237-253.

Keating, B., Keramat, S., Waller, I. and Hashmi, R. (2022) Job satisfaction dynamics: How do they impact employment mode choice for people with a disability? *Helyon* 8(10): 1-9.

Lavaysse, L., Probst, T. and Arena, D. (2018) Is more always merrier? Intersectionality as an antecedent of job insecurity. *International Journal of Environmental Research and Public Health* 15(11):2559. doi: 10.3390/ijerph15112559

Lavetti, K. (2023) Compensating wage differentials in labor markets: Empirical challenges and applications. *Journal of Economic Perspectives* 37 (3): 189–212.

Lewis, D. and Gunn, R. W. (2007) Workplace bullying in the public sector: Understanding the racial dimension. *Public Administration and International Quarterly* 83, 641–665.

Liu, X. and Zhang, L. (2014) Examining the relationship among part-time work arrangements, job satisfaction, and work effort. *Advances in Industrial and Labor Relations* 21: 87-117.

Maddala, G.S. (1992) *Introduction to Econometrics*. Second Edition. Prentice Hall: New Jersey.

McFadden, C. (2015) Lesbian, gay, bisexual and transgender careers and human resource development: A systematic literature review. *Human Resource Development Review* 14(2): 125-162.

Mumford, K. and Sechel, C. (2019) Job satisfaction amongst academic economists in the UK. *Economic Letters* 182: 55-58.

Mumford, K.A. and Smith, P.N. (2015) Peer salaries and employee satisfaction in the workplace. *Manchester School* 83 (3): 307-313.

Meyer, I. (2013) Prejudice, social stress, and mental health in lesbian, gay and bisexual populations: conceptual issues and research evidence. *Psychology Bulletin* 129: 674-697.

O'Brien, G. and Dowling, P. (2011) Age and job satisfaction. *Australian Psychologist* 16: 49-61.

NHS Digital. (2018) Electronic Trust Record (etr), Electronic Dataset, NHS Digital. Retrieved from https://digital.nhs.uk/services/organisation-data-service/data-downloads/other-nhs-organisations.

Perales, F. (2022) Improving the wellbeing of LGBTQ+ employees: Do workplace diversity training and ally networks make a difference? *Preventive Medicine* 161: 1-8.

Perrin, P., Sutter, M., Trujillo, M., Henry, R. and Pugh J (2020) The minority strengths model: Development and initial path analytical validation in racially/ethnically diverse LGBTQ+ individuals. *Journal of Clinical Psychology* 76: 118-136.

Raver, J.L. and Nishii L.H. (2010) Once, twice, or three times as harmful? Ethnic harassment, gender harassment, and generalized workplace harassment. *Journal of Applied Psychology* 95: 236–254.

Salin, D. and Hoel, H. (2013) Workplace bullying as a gendered phenomenon. *Journal of Managerial Psychology* 28, 235–51.

Shields, M. and Ward, M. (2001) Improving nurse retention in the National Health Service in England: The impact of job satisfaction on intentions to quit. *Journal of Health Economics* 20: 677-701.

Shields, M. and Wheatley Price, S. (2002) Racial harassment, job satisfaction and intentions to quit: Evidence from the British nursing profession. *Economica* 69(274): 295-326.

Solomon, B., Nikolaev, B. and Shepherd, D. (2022) Does educational attainment promote job satisfaction? The bittersweet trade-offs between job resources, demands, and stress. *Journal of Applied Psychology* 107(7):1227-1241.

Swann, W. (2011) Self-verification theory. *Handbook of Theories of Social Psychology* 2: 23-42.

Tajfel, H. and Turner, J. (1979) An integrative theory of intergroup conflict. In Austin W, Worchel S, (eds) *The Social Psychology of Intergroup Relations*, pages 33-47. (Monterey, CA: Wadsworth)

Yean, T., Johari, J., Yahya, K. and Chin, T. (2022) Determinants of job dissatisfaction and its impact on the counterproductive work behavior of university staff. *Sage Open* 12(3): 1–11. https://doi.org/10.1177/21582440221123289

Table 1. Variable means

	(1) Total	(2) Women	(3) LGB+	(4) Ethnic
Job satisfaction				
extremely dissatisfied	0.08	***0.08-***	***0.05-***	***0.13+***
somewhat dissatisfied	0.22	0.22	0.20	0.19
neither satisfied nor dissatisfied	0.15	**0.14-**	0.15	**0.18+**
somewhat satisfied	0.41	**0.42+**	0.44	0.37
extremely satisfied	0.11	0.11	0.12	0.10
job satisfaction binary	0.53	***0.54+***	**0.58+**	***0.49-***
Demographics				
woman	0.79		***0.52-***	***0.72-***
LGB+	0.12	***0.08-***		0.10
disclose	0.06	***0.03-***	0.51	**0.04-**
ethnic	0.12	***0.11-***	0.09	
age	46.14	***46.45+***	***41.32-***	**45.09-**
dependent children	0.32	0.32	***0.15-***	***0.41+***
living together	0.69	0.68	***0.57-***	***0.62-***
disability	0.36	0.35	***0.45+***	***0.29-***
Qualifications, highest acquired				
minimal	0.05	0.05	***0.02-***	***0.02-***
GCSE grades A-C	0.08	**0.09+**	***0.05-***	***0.05-***
Trade	0.004	***0.002-***	0.002	0.002
A levels	0.09	0.09	0.10	***0.05-***
Diploma	0.16	0.16	0.16	***0.12-***
First degree	0.30	0.30	0.31	***0.34+***
Higher degree	0.28	0.28	***0.32***	***0.37+***
Job characteristics				
own-wage	2.74	***2.73-***	2.75	***2.81+***
relative-wage	0.001	***-0.01-***	**0.02+**	***0.05+***
part-time	0.24	***0.28+***	***0.11-***	***0.15-***
job permanent	0.93	0.93	0.93	0.91
training	0.52	***0.54+***	0.53	0.50
promotions	0.36	0.36	***0.42+***	0.38
trade union	0.57	***0.58+***	0.57	0.58
pressure	0.56	0.56	0.54	0.59
overwhelmed	0.38	***0.39+***	0.36	***0.42+***
wants to works less	0.50	0.51	0.49	0.51
can maintain work-life balance	0.59	0.59	0.60	0.55

continued

Table 1. Variable means

	(1) Total	(2) Women	(3) LGB+	(4) Ethnic
Supportive workplace measures				
supportive colleagues	0.77	0.77	0.80	***0.64-***
cooperative workplace	0.39	0.39	**0.43+**	***0.33-***
supervisor responsive	0.61	0.62	0.60	0.57
mentor	0.46	***0.48+***	0.46	***0.38-***
part of the family	0.44	***0.45+***	0.44	***0.35-***
influence	0.52	**0.53+**	0.51	0.50
quality of care	0.64	***0.65+***	0.67	0.65
network exists	0.48	***0.46-***	***0.74+***	***0.70+***
measures to prevent bullying				
not effective	0.22	0.22	0.20	***0.30+***
slightly effective	0.17	0.18	0.17	***0.22+***
moderately effective	0.37	0.37	0.39	***0.30-***
very or extremely effective	0.21	0.21	0.23	***0.17-***
bullying index	1.14	**1.12-**	1.17	***1.60+***
NHS England Region				
North of England	0.23	0.23	0.23	***0.10-***
Midlands and East of England	0.32	***0.34+***	***0.25-***	***0.19-***
London	0.15	***0.15-***	***0.24+***	***0.46+***
South West	0.11	0.11	***0.06-***	***0.05-***
South East	0.16	**0.15-**	**0.19+**	0.17
Trust type				
Acute Specialist	0.02	***0.01-***	***0.05+***	0.01
Acute Trusts	0.50	0.50	***0.38-***	0.48
Ambulance	0.01	***0.009-***	***0.04+***	0.007
Combined Acute and Community	0.12	0.12	0.11	***0.08-***
Combined Mental Health/Learning Dis	0.08	**0.09+**	***0.11+***	***0.16+***
Community	0.10	***0.11+***	0.08	***0.05-***
Mental Health/Learning Disability	0.14	0.13	***0.20***	***0.18+***
Observations	3497	2752	435	403

Table 2. Probability of job satisfaction excluding supportive workplace measures, selected results, marginal effects

	(1) Total	(2) Women	(3) LGB+	(4) Ethnic
Demographics				
woman	0.064***		0.043	0.058
LGB+	0.036	0.033		0.201***
disclose			0.0625	
ethnic	-0.022	-0.025	0.130**	
age	-0.001***	-0.002***	-0.004**	-0.002
dependent children	0.004	0.007	0.108*	0.028
living together	0.045***	0.035**	0.002	0.090**
disability	0.009	0.002	0.031	0.047
Job characteristics				
own-wage	0.166***	0.208***	0.077	0.169
relative-wage	-0.056	-0.101**	0.024	-0.032
job permanent	-0.027	-0.025	0.007	-0.183**
training	0.308***	0.307***	0.252***	0.284***
promotion	0.087***	0.088***	0.125**	0.106***
trade union	-0.023	-0.037**	-0.032	0.034
part-time	-0.007	-0.005	-0.092	-0.086*
works less	-0.045***	-0.039**	-0.113**	-0.031
work-life balance	0.205***	0.215***	0.183***	0.245***
pressure	-0.038**	-0.043*	-0.028	-0.045
overwhelmed	-0.157***	-0.135***	-0.177***	-0.162***
Qualifications	√	√	√	√
Region	√	√	√	√
Trust type	√	√	√	√
Observations	3497	2752	435	403
Pseudo R-squared	0.2726	0.2736	0.2667	0.3321

Standard errors in parentheses
*p<0.10, **p<0.05, ***p<0.001

Table 3. Probability of job satisfaction, full model, selected results, marginal effects

	(1) Total	(2) Women	(3) LGB+	(4) Ethnic
Demographics				
woman	0.046***		0.019	0.047
LGB+	0.025	0.019		0.201***
disclose			-0.023	
ethnic	0.009	0.005	0.146***	
age	-0.001**	-0.001**	-0.005***	-0.002
dependent children	-0.008	-0.005	0.074	0.014
living together	0.025**	0.023*	0.002	0.026
disability	0.015	0.011	0.010	0.025
Job characteristics				
own-wage	0.058*	0.065*	0.060	0.0516
relative-wage	-0.024	-0.044	-0.031	0.057
job permanent	-0.038*	-0.045**	0.031	-0.148**
training	0.078***	0.076***	0.073*	0.062*
promotion	0.039**	0.041**	0.067*	0.046
trade union	-0.0007	-0.009	-0.009	0.026
part-time	0.001	0.0005	-0.065	-0.047
works less	-0.022*	-0.018	-0.089**	0.008
work-life balance	0.093***	0.097***	0.075**	0.110***
pressure	0.004	-0.009	0.048	0.014
overwhelmed	-0.088***	-0.068***	-0.114***	-0.101**
Supportive workplace measures				
supportive colleagues	0.042**	0.049***	0.035	0.023
cooperative workplace	0.093***	0.094***	0.122***	0.019
supervisor responsive	0.056***	0.045**	0.102**	0.075
mentor	0.074***	0.073***	0.101***	0.102***
part of the family	0.104***	0.094***	0.0985**	0.152**
influence	0.150***	0.151***	0.093**	0.181***
quality of care	0.093***	0.096***	0.075*	0.048
network exists	0.065***	0.062***	0.133***	0.008
Measures to prevent bullying (omitted category: not effective)				
slightly effective	0.032*	0.036*	-0.064	0.054
moderately effective	0.047***	0.051***	-0.011	0.054
very or extremely	0.076***	0.085***	0.109*	0.122*
bullying index	-0.020***	-0.022***	-0.031*	-0.026
Qualifications	√	√	√	√
Region	√	√	√	√
Trust type	√	√	√	√
Observations	3497	2752	435	403
Pseudo R-squared	0.4407	0.4412	0.4383	0.4879

Standard errors in parentheses
*p<0.10, **p<0.05, ***p<0

Appendix

Table A1. Variable definitions

Variable	Definition
Job satisfaction	
Overall, how satisfied are you with your job these days?	
	extremely dissatisfied
	somewhat dissatisfied
	neither satisfied nor dissatisfied
	somewhat satisfied
	extremely satisfied
job satisfaction	Somewhat satisfied or extremely satisfied with job.
Demographics	
gender	Response to 'What best describes your gender?'
LGB+	Respondent chose Lesbian, Gay, Bisexual or Plus to 'Which of the following best describes how you think of yourself?'.
disclose	Open about sexuality at workplace.
ethnic	Member of an ethnic minority. Selected Asian; Black; Arab; Other Non White; or Mixed to 'What is your ethnic background?'
age	Age of respondent.
dependent children	Has dependent children (aged 0-18).
living together	Living with partner.
disability	Long-standing illness, health problem or disability.
Qualifications, highest acquired	
minimal	O level or GCSE grades D-G, or below, or not recognised. These national exams are typically taken at age 16.
GCSE grades A-C	O level or GCSE grades A-C. These national exams are typically taken at age 16.
Trade	Trade apprenticeships.
A levels	A levels. These national exams are taken at the completion of secondary school, typically at age 18.
Diploma	Diploma in higher education, teaching qual. and others.
First degree	First degree and PGCE.
Higher degree	Higher degree or postgraduate.
Job characteristics	
own-wage	Average hourly pay in GBP, full time equivalent, (log).
relative-wage	Relative wage by occupation and age category (log).
part-time	Part-time work.
job permanent	Has a permanent contract.
training	Somewhat or extremely satisfied with training received.

continued

Table A1. Variable definitions

Variable	Definition
Job characteristics cont.	
promotion	Has been promoted.
trade union	Member of a trade union.
pressure	Job makes feel pressure always or often.
overwhelmed	Job makes feel overwhelmed always or often.
work less	Prefers to work less hours.
work-life balance	Somewhat or strongly agrees maintain work-life balance.
Supportive workplace measures	
supportive colleagues	Somewhat or strongly agrees the people worked with are supportive.
cooperative workplace	Feels workplace is cooperative.
supervisor responsive	Somewhat or strongly agrees supervisor responds to their suggestions.
mentor	Has mentor/coach for work advice.
part of the family	Somewhat or strongly agrees they feel part of the family at this organisation.
influence	Somewhat or extremely satisfied with the amount of influence has over job.
quality of care	Somewhat or strongly agrees is satisfied with the quality of care given to patients/service-users.
network exists	Staff network exists in their Trust
bullying index	Bullied, witnessed bullying and/or subject to discrimination.
Measures to prevent bullying	
Do you think the measure your organisation takes to prevent bullying/discrimination are effective?	
	not effective
	slightly effective
	moderately effective
	very or extremely effective
Workplace controls	
NHS England region	
North of England	
Midlands and East of England	
London	
South West	
South East	
Trust type	
Acute Specialist Trusts	
Acute Trusts	
Ambulance Trusts	
Combined Acute and Community Trusts	
Combined. Mental Health / Learning Disability	
Community Trusts	
Mental Health / Learning Disability Trusts	

Table A2. Means of variables by gender, nursing occupation, sexual orientation and ethnicity

		Gender		Nurse		Sexual orientation		Ethnicity	
	Total	Men	Women	Yes	No	LGB+	Hetero	Ethnic	Non-Ethnic
Job satisfaction									
extremely dissatisfied	0.08	***0.11***	***0.08***	0.08	0.09	***0.05***	***0.09***	***0.13***	***0.08***
somewhat dissatisfied	0.22	0.21	0.22	0.21	0.22	0.20	0.22	0.19	0.22
neither satisfied nor dissatisfied	0.15	**0.17**	**0.14**	0.15	0.15	0.15	0.15	**0.18**	**0.15**
somewhat satisfied	0.41	**0.38**	**0.42**	0.42	0.41	0.44	0.41	0.37	0.42
extremely satisfied	0.11	0.10	0.11	0.11	0.11	0.12	0.11	0.10	0.11
job satisfaction binary	0.53	***0.49***	***0.54***	0.54	0.53	**0.58**	**0.53**	***0.49***	***0.54***
Demographics									
woman	0.79			***0.88***	***0.76***	***0.52***	***0.82***	***0.72***	***0.80***
LGB+	0.12	***0.28***	***0.08***	**0.11**	**0.13**			0.10	0.13
disclose	0.06	***0.17***	***0.03***	0.06	0.07			**0.04**	**0.07**
ethnic	0.12	***0.15***	***0.11***	0.13	0.11	0.09	0.12		
age	46.14 (11.44)	***44.98*** (11.78)	***46.45*** (11.33)	***47.35*** (10.45)	***45.73*** (11.73)	***41.32*** (11.24)	***46.82*** (11.31)	**45.09** (10.87)	**46.27** (11.51)
dependent children	0.32	0.30	0.32	***0.38***	***0.30***	***0.15***	***0.34***	***0.41***	***0.31***
living together	0.69	0.69	0.68	0.70	0.68	***0.57***	***0.70***	***0.62***	***0.70***
disability	0.36	0.37	0.35	***0.33***	***0.36***	***0.45***	***0.34***	***0.29***	***0.36***

continued

Table A2. Means of variables by gender, nursing occupation, sexual orientation and ethnicity

		Gender		Nurse		Sexual orientation		Ethnicity	
	Total	Men	Women	Yes	No	LGB+	Hetero	Ethnic	Non-Ethnic
Qualifications, highest acquired									
minimal	0.05	0.05	0.05	***0.007***	***0.07***	***0.02***	***0.06***	***0.02***	***0.06***
GCSE grades A-C	0.08	**0.07**	**0.09**	***0.01***	***0.11***	***0.05***	***0.09***	***0.05***	***0.09***
Trade	0.004	***0.01***	***0.002***		0.006	0.002	0.004	0.002	0.004
A levels	0.09	0.10	0.09	***0.009***	***0.12***	0.10	0.09	***0.05***	***0.10***
Diploma	0.16	0.17	0.16	***0.27***	***0.12***	0.16	0.16	***0.12***	***0.17***
First degree	0.30	0.31	0.30	***0.38***	***0.27***	0.31	0.30	**0.34**	**0.29**
Higher degree	0.28	0.27	0.28	***0.30***	***0.27***	0.32	0.27	***0.37***	***0.26***
Job characteristics									
own-wage	2.74 (0.36)	***2.78*** (0.38)	***2.73*** (0.35)	***2.89*** (0.22)	***2.69*** (0.38)	2.75 (0.36)	2.74 (0.36)	***2.81*** (0.35)	***2.73*** (0.36)
relative-wage	0.001 (0.29)	***0.04*** (0.32)	***-0.01*** (0.28)	0.00 (0.21)	0.00 (0.31)	**0.02** (0.29)	**-0.002** (0.29)	***0.05*** (0.29)	***-0.01*** (0.29)
part-time	0.24	***0.08***	***0.28***	0.24	0.23	***0.11***	***0.25***	***0.15***	***0.25***
job permanent	0.93	0.92	0.93	***0.95***	***0.92***	0.93	0.93	0.91	0.93
training	0.52	***0.46***	***0.54***	***0.61***	***0.49***	0.53	0.52	0.50	0.52
promotions	0.36	0.36	0.36	***0.42***	***0.34***	0.42	0.36	0.38	0.36
trade union	0.57	***0.53***	***0.58***	***0.90***	***0.46***	0.57	0.57	0.58	0.57
pressure	0.56	0.55	0.56	***0.63***	***0.53***	0.54	0.56	0.59	0.55
overwhelmed	0.38	***0.33***	***0.39***	***0.44***	***0.36***	0.36	0.38	***0.42***	***0.37***

continued

Table A2. Means of variables by gender, nursing occupation, sexual orientation and ethnicity

		Gender		Nurse		Sexual orientation		Ethnicity	
	Total	Men	Women	Yes	No	LGB+	Hetero	Ethnic	Non-Ethnic
Job characteristics cont.									
works less	0.50	0.48	0.51	0.52	0.50	0.49	0.50	0.51	0.50
work-life balance	0.59	0.59	0.59	***0.53***	***0.61***	0.60	0.58	0.55	0.59
Supportive workplace measures									
supportive colleagues	0.77	0.75	0.77	0.79	0.76	0.80	0.76	***0.64***	***0.78***
cooperative workplace	0.39	0.41	0.39	0.40	0.39	**0.43**	**0.39**	***0.33***	***0.40***
supervisor responsive	0.61	0.59	0.62	***0.58***	***0.62***	0.60	0.61	0.57	0.62
mentor	0.46	0.42	0.48	***0.51***	***0.45***	0.46	0.47	***0.38***	***0.47***
part of the family	0.44	0.40	0.45	0.44	0.44	0.44	0.44	***0.35***	***0.45***
influence	0.52	0.49	0.53	0.52	0.52	0.51	0.52	0.50	0.53
quality of care	0.64	0.60	0.65	***0.68***	***0.63***	0.67	0.64	0.65	0.64
bullying index	1.14	1.20	1.12	***1.22***	***1.11***	1.17	1.13	***1.60***	***1.08***
measures to prevent bullying									
not effective	0.22	0.24	0.22	0.23	0.22	0.20	0.23	***0.30***	***0.21***
slightly effective	0.17	0.17	0.18	***0.20***	***0.17***	0.17	0.18	***0.22***	***0.17***
moderately effective	0.37	0.35	0.37	0.35	0.38	0.39	0.37	***0.30***	***0.38***
very or extremely effective	0.21	0.22	0.21	0.20	0.22	0.23	0.21	***0.17***	***0.22***
network exists	0.48	***0.54***	***0.46***	**0.50**	**0.47**	***0.74***	***0.44***	***0.70***	***0.45***
LGB+ network membership						***0.36***	***0.04***		

continued

Table A2. Means of variables by gender, nursing occupation, sexual orientation and ethnicity

		Gender		Nurse		Sexual orientation		Ethnicity	
	Total	Men	Women	Yes	No	LGB+	Hetero	Ethnic	Non-Ethnic
Supportive workplace measures cont.									
women's network membership		0.01	0.03						
ethnic network membership								***0.34***	***0.03***
Workplace controls									
NHS England region									
North of England	0.23	0.23	0.23	0.21	0.24	0.23	0.23	***0.10***	***0.25***
Midlands and East of England	0.32	***0.28***	***0.34***	**0.35**	**0.32**	***0.25***	***0.33***	***0.19***	***0.34***
London	0.15	***0.18***	***0.15***	0.16	0.15	***0.24***	***0.14***	***0.46***	***0.11***
South West	0.11	0.11	0.11	0.10	0.12	***0.06***	***0.12***	***0.05***	***0.12***
South East	0.16	**0.18**	**0.15**	0.15	0.16	**0.19**	**0.15**	0.17	0.15
Trust type									
Acute Specialist	0.02	***0.04***	***0.01***	0.02	0.02	***0.05***	***0.01***	0.01	0.02
Acute Trusts	0.50	0.51	0.50	**0.53**	**0.49**	***0.38***	***0.52***	0.48	0.50
Ambulance	0.01	***0.03***	***0.009***	***0.003***	***0.01***	***0.04***	***0.008***	0.007	0.01
Combined Acute and Community	0.12	0.12	0.12	0.12	0.12	0.11	0.12	***0.08***	***0.13***
Combined Mental Health/Learning Disability	0.08	**0.07**	**0.09**	0.07	0.09	***0.11***	***0.08***	***0.16***	***0.07***
Community	0.10	***0.06***	***0.11***	0.09	0.10	0.08	0.10	***0.05***	***0.10***
Mental Health/Learning Disability	0.14	0.15	0.13	0.13	0.14	***0.20***	***0.13***	***0.18***	***0.13***
Observations	3497	745	2752	876	2621	435	3062	403	3094

Mean pair difference (Men Vs. Women, Nurses Vs. Non-nurses, LGB+ Vs. Heterosexual, Ethnic Vs. Non-Ethnic): bold and italic p<0.05, bold p<0.10.

Table A3. Probability of job satisfaction, full model, marginal effects

	(1) Total	(2) Nurses	(3) LGB+	(4) Women	(4) Ethnic
Demographics					
woman	0.0456*** (0.0141)	-0.0618 (0.0430)	0.0189 (0.0351)		0.0465 (0.0341)
LGB+	0.0252 (0.0196)	-0.0236 (0.0359)		0.0187 (0.0220)	0.201*** (0.0500)
disclose			0.0230 (0.0391)		
ethnic	0.00880 (0.0184)	-0.0441 (0.0383)	0.146*** (0.0495)	0.00519 (0.0225)	
age	-0.00121** (0.000562)	-0.00256 (0.00270)	-0.0050*** (0.0019)	-0.00138** (0.000652)	-0.00244 (0.00175)
dependent children	-0.00840 (0.0179)	-0.0204 (0.0273)	0.0739 (0.0496)	-0.00479 (0.0208)	0.0138 (0.0401)
living together	0.0254** (0.0112)	0.00314 (0.0237)	0.0021 (0.0435)	0.0225* (0.0120)	0.0263 (0.0321)
disability	0.0146 (0.0116)	0.0132 (0.0214)	0.0094 (0.0439)	0.0111 (0.0123)	0.0248 (0.0346)
Qualifications (omitted category: minimal)					
GCSE grades A-C[4]	-0.0403 (0.0327)	-0.0102 (0.132)	-0.0746 (0.0961)	-0.0500 (0.0378)	-0.0165 (0.126)
Trade	0.0566 (0.0543)			0.0217 (0.100)	
A levels	-0.0104 (0.0283)	-0.00594 (0.176)	-0.0963 (0.0980)	-0.0174 (0.0300)	-0.0651 (0.100)
Diploma	0.00350 (0.0291)	-0.0206 (0.108)	-0.0635 (0.0929)	-0.0123 (0.0342)	-0.0936 (0.0849)
First degree	0.0375 (0.0346)	0.0124 (0.108)	-0.0748 (0.0882)	0.0351 (0.0378)	-0.00432 (0.0830)
Higher degree	0.0337 (0.0295)	0.0507 (0.109)	-0.112 (0.0928)	0.0309 (0.0320)	-0.0296 (0.0842)
Job characteristics					
own-wage	0.0581* (0.0314)	-0.106 (0.524)	0.0601 (0.0906)	0.0645* (0.0358)	0.0506 (0.0986)
relative-wage	-0.0244 (0.0355)	0.0983 (0.506)	-0.0309 (0.0946)	-0.0437 (0.0394)	0.0566 (0.0956)
part-time	0.00103 (0.0164)	-0.0216 (0.0367)	-0.0652 (0.0617)	0.000532 (0.0181)	-0.0462 (0.0441)
job permanent	-0.0383* (0.0199)	-0.0283 (0.0513)	0.0314 (0.0807)	-0.0448** (0.0209)	-0.148** (0.0690)
training	0.0778*** (0.0164)	0.0714*** (0.0268)	0.0730* (0.0400)	0.0763*** (0.0187)	0.0618* (0.0332)
promotions	0.0394** (0.0165)	0.0790*** (0.0249)	0.0666 (0.0406)	0.0407** (0.0168)	0.0461 (0.0292)
trade union	-0.000726 (0.0134)	-0.0520 (0.0457)	-0.00920 (0.0400)	-0.00914 (0.0151)	0.0260 (0.0336)

continued

Table A3. Probability of job satisfaction, full model, marginal effects

	(1) Total	(2) Nurses	(3) LGB+	(4) Women	(4) Ethnic
Job characteristics cont.					
pressure	0.00443 (0.0161)	-0.0298 (0.0270)	0.0476 (0.0373)	-0.00900 (0.0223)	0.0141 (0.0390)
overwhelmed	-0.0881*** (0.0134)	-0.0424 (0.0271)	-0.114*** (0.0418)	-0.0680*** (0.0140)	-0.101** (0.0493)
works less	-0.0223* (0.0123)	-0.00462 (0.0253)	-0.0892** (0.0368)	-0.0175 (0.0115)	0.00755 (0.0313)
work-life balance	0.0929*** (0.0120)	0.117*** (0.0242)	0.0751** (0.0351)	0.0970*** (0.0148)	0.110*** (0.0402)
Supportive workplace measures					
supportive colleagues	(0.0168) 0.0930***	(0.0264) 0.0980**	(0.0455) 0.122***	(0.0179) 0.0941***	(0.0499) -0.0185
cooperative workplace	(0.0145) 0.0561***	(0.0407) 0.0610*	(0.0463) 0.102**	(0.0145) 0.0452**	(0.0376) 0.0749
supervisor responsive	(0.0162) 0.0742***	(0.0352) 0.0859***	(0.0425) 0.101***	(0.0197) 0.0726***	(0.0509) 0.102***
mentor	(0.0125) 0.104***	(0.0245) 0.110***	(0.0321) 0.0985**	(0.0149) 0.0941***	(0.0325) 0.152**
part of the family	(0.0142) 0.150***	(0.0258) 0.115***	(0.0442) 0.0927**	(0.0160) 0.151***	(0.0597) 0.181***
influence	(0.0173) 0.0926***	(0.0310) 0.0816***	(0.0437) 0.0749*	(0.0167) 0.0959***	(0.0564) 0.0477
quality of care	(0.0200) -0.0202***	(0.0275) -0.0173	(0.0444) -0.0311*	(0.0220) -0.0224***	(0.0381) -0.0261
bullying index	(0.00666) 0.056***	(0.0126) 0.045**	(0.0187) 0.102**	(0.00779) 0.075	(0.0214) 0.014
measures to prevent bullying (omitted category: not effective)					
slightly effective	0.0317* (0.0172)	0.0559* (0.0335)	-0.0638 (0.0552)	0.0361* (0.0192)	0.0538 (0.0499)
moderately effective	0.0471*** (0.0158)	0.0735** (0.0363)	-0.0105 (0.0401)	0.0507*** (0.0193)	0.0541 (0.0458)
very or extremely effective	0.0758*** (0.0203)	0.135*** (0.0447)	0.109* (0.0597)	0.0849*** (0.0231)	0.122* (0.0626)
network exists	0.0646*** (0.0129)	0.0768*** (0.0215)	0.133*** (0.0442)	0.0620*** (0.0151)	-0.00771 (0.0408)
NHS England region (omitted category: London)					
North of England	0.0407* (0.0229)	0.0765** (0.0373)	0.0579 (0.0498)	0.0450* (0.0238)	0.125** (0.0613)
Midlands & East of England	0.0170 (0.0242)	0.0630* (0.0363)	-0.0099 (0.0472)	0.0303 (0.0258)	-0.0158 (0.0484)
South West	0.0511** (0.0227)	0.122*** (0.0432)	0.0985 (0.0639)	0.0576** (0.0238)	-0.00450 (0.0976)
South East	0.0325 (0.0265)	0.0421 (0.0332)	-0.0277 (0.0526)	0.0358 (0.0268)	-0.0408 (0.0463)

continued

Table A3. Probability of job satisfaction, full model, marginal effects

	(1) Total	(2) Nurses	(3) LGB+	(4) Women	(4) Ethnic
Trust type (omitted category: Acute)					
Acute Specialist²	-0.0278 (0.0359)	-0.143*** (0.0431)	-0.164*** (0.0590)	-0.0196 (0.0303)	0.0433 (0.137)
Ambulance	0.0223 (0.0647)	-0.0080 (0.107)	-0.147* (0.0847)	-0.0192 (0.0929)	
Acute & Community	-0.00345 (0.0207)	0.00133 (0.0358)	-0.176*** (0.0542)	-0.00248 (0.0198)	-0.154*** (0.0541)
Comb. Mental Health/Learning Disability	-0.0174 (0.0216)	0.0315 (0.0381)	-0.172*** (0.0566)	-0.0213 (0.0234)	-0.0555 (0.0659)
Community	-0.0502*** (0.0192)	-0.0396 (0.0386)	-0.0852 (0.0538)	-0.0482** (0.0196)	0.0247 (0.0662)
Mental Health / Learning Disability	-0.0429** (0.0206)	-0.0346 (0.0298)	-0.0521 (0.0422)	-0.0344* (0.0202)	-0.0591 (0.0440)
Observations	3497	876	435	2752	403
Pseudo R-squared	0.4407	0.4492	0.4383	0.4412	0.4879

Standard errors in parentheses *p<0.10, **p<0.05, ***p<0.01
For the ethnic sample, category 'Acute Specialist Trusts' also includes category 'Ambulance Trusts'; For LGB+ and ethnic samples, category 'GCSE grades A-C' also includes category 'trade'.

Table A4. Probability of job satisfaction, full model, selected results, marginal effects

	(1) Total	(2) 2-way interactions (woman, LGB+, ethnic)	(3) 3-way interactions (woman*LGB+*ethnic)
Demographics			
woman	0.046***	0.047***	0.048***
LGB+	0.025	0.026	0.028
ethnic	0.009	0.012	0.013
age	-0.001**	-0.001**	-0.001**
dependent children	-0.008	-0.007	-0.007
living together	0.025**	0.025**	0.025**
disability	0.015	0.015	0.015
Job characteristics			
own-wage	0.058*	0.057*	0.057*
relative-wage	-0.024	-0.024	-0.024
job permanent	-0.038*	-0.039*	-0.039**
training	0.078***	0.077***	0.077***
promotion	0.039**	0.040**	0.040**
trade union	-0.0007	-0.0009	-0.0009
part-time	0.001	0.00006	-0.00005
works less	-0.022*	-0.021*	-0.021*
work-life balance	0.092***	0.092***	0.092***
pressure	0.004	0.005	0.004
overwhelmed	-0.088***	-0.088***	-0.088***
Supportive workplace measures			
supportive colleagues	0.041**	0.041**	0.041**
cooperative workplace	0.093***	0.092***	0.092***
supervisor responsive	0.056***	0.056***	0.056***
mentor	0.074***	0.073***	0.073***
part of the family	0.104***	0.103***	0.103***
influence	0.150***	0.151***	0.151***
quality of care	0.092***	0.093***	0.093***
network exists	0.064***	0.065***	0.065***
Measures to prevent bullying (omitted category: not effective)			
slightly effective	0.031*	0.030*	0.030*
moderately effective	0.047***	0.046***	0.046***
very or extremely	0.075***	0.074***	0.074***
bullying index	-0.0202***	-0.020***	-0.020***
Qualifications	√	√	√
Region	√	√	√
Trust type	√	√	√

continued

Table A4. Probability of job satisfaction, full model, selected results, marginal effects

	(1) Total	(2) 2-way interactions (woman, LGB+, ethnic)	(3) 3-way interactions (woman*LGB+*ethnic)
2-way interactions between woman, LGB+, ethnic included		√	√
3-way interaction (woman, LGB+, ethnic) included			√
Observations	3497	3497	3497
Pseudo R-squared	0.4407	0.4419	0.4420

Additional notes for Table A4.

Table A4 presents the combined marginal effects associated with the variables of interest. To further explore the interaction **Female and LGB+** in the model presented in column 2 of Table A4, the marginal effects can be evaluated at particular combinations of these variables (e.g., 0 and 1). This approach allows us to assess, ceteris paribus, the marginal effect of being LGB+ on the probability of job satisfaction separately for males and females, as shown in the following table:

	Marginal effect	Std. Error	P-value
Male	0.042	0.029	0.147
Female	0.022	0.021	0.303

We observe:

- Among men, identifying as LGB+ is associated with a 4.2 percentage point higher probability of job satisfaction compared to non-LGB+ men, but this difference is not statistically significant ($p = 0.15$).
- Among women, identifying as LGB+ is associated with a 2.2 percentage point higher probability of job satisfaction compared to non-LGB+ women, also not statistically significant ($p = 0.25$).

As both effects are small and not statistically significant, we cannot conclude that LGB+ status affects job satisfaction differently for men vs. women.

For the other interactions in column 2 of Table A4, the results obtained are:

a. Female*ethnic

	Marginal effect	Std. Error	P-value
Not ethnic minority	0.049	0.015	0.001
Ethnic minority	0.031	0.034	0.365

b. Ethnic*LGB+

	Marginal effect	Std. Error	P-value
Not ethnic minority	0.009	0.020	0.630
Ethnic minority	0.153	0.048	0.002

Considering the interactions in the model presented in column (3) of Table A4:

a. Female*LGB+

	Marginal effect	Std. Error	P-value
Not ethnic minority	0.041	0.029	0.152
Ethnic minority	0.024	0.021	0.251

b. Female*ethnic

	Marginal effect	Std. Error	P-value
Not ethnic minority	0.050	0.015	0.001
Ethnic minority	0.032	0.034	0.346

c. Ethnic*LGB+

	Marginal effect	Std. Error	P-value
Not ethnic minority	0.008	0.020	0.671
Ethnic minority	0.173	0.055	0.002

d. Similarly, with the triple interaction **Female*LGB+*ethnic**, there is no single "marginal effect" of the triple interaction. It depends on all three variables and where they are evaluated. For example:

- Average marginal effect of being LGB+ (compared to non-LGB+) on the predicted probability of job satisfaction, separately for each combination of sex (female = 0/1) and ethnic minority status (ethnic minority = 0/1).

	Marginal effect	Std. Error	P-value
Male, not ethnic minority	0.029	0.030	0.338
Male, ethnic minority	0.132	0.085	0.121
Female, not ethnic minority	0.003	0.023	0.893
Female, ethnic minority	0.185	0.069	0.008

- Average marginal effect of being female (vs. male) on the predicted probability of job satisfaction, separately for each combination of ethnic minority status (0/1), and LGB+ status (0/1).

	Marginal effect	Std. Error	P-value
Non-LGB+, not ethnic minority	0.053	0.016	0.001
Non-LGB+, ethnic minority	0.026	0.038	0.501
LGB+, not ethnic minority	0.027	0.034	0.421
LGB+, ethnic minority	0.078	0.104	0.452

- Average marginal effect of being Ethnic minority (vs. non-minority) on the predicted probability of job satisfaction, separately for each combination of gender (female = 0/1), and LGB+ status (LGB+ = 0/1).

	Marginal effect	Std. Error	P-value
Male, non-LGB+	0.016	0.037	0.662
Male, LGB+	0.119	0.081	0.142
Female, non-LGB+	−0.011	0.021	0.610
Female, LGB+	0.170	0.076	0.026

Table A5. Reporting coefficients of interactions, full model, not marginal effects

	(1) Total	(2) 2-way interactions (woman, LGB+, ethnic)	(3) 3-way interactions (woman*LGB+*ethnic)
Interactions			
Female*LGB+		-0.090	-0.121
(Robust Standard error)		(0.1494)	(0.1626)
[p-value]		[0.545]	[0.456]
Female*ethnic		-0.084	-0.127
(Robust Standard error)		(0.1827)	(0.2003)
[p-value]		[0.643]	[0.525]
Ethnic*LGB+		0.686**	0.488
(Robust Standard error)		(0.2554)	(.4368)
[p-value]		0.007	[0.263]
Female*LGB+*ethnic			0.394
(Robust Standard error)			(0.6255)
[p-value]			[0.528]

Notes to Authors

The Australian Journal of Labour Economics (AJLE) is a forum for the analysis of labour economics and labour relations. It is particularly focused on theoretical and policy developments in respect of Australian labour markets. Interdisciplinary approaches are particularly encouraged. The AJLE invites submissions of articles in the following areas:

- Time allocation, work behaviour, and employment determination
- Wages, compensation, and labour costs
- Labour-management relations, trade unions, and collective bargaining
- Work organisation and the sociology of work
- Productivity
- Income and wealth distribution
- Mobility, unemployment, labour force participation and vacancies
- Gender, ethnicity, labour market segmentation and discrimination
- Population and demography in respect of the labour market

While contributors to the AJLE are expected to demonstrate theoretical or empirical originality – and preferably both – they should make their work accessible to readers from a non-technical background. Survey articles are also encouraged. Further, as a means of strengthening the integration of theory and practice the AJLE welcomes reflective contributions from practitioners. The AJLE recognises that the areas of labour economics and labour relations are subject to controversy and aims to provide an arena for debate.

Submission of Papers

Contributors should submit manuscripts online at https://ajle.org or via email to:

ajle@curtin.edu.au
Sandie Rawnsley
Australian Journal of Labour Economics
C/o Bankwest Curtin Economics Centre
Faculty of Business and Law
Curtin University
GPO Box U1987
Perth WA 6845

Manuscripts should not normally exceed 12,000 words and should contain an abstract of approximately 200 words, along with at least three EconLit subject descriptors. They should be double-spaced and should include a separate title sheet which contains the author's name, affiliation and contact details.

References in Harvard style. A detailed style guide is accessible through the Centre's website at https://ajle.org/index.php/ajle_home/about/submissions#authorGuidelines

Refereeing Procedure

It is the policy of the editors to send submitted papers to two referees. The names of authors are not disclosed to referees.